

Dual equivalence graphs revisited and the explicit Schur expansion of a family of LLT polynomials

Austin Roberts

Received: 1 February 2013 / Accepted: 8 May 2013 / Published online: 29 May 2013
© Springer Science+Business Media New York 2013

Abstract 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 so that the sum of the weights of a connected component is a single Schur function. In this paper, we improve on Assaf’s axiomatization of such 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.

Keywords Dual equivalence graph · LLT polynomials · Macdonald polynomials · Schur expansion · Quasisymmetric functions

Electronic supplementary material The online version of this article (doi:[10.1007/s10801-013-0452-y](https://doi.org/10.1007/s10801-013-0452-y)) contains supplementary material, which is available to authorized users.

A. Roberts (✉)
Department of Mathematics, University of Washington, Seattle, WA 98195-4350, USA
e-mail: austinis@math.washington.edu

1 Introduction

Dual equivalence was developed and applied by Mark Haiman in [16] as an extension of work done by Donald Knuth in [17]. Sami Assaf then introduced the theory of dual equivalence graphs in her Ph.D. dissertation [1] and subsequent preprint [3]. 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 [4] and for the product of a Schubert polynomial with a Schur polynomial in [5].

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 [9] 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 [20], for Lascoux–Leclerc–Thibon (LLT) polynomials in [14], for Macdonald polynomials in [13], and conjecturally for the composition of the nabla operator with an elementary symmetric function in [14]. 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 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 2.6, are necessarily Schur positive. In Theorems 3.13 and 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 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 4.3. LLT polynomials were first introduced in [19] as a q -analogue to products of Schur functions and were later given a description in terms of tuples of skew tableaux in [14]. Corollary 4.4 then applies the results of [13] to give an explicit combinatorial description for a family of modified Macdonald polynomials. First introduced in [21], 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 [22]. 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 [22] for details). In [15], 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 [6] and [28], and modified Macdonald polynomials indexed by shapes with strictly less than three columns was described in [13] (which in turn drew

on the earlier work in [6] and [28]). The first combinatorial description of the two column case was given in [7], but others were subsequently given in [18, 29], and [2]. 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 [30]. Finding a combinatorial interpretation for the three column case is still an open problem, though there is a conjectured formula in [11].

This paper is broken into sections as follows. Section 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 3 is dedicated to further developing the theory of dual equivalence graphs, culminating in a new axiomatization for dual equivalence graphs in Theorem 3.17. Section 4 applies the results of Sect. 3 to LLT polynomials and Macdonald polynomials. The graph structure given to LLT polynomials in [3] is reviewed before Theorem 4.11 classifies the set of LLT polynomials that correspond to dual equivalence graphs. Theorem 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 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 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.

2.1 Tableaux

A *partition* λ is a weakly decreasing finite sequence of nonnegative integers $\lambda_1 \geq \dots \geq \lambda_k \geq 0$. 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 Fig. 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 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 right skew diagram of Fig. 1. Here, λ/ρ is referred to as the



Fig. 1 The diagrams for $(4, 3, 2, 2)$ and $(4, 3, 1, 1)/(2, 2)$



Fig. 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

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 $[n] = \{1, \dots, n\}$ 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 Fig. 2). The set of all standard Young tableaux of shape λ is denoted $\text{SYT}(\lambda)$. The union of all $\text{SYT}(\lambda)$ over $\lambda \vdash n$ is denoted $\text{SYT}(n)$. Similarly, the set of all skew tableaux of shape λ/ρ is denoted $\text{SYT}(\lambda/\rho)$. 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 Fig. 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 [8, Part I], [24, Ch. 3], or [25, Ch. 7]. Given a (possibly skew) tableau T , define $\text{sh}(T)$ 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 Fig. 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.

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 Fig. 2 are $+ - - + - + - +$ and $- - +$, respectively. Given a fixed tableau, the row reading

word and the content reading word always have the same signature (see [25, Ch. 7] for details).

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 [8, Ch. 2–4], [24, Ch. 3], or [25, 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, $P: S_n \rightarrow \text{SYT}(n)$ and $Q: S_n \rightarrow \text{SYT}(n)$ will be the functions taking a permutation to its insertion tableau and recording tableau, respectively. These two functions are related by

$$Q(w) = P(w^{-1}). \tag{2.1}$$

A detailed proof of this fact can be found in [8, Ch. 4.1]. We then write $\text{sh}(w)$ to mean $\text{sh}(P(w))$.

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

$$\begin{aligned} K_j(\dots xyz \dots) &= (\dots xyz \dots), & K_j(\dots zyx \dots) &= (\dots zyx \dots), \\ K_j(\dots yxz \dots) &= (\dots yxz \dots), & K_j(\dots xzy \dots) &= (\dots xzy \dots), \end{aligned} \tag{2.2}$$

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 $\text{sh}(v) = \text{sh}(w)$. 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 [8, Ch. 2.1] for the details of the proof).

Next, we comment on the relationship between $\text{sh}(w)$ and subwords of w , as is well presented in [8, Ch. 3]. If $\text{sh}(w) = \lambda$, then the longest increasing subword of w has length λ_1 , and the longest decreasing subword of w has length λ_1 . For instance, if $w = 15342$, then $\text{sh}(w) = (3, 1, 1)$, 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 $S = \{a, a + 1, \dots, b\}$, 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 [8, 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 $\text{SYam}(\lambda)$ 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 $\text{Yam}(\lambda)$ 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 $\text{Yam}(\lambda)$ by replacing all of the 1’s with $1, \dots, \lambda_1$ in increasing reading order, all of the 2’s with $\lambda_1 + 1, \dots, \lambda_2$ 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 $\text{Yam}(\lambda)$ is precisely $\text{SYam}(\lambda)$.

2.3 Symmetric functions

The ring of symmetric functions has several well-known bases with ties to tableaux, as is well laid out in [25, Ch. 7], [8, Part I], or [24, 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 2.1 Given any signature $\sigma \in \{\pm 1\}^{n-1}$, define the *fundamental quasisymmetric function* $F_\sigma(X) \in \mathbb{Z}[x_1, x_2, \dots]$ by

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

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 ([9]) Given any skew shape λ/ρ , define

$$s_{\lambda/\rho}(X) := \sum_{T \in \text{SYT}(\lambda/\rho)} F_{\sigma(T)}(X), \tag{2.3}$$

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



Fig. 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 $---+++-$

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 [19] and [21], respectively) may also be expressed using the sum of fundamental quasisymmetric polynomials. We now present these combinatorial definitions, as they will be needed in Sect. 4. The LLT polynomials, denoted $G_{\mathbf{v}}(X; q)$, were originally described in terms of ribbon tableaux in [19]. We will instead use the equivalent definition given in [14, Cor. 5.2.4], which defines $G_{\mathbf{v}}(X; q)$ by using a k -tuple of skew shapes \mathbf{v} .

Given a k -tuple of skew shapes $\mathbf{v} = (v^{(0)}, \dots, v^{(k-1)})$, we write $|\mathbf{v}| = n$ if $\sum_{i=0}^{k-1} |v^{(i)}| = n$. A standard filling $\mathbf{T} = (T^{(0)}, \dots, T^{(k-1)})$ of \mathbf{v} is a bijective filling of the diagram of \mathbf{v} 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 \mathbf{v} as $\text{SYT}(\mathbf{v})$. Define the *shifted content* of a cell x in $v^{(i)}$ as

$$\tilde{c}(x) = k \cdot c(x) + i, \tag{2.4}$$

where $c(x)$ is the content of x in $v^{(i)}$. 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 Fig. 3.

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

$$\text{Inv}_k(\mathbf{T}) := \{(x, y) \mid k > \tilde{c}(y) - \tilde{c}(x) > 0 \text{ and } \mathbf{T}(x) > \mathbf{T}(y)\}. \tag{2.5}$$

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

$$\text{inv}_k(\mathbf{T}) := |\text{Inv}_k(\mathbf{T})|. \tag{2.6}$$

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

$$\text{inv}_{\mathbf{v}}(w) := \text{inv}_k(\mathbf{T}). \tag{2.7}$$

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

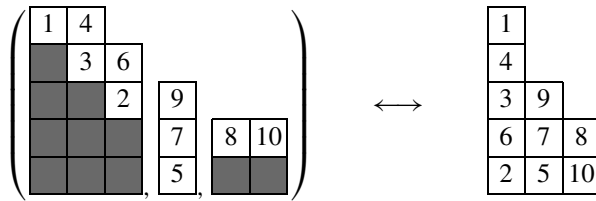


Fig. 4 An example of the bijection between standard fillings of shapes in $\text{TR}(\mu/\rho)$ and bijective fillings of μ/ρ

Now define the set of LLT polynomials by

$$G_{\mathbf{v}}(X; q) := \sum_{\mathbf{T} \in \text{SYT}(\mathbf{v})} q^{\text{inv}_k(\mathbf{T})} F_{\sigma(\mathbf{T})}(X). \tag{2.8}$$

Though LLT polynomials are known to be symmetric, with proofs in [19, Thm. 6.1] and [13, Thm. 3.3], it is still challenging to expand them in terms of Schur functions. A partial solution to this problem is given in Sect. 4.

We now move on to the definition of the modified Macdonald polynomials $\tilde{H}_{\mu/\rho}(X; q, t)$. We will use [13, Thm. 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 $\text{TR}(\mu/\rho)$ be the set of tuples of ribbons $\mathbf{v} = (v^{(0)}, \dots, v^{(k-1)})$ such that $v^{(i)}$ 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 $\text{TR}(\mu/\rho)$ and bijective fillings of μ/ρ given by turning each ribbon into a column of μ/ρ as demonstrated in Fig. 4.

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

$$a(x) := |\{v^{(j)}: j > i \text{ and there exists some } y \in v^{(j)} \text{ such that } c(x) = c(y)\}|, \tag{2.9}$$

$$l(x) := |\{y: y \in v^{(i)} \text{ and } c(y) < c(x)\}|. \tag{2.10}$$

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 Fig. 4, then $a(x) = 2$ and $l(x) = 3$.

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

$$\text{Des}(\mathbf{v}) := \{x \in \mathbf{v}: x \text{ is a descent of } \mathbf{v}\}. \tag{2.11}$$

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

$$a(\mathbf{v}) := \sum_{x \in \text{Des}(\mathbf{v})} a(x), \tag{2.12}$$

$$\text{inv}(\mathbf{T}) := \text{inv}_k(\mathbf{T}) - a(\mathbf{v}), \tag{2.13}$$

$$\text{maj}(\mathbf{T}) := \text{maj}(\mathbf{v}) := \sum_{x \in \text{Des}(\mathbf{v})} 1 + l(x). \tag{2.14}$$

Using the left diagram in Fig. 4 as an example again, we have $a(\mathbf{v}) = 3$, $\text{inv}(\mathbf{T}) = 4 - 3 = 1$, and $\text{maj}(\mathbf{T}) = 9$. A simple proof that $\text{inv}(\mathbf{T})$ is always nonnegative can be found in [13, Sect. 2].

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

$$\begin{aligned} \tilde{H}_{\mu/\rho}(X; q, t) &:= \sum_{\substack{\mathbf{v} \in \text{TR}(\mu/\rho) \\ \mathbf{T} \in \text{SYT}(\mathbf{v})}} q^{\text{inv}(\mathbf{T})} t^{\text{maj}(\mathbf{T})} F_{\sigma(\mathbf{T})} \\ &= \sum_{\mathbf{v} \in \text{TR}(\mu/\rho)} q^{-a(\mathbf{v})} t^{\text{maj}(\mathbf{v})} G_{\mathbf{v}}(X; q). \end{aligned} \tag{2.15}$$

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 [22] (see also [12, Eq. (2.30)]) that

$$\tilde{H}_{\mu/\rho}(X; q, t) = \tilde{H}_{\tilde{\mu}/\tilde{\rho}}(X; t, q). \tag{2.16}$$

2.4 Dual equivalence graphs

We now provide the necessary definitions and results from [3]. We begin by recalling Mark Haiman’s dual to the fundamental Knuth equivalences defined in (2.2).

Definition 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

$$\begin{aligned} d_i(\dots i - 1 \dots i \dots i + 1 \dots) &= (\dots i - 1 \dots i \dots i + 1 \dots), \\ d_i(\dots i + 1 \dots i \dots i - 1 \dots) &= (\dots i + 1 \dots i \dots i - 1 \dots), \\ d_i(\dots i \dots i - 1 \dots i + 1 \dots) &= (\dots i + 1 \dots i - 1 \dots i \dots), \\ d_i(\dots i - 1 \dots i + 1 \dots i \dots) &= (\dots i \dots i + 1 \dots i - 1 \dots). \end{aligned} \tag{2.17}$$

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) and (2.17) that d_i is related to K_i by

$$d_i(w) = (K_i(w^{-1}))^{-1}. \tag{2.18}$$

By (2.1) and (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 2.4 ([16, 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 2.4 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 [16, Lemma 2.3] that for all $1 < i, j < n$,

$$Q(K_j \circ d_i(w)) = Q(K_j(w)) = Q(d_i \circ K_j(w)). \tag{2.19}$$

Applying (2.18) and (2.19), yields

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

In particular,

$$P(d_i \circ K_j(w)) = P(K_j \circ d_i(w)) = P(d_i(w)). \tag{2.20}$$

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 2.4, 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}_λ (see Fig. 5).

Definition 2.2 and Theorem 2.4 determine the connection between Schur functions and dual equivalence graphs as highlighted in [3, Cor. 3.10]. Given any standard dual equivalence graph $\mathcal{G}_\lambda = (V, \sigma, E)$,

$$\sum_{v \in V} F_{\sigma(v)} = s_\lambda. \tag{2.21}$$

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

Definition 2.5 An *edge colored graph* consists of the following data:

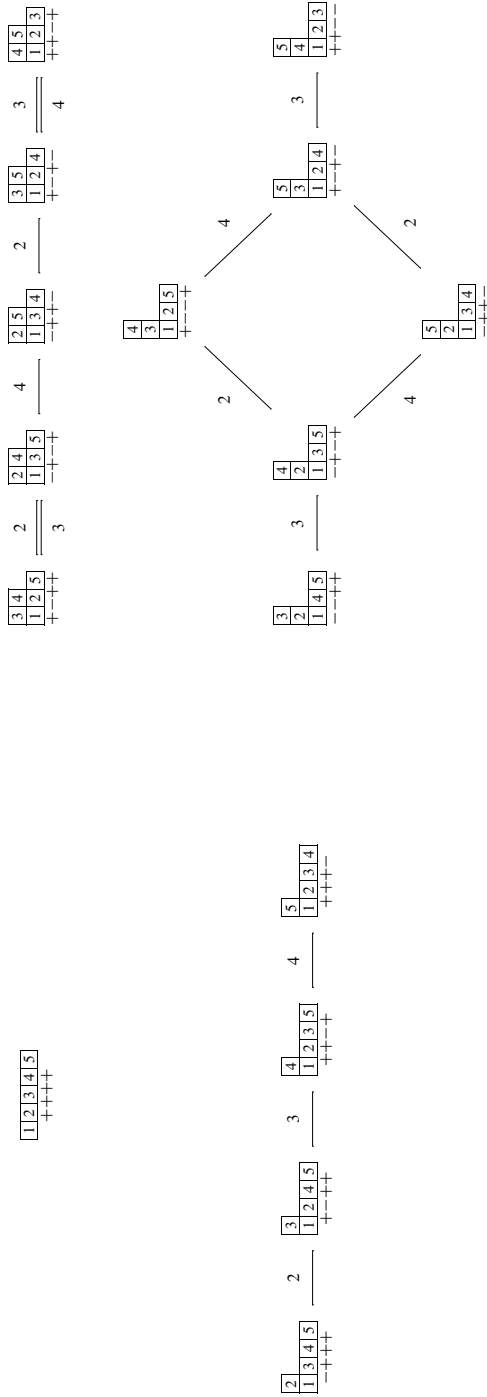


Fig. 5 The standard dual equivalence graphs on partitions of 5 up to conjugation

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 $\sigma : V \rightarrow \{\pm 1\}^{N-1}$ for some positive integer $N \geq n$.

We denote a signed colored graph by $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 Sect. 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 2.6 A map $\phi : \mathcal{G} \rightarrow \mathcal{H}$ between edge colored graphs $\mathcal{G} = (V, E_{m+1} \cup \dots \cup E_{n-1})$ and $\mathcal{H} = (V', E'_{m+1} \cup \dots \cup E'_{n-1})$ 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 $\phi : \mathcal{G} \rightarrow \mathcal{H}$ between signed colored graphs $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ and $\mathcal{H} = (V', \sigma', E'_{m+1} \cup \dots \cup E'_{n-1})$ 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 2.7 ([3, Proposition 3.11]) *If $\phi : \mathcal{G}_\lambda \rightarrow \mathcal{G}_\mu$ 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 2.8 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 2.9 A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 Figs. 6 and 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 $(V, \sigma, E_{m+1} \cup \dots \cup E_i)$ 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*.

Remark 2.10 The following are immediate consequences of Definition 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 [3].
4. It is an instructional exercise to check that if Axioms 1, 2, and 6 are obeyed, then Axiom 4 on two consecutive colors implies Axiom 4 on three consecutive colors.
5. A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 2.11 ([3, Theorem 3.5]) *For any $\lambda \vdash n$, \mathcal{G}_λ is an (n, n) -dual equivalence graph.*



Fig. 6 The allowable $E_{i-1} \cup E_i$ components of Axiom 4

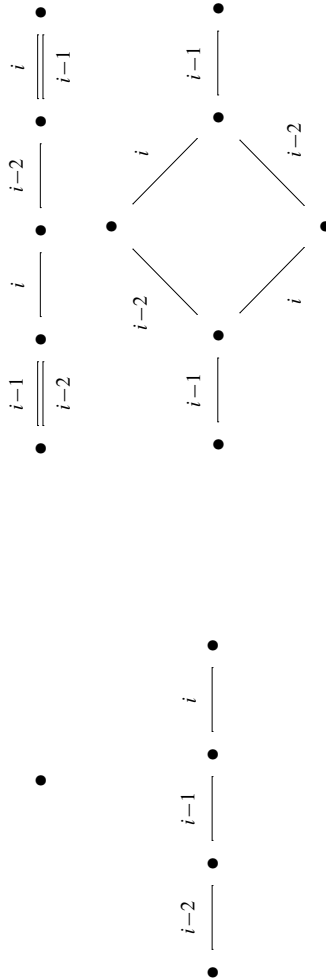


Fig. 7 The allowable $E_{i-2} \cup E_{i-1} \cup E_i$ components of Axiom 4

The converse is also true.

Theorem 2.12 ([3, Theorem 3.9]) *Every connected component of an (n, n) -dual equivalence graph is isomorphic to a unique \mathcal{G}_λ .*

The key to proving Theorem 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 2.13 Fix any partitions $\lambda \subset \mu$ with $|\lambda| = n$ and $|\mu| = N$ and a skew tableau A of shape μ/λ with entries $n + 1, \dots, N$. Define the set of standard Young tableaux *augmented by A* , denoted $\text{ASYT}(\lambda, A)$, as the set of $T \in \text{SYT}(\mu)$ such that the restriction of T to μ/λ is A . Further, define a signed colored graph $\mathcal{G}_{\lambda,A}$ on $\text{ASYT}(\lambda, A)$ with signature given by inverse descents and edges given by the non-trivial orbits of d_i for $1 < i < n$.

Remark 2.14 The graph $\mathcal{G}_{\lambda,A}$ is isomorphic to an (n, N) -component of \mathcal{G}_μ , and every (n, N) -component of \mathcal{G}_μ is isomorphic to some $\mathcal{G}_{\lambda,A}$, as is clear from the definition of $\mathcal{G}_{\lambda,A}$. By Part 1 of Remark 2.10 and Theorem 2.11, $\mathcal{G}_{\lambda,A}$ is therefore a dual equivalence graph with (n, n) -restriction isomorphic to \mathcal{G}_λ . Applying Theorem 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 2.15 ([3, 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}_λ for some partition λ of n . Then ϕ extends to an isomorphism $\tilde{\phi}: \mathcal{G} \rightarrow \mathcal{G}_{\lambda,A}$, 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 2.15, 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 2.15, then the unique extension is isomorphic to \mathcal{G}_μ , where μ is the union of λ and the cell of A containing $(n + 1)$.

Definition 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 [3], the following is an immediate consequence of the proof of Theorem 2.12.

Corollary 2.17 ([3, 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}_μ . Then $\mathcal{C} \cup (\cup \mathcal{B})$ is isomorphic to the $(n, n + 1)$ restriction of \mathcal{G}_μ , where the union is over all \mathcal{B} that are neighbors of \mathcal{C} in \mathcal{G} . Furthermore, there exists a morphism $\phi: \mathcal{G} \rightarrow \mathcal{G}_\mu$.

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 3.13 and 3.14, the improved axiomatization of dual equivalence graphs given in Theorem 3.17, and the more specific criterion for satisfying the dual equivalence axioms given in Corollary 3.20. In the process, a number of smaller results about the structure of dual equivalence graphs are highlighted.

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 [16, Lemma 2.3]. It lays out a fundamental relationship between the dual equivalence and Knuth equivalence maps defined in (2.2) and (2.17).

Lemma 3.1 *Given any $w \in S_n$ and any $1 < i < n, 1 < j < n$, we have $K_j \circ d_i(w) = d_i \circ K_j(w)$ 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). 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 (2.19) and (2.20),

$$(P(d_i \circ K_j(w)), Q(d_i \circ K_j(w))) = (P(K_j \circ d_i(w)), Q(K_j \circ d_i(w))). \tag{3.1}$$

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 Fig. 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. □

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

Theorem 3.2 *The function $P: S_n \rightarrow \text{SYT}(n)$ 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*

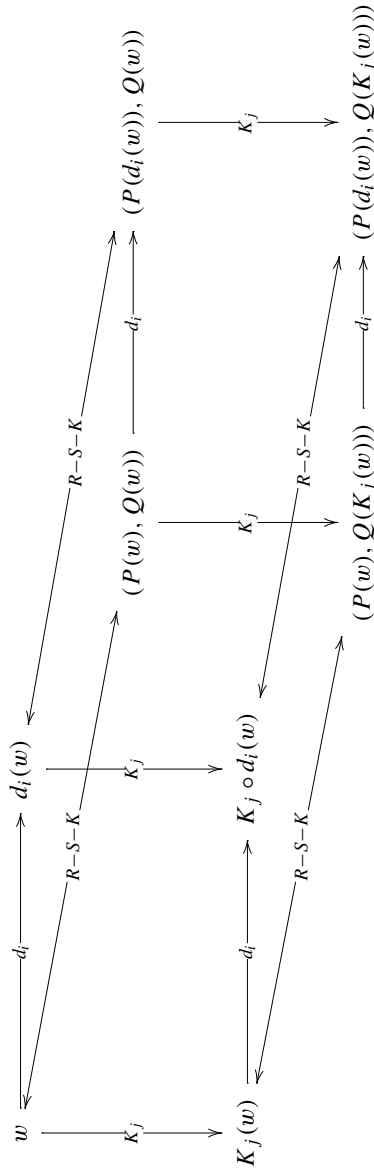


Fig. 8 The commutativity of d_i , K_j , and the R-S-K correspondence acting on a word $w \in S_n$

\mathcal{G}_n to \mathcal{G}_λ 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 $\text{SYT}(\lambda/\rho)$, 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 Sect. 2.2, each $w \in S_n$ is in a Knuth class with the reading word of some $T \in \text{SYT}(\lambda)$. 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 Sects. 2.1, 2.2, and 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 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 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 that every $T \in \text{SYT}(n)$ 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}_λ 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 (2.17) that $\text{SYT}(\lambda/\rho)$ 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 3.2, it makes sense to extend the notation \mathcal{G}_λ to the skew case $\mathcal{G}_{\lambda/\rho}$. That is, $\mathcal{G}_{\lambda/\rho}$ is the dual equivalence graph with vertices in $\text{SYT}(\lambda/\rho)$, 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 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 3.2 leads to a number of simple corollaries. The first is a well-known fact (see [8, 24], or [25]), while the rest help to illuminate the structure of dual equivalence graphs.

Corollary 3.4 (Littlewood–Richardson Rule) *The skew Schur functions $s_{\nu/\lambda}$ are Schur positive. Moreover, for all λ/ρ ,*

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

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

Proof By (2.21) and Theorem 3.2, we can interpret $c_{\rho, \mu}^{\lambda}$ as the number of connected components isomorphic to \mathcal{G}_{μ} in $\mathcal{G}_{\lambda/\rho}$. Applying Theorem 3.2, any component that is isomorphic to \mathcal{G}_{μ} has exactly one vertex v_{μ} whose row reading word is mapped to U_{μ} by P . Recall that $\text{SYam}(\mu)$ is defined to be the set of words w such that $P(w) = U_{\mu}$. Thus, $c_{\rho, \mu}^{\lambda} = |\{w \in \text{SYam}(\mu) : w \text{ is the row reading word of a filling in } \text{SYT}(\lambda/\rho)\}|$. □

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

Definition 3.5 Fix some signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 $(V, \sigma, E_{m+1} \cup \dots \cup E_{n-2})$ is termed *the downward restriction of \mathcal{G}* .

Corollary 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 $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ is connected. Let \mathcal{H} be the h -upward restriction of \mathcal{G} . It follows immediately from Definition 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 2.9 and 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 2.12. Then \mathcal{H} can be obtained by removing the 1 through h boxes from each vertex in \mathcal{G}_{λ} 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 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.

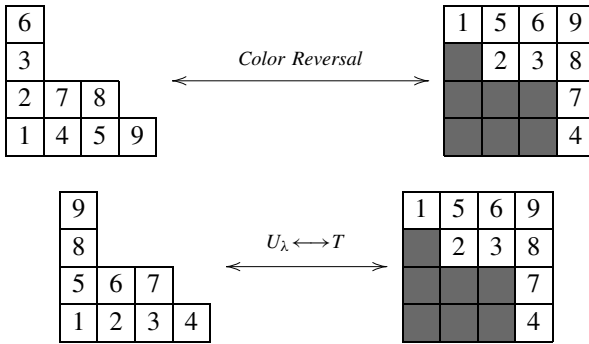


Fig. 9 At left, a tableau and its color reversal. At right, U_λ for $\lambda = (4, 3, 1, 1)$ and its corresponding skew tableau $T \in \text{SYT}(\mu/\rho)$

Definition 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 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 2.12 allows us to further reduce to the case where \mathcal{G} is an arbitrary \mathcal{G}_λ . 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 $\text{SYT}(\mu/\rho)$. The reader can check that $\mathcal{G}_{\mu/\rho}$ is isomorphic to the color reversal of \mathcal{G}_λ by simply rotating any filling of λ and then reversing the order of the numbers in the filling as in Fig. 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 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 $\text{sh}(v) = \lambda$.

Let $T \in \text{SYT}(\mu/\rho)$ 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 Fig. 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 Sect. 2.2. Therefore, the row reading word of T has shape λ , completing the proof. \square

Remark 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 [27, Sect. 3.20].

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 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 $\phi: \mathcal{G} \rightarrow \mathcal{H}$. 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 3.11 Let $\phi: \mathcal{G} \rightarrow \mathcal{H}$ be as in Lemma 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 3.10 are well-known properties of covering maps. For more details, see [10, Sect. 6.8].

Corollary 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 3.10, we need only show that ϕ is bijective. By Parts 2 and 5 of Lemma 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 2.12 allows us to assume that the domain of ϕ is some \mathcal{G}_λ . Now notice that in \mathcal{G}_λ there is only one vertex with the signature of U_λ , implying that $m = 1$. Thus, ϕ is a bijection. \square

In light of (2.21) and Part 5 of Lemma 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 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 $P: V \rightarrow \text{SYT}(n)$ induces a morphism $\phi: \mathcal{G} \rightarrow \bigcup_{\lambda \vdash n} \mathcal{G}_\lambda$. Furthermore,

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

Proof To show that P induces a morphism on \mathcal{G} , first notice that Theorem 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 3.1 and Theorem 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 (2.21), we need only show that $|\{V \cap \text{SYam}(\lambda)\}|$ is equal to the value of the index m of this restriction, as stated in Part 5 of Lemma 3.10. By definition, $\text{SYam}(\lambda)$

is the set of words mapped to U_λ by P , so $m = |\text{SYam}(\lambda) \cap V|$, completing the proof. \square

Theorem 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 3.14 *Let \mathcal{G} be an nonempty (n, n) -signed colored graph satisfying Axiom 1 and admitting a morphism $\phi: \mathcal{G} \rightarrow \mathcal{G}_\lambda$. 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 $\sigma: V \rightarrow \{\pm 1\}^{N-1}$ 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 3.10, ϕ is a surjective m to 1 map. Choose any $\mu \supset \lambda$ such that $|\text{SYT}(\mu)| \geq m$. Here, $N = |\mu|$. Let A be any skew tableau of shape μ/λ with the values $n + 1, \dots, N$. 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 $\tilde{\phi}: \tilde{\mathcal{G}} \rightarrow \mathcal{G}_{\lambda, A}$. Because $|\text{SYT}(\mu)|$ 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 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 3.10 and Theorem 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 $d_i(v')$. \square

Remark 3.15 Theorems 3.13 and 3.14 can both be extended to statements about (n, N) -signed colored graphs with morphisms to augmented dual equivalence graphs.

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 3.16 A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \cdots \cup E_{n-1})$ 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}, E_{i-2} is omitted if $i \leq m + 3$ or $i = m + 2$, respectively.

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

Theorem 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 3.18

1. We may readily classify the set of edge colored graphs described in Definition 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 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 $\mathcal{G} = (V, \sigma, E_{m+1} \cup \cdots \cup E_{n-1})$ 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 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 2.17, 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, $\mathcal{C} \cup (\cup \mathcal{B})$ 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 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 3.6. We now proceed by induction on n . Axiom 4⁺, when considered with Axioms 1, 2, and 3, implies that 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 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 2.12, \mathcal{C} is isomorphic to some \mathcal{G}_λ , and by Proposition 2.15, there is some \mathcal{G}_μ isomorphic to the unique extension of \mathcal{C} . Furthermore, Corollary 2.17 guarantees the existence of a morphism from \mathcal{G} to \mathcal{G}_μ . Label the vertices of \mathcal{G} , as well as the vertices of the downward restrictions of \mathcal{G} , by the set of tableaux in $\text{SYT}(\mu)$ as given by the image of this morphism. By Corollary 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}_μ .

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 Fig. 10. It follows from Theorem 2.4 that the number of $(n, n + 1)$ -components in \mathcal{G}_μ 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 2.17, \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 $n + 1$.

To apply Lemma 3.19, it suffices to show that there is a tableau $T \in \text{SYT}(\mu)$ 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 $S \in \text{SYT}(\mu)$ be some vertex of \mathcal{G}_μ . As in

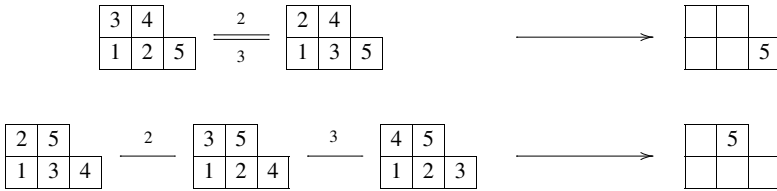


Fig. 10 The $(4, 4)$ -components of $\mathcal{G}_{(3,2)}$, at left, are represented via the cells containing 5, at right. Signatures are omitted

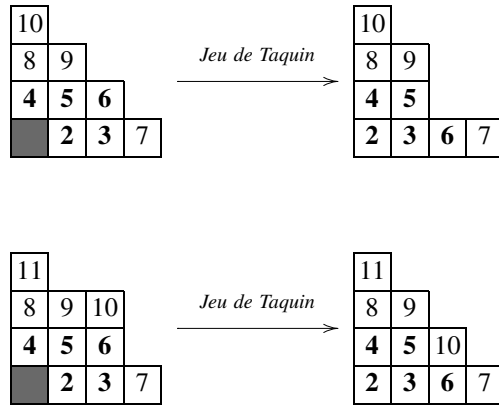


Fig. 11 On top, a filling chosen to lose the corner containing 6 to jeu de taquin. On bottom, a filling of a nonstaircase shape that does not lose any Northeast corners to jeu de taquin

the proof of Corollary 3.6, the set of vertices in the same component of the upward restriction of \mathcal{G}_μ 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 $\text{sh}(J(S))$ 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 3.1. Theorem 2.4 guarantees that $n + 1$ may then be moved to any Northeast corner of $\text{sh}(J(S))$. 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}_μ with no 2-edges that connects S to a tableau with the value $n + 1$ in c . Therefore, to apply Lemma 3.19, we need only find $T \in \text{SYT}(\mu)$ 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 that 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 Fig. 11.

As a second case, suppose that μ is not a staircase shape, i.e., μ is not of shape $(\mu_1, \mu_1 - 1, \dots, 2, 1)$. Then we claim that there exists some $T \in \text{SYT}(\mu)$ 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 Fig. 11 for an example.

We have shown that \mathcal{G} is a dual equivalence graph if μ is not a staircase 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 3.18, we may assume that \mathcal{G} has type (n, n) . By Theorem 2.12, \mathcal{G} is isomorphic to some \mathcal{G}_λ , 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 2.17, each of these components is connected to two neighboring components. If Axiom 6 is not satisfied, then Lemma 3.19 implies that these two components cannot be neighbors of each other. Rather, the $(5, 6)$ -components must form a loop. The smallest example is presented in Fig. 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 Fig. 13. We may then omit signatures and relabel edges with elements in $\{i - 3, i - 2, i - 1, i\}$, as mentioned in Remark 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 $E_{i-3} \cup \dots \cup E_i$ that satisfy Axioms 1, 2, 3, 4, and 5 but not Axiom 6. The following corollary reformulates Theorem 3.17 in terms of \mathcal{F} .

Corollary 3.20 *Let $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 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 equiv-

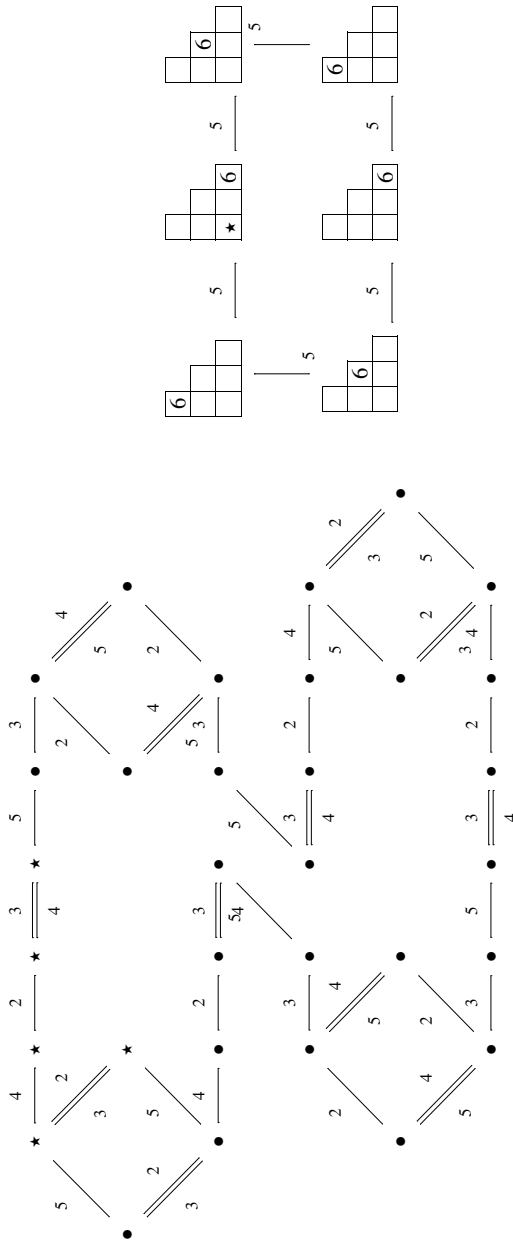


Fig. 12 Two representations of the smallest graph in \mathcal{F} with edge labels in $\{2, 3, 4, 5\}$. Starred vertices on the *left* correspond to the starred shape on the *right*

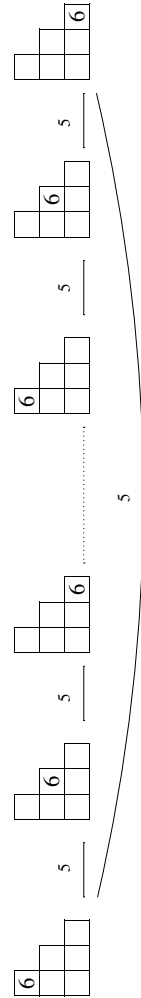


Fig. 13 A generic element of \mathcal{F} in the case where edge labels are in $\{2, 3, 4, 5\}$

alence graphs. Let $\mathcal{G} = (V, \sigma, E_{m+1} \cup \dots \cup E_{n-1})$ 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 4.11.

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.

4.1 The Schur expansion of $G_{\mathbf{v}}(X; q)$ when $\text{diam}(\mathbf{v}) \leq 3$

Recall the notation and vocabulary introduced in Sect. 2.3 on LLT polynomials and Macdonald polynomials. To state the main theorem of this section, we will also need the following definition.

Definition 4.1 Given a k -tuple of skew shapes \mathbf{v} , let $S(\mathbf{v})$ be the set of distinct shifted contents of the cells in \mathbf{v} . Define the diameter of \mathbf{v} , denoted $\text{diam}(\mathbf{v})$, as

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

See Fig. 14 for an example.

Remark 4.2 If $\mathbf{v} = (v^{(0)}, \dots, v^{(k-1)})$ is a tuple of skew shapes, it follows from Definition 4.1 that $\text{diam}(\mathbf{v}) \leq k + 1$. There are many examples where this bound is sharp. For instance, if \mathbf{v} is a k -tuple of straight shapes where $v^{(0)}$ has as least two columns and $|v^{(i)}| \geq 1$ for all $1 \leq i \leq k - 1$, then $\text{diam}(\mathbf{v}) = k + 1$. In Fig. 14, \mathbf{v}_3 is an example of such a tuple.

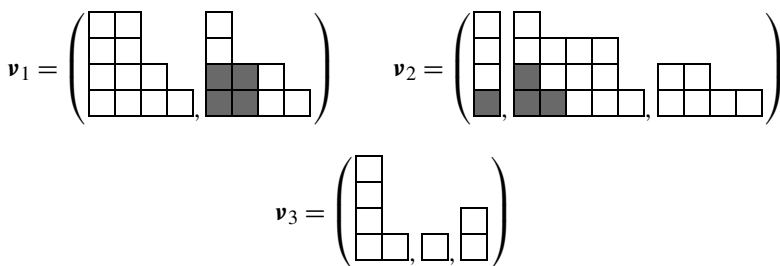


Fig. 14 The tuples \mathbf{v}_1 and \mathbf{v}_2 have diameter 3. The tuple \mathbf{v}_3 has diameter 4

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

$$\mathbf{T}(\mathbf{v}, \lambda) := \{ \mathbf{T} \in \text{SYT}(\mathbf{v}) : \text{the shifted content word of } \mathbf{T} \text{ is in } \text{SYam}(\lambda) \}, \tag{4.1}$$

$$T(\mu/\rho, \lambda) := \{ T \in \text{SYT}(\mu/\rho) : \text{the content reading word of } T \text{ is in } \text{SYam}(\lambda) \}. \tag{4.2}$$

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

Theorem 4.3 *Let \mathbf{v} be any k -tuple of skew shapes with $\text{diam}(\mathbf{v}) \leq 3$. Then*

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

The proof of Theorem 4.3 is postponed until Sect. 4.2. As mentioned in Remark 4.2, the set of \mathbf{v} such that $\text{diam}(\mathbf{v}) \leq 3$ properly contains the set of \mathbf{v} that are 2-tuples.

The next corollary follows immediately by applying Theorem 4.3 to the definition of modified Macdonald polynomials in (2.15). We also use the fact that tuples of ribbons in $\text{TR}(\mu/\rho)$ 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 $\text{TR}(\mu/\rho)$ have the same set of distinct shifted contents.

Corollary 4.4 *Let μ/ρ be a skew shape not containing (3, 3) or (4) as a subdiagram. Then*

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

In particular, Corollary 4.4 applies to all $\tilde{H}_{\mu}(X; q, t)$ with $\mu_1 \leq 3$ where $\mu_2 \leq 2$.

Remark 4.5 The conditions on \mathbf{v} and μ/ρ in Theorem 4.3 and Corollary 4.4, respectively, are sharp in the following sense. Let $\lambda = (2, 2)$ and $\mathbf{v} = ((2), (1), (1))$ or $((1), (1), (1), (1))$. In particular, $\text{diam}(\mathbf{v}) = 4$, and \mathbf{v} has the diagram

$$\left(\begin{array}{cccc} \square & \square & \square & \square \\ & \square & \square & \square \end{array} \right) \quad \text{or} \quad \left(\begin{array}{cccc} \square & \square & \square & \square \\ & \square & \square & \square \\ & & \square & \square \\ & & & \square \end{array} \right).$$

Then

$$G_{\mathbf{v}}(X; q)|_{q^2 s_{\lambda}} = 1 \quad \text{and} \quad \sum_{\mathbf{T} \in \mathbf{T}(\mathbf{v}, \lambda)} q^{\text{inv}_k(\mathbf{T})} s_{\lambda}|_{q^2 s_{\lambda}} = 0. \tag{4.3}$$

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

$$\tilde{H}_{\mu}(X; q, t)|_{q^2 s_{\rho}} = 1 \quad \text{and} \quad \sum_{T \in T(\mu, \lambda)} q^{\text{inv}(T)} t^{\text{maj}(T)} s_{\lambda}|_{q^2 s_{\lambda}} = 0. \tag{4.4}$$

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

$$\tilde{H}_{\mu}(X; q, t)|_{q^2 s_{\rho}} = 1 \quad \text{and} \quad \sum_{T \in T(\mu, \lambda)} q^{\text{inv}(T)} t^{\text{maj}(T)} s_{\lambda}|_{q^2 s_{\lambda}} = 0. \tag{4.5}$$

Corollary 4.6 *Let μ/ρ be a skew shape not containing $(2, 2, 2)$ or $(1, 1, 1, 1)$ as a subdiagram. Then*

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

Proof The corollary follows from (2.16) and Corollary 4.4. □

In particular, Corollary 4.6 applies to all $\tilde{H}_{\mu}(X; q, t)$ where μ has at most three rows and $\mu_3 \leq 1$.

4.2 LLT graphs

The goal of this section is to prove Theorem 4.3. We begin by following [3] in defining an involution that will provide the edge sets of a signed colored graph. In this section, \mathbf{v} will always denote a k -tuple of skew shapes whose sizes sum to $|\mathbf{v}| = n$. Also, w will always denote a permutation in S_n .

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

$$\begin{aligned} \tilde{d}_i(\dots i - 1 \dots i \dots i + 1 \dots) &= (\dots i - 1 \dots i \dots i + 1 \dots), \\ \tilde{d}_i(\dots i + 1 \dots i \dots i - 1 \dots) &= (\dots i + 1 \dots i \dots i - 1 \dots), \\ \tilde{d}_i(\dots i \dots i - 1 \dots i + 1 \dots) &= (\dots i - 1 \dots i + 1 \dots i \dots), \\ \tilde{d}_i(\dots i \dots i + 1 \dots i - 1 \dots) &= (\dots i + 1 \dots i - 1 \dots i \dots). \end{aligned} \tag{4.6}$$

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 \mathbf{v} 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 $\tau = \tau_1 \tau_2 \dots \tau_n$ by

$$\tau_i = \max\{j \in [n]: \tilde{c}_j - \tilde{c}_i \leq k\}. \tag{4.7}$$

See Fig. 15 for an example. To emphasize the relationship between τ and \mathbf{v} , we will sometimes write $\tau = \tau(\mathbf{v})$. 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 [26]). Next, let $m(i)$ be the index of 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,



Fig. 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

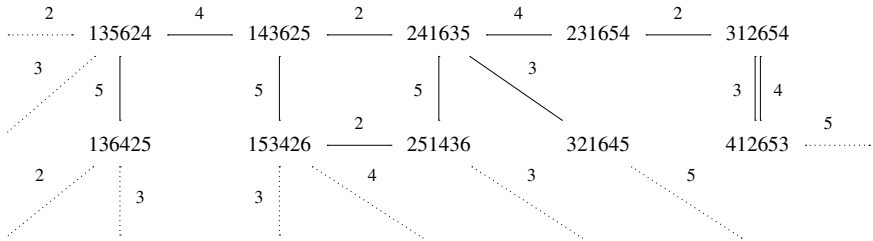


Fig. 16 A portion of \mathcal{L}_ν with signatures omitted. Here $\nu = ((2), (2), (1), (1))$

$$D_i^{(\tau)}(w) := \begin{cases} d_i(w), & \tau_{m(i)} < M(i), \\ \tilde{d}_i(w), & \tau_{m(i)} \geq M(i). \end{cases} \tag{4.8}$$

As an example, we may take $\tau = 456667899$ and $w = 534826179$, as in Fig. 15. Then $D_3^{(\tau)}(w) = \tilde{d}_3(w) = 542836179$ and $D_5^{(\tau)}(w) = d_5(w) = 634825179$.

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 $\tau = \tau(\nu)$, then $D_i^{(\tau)}$ takes shifted content words of standard fillings of ν to shifted content words of other standard fillings of ν . Thus, $D_i^{(\tau)}$ has a well-defined action on $\text{SYT}(\nu)$ 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 4.7 Given some tuple of skew shapes ν , the LLT graph $\mathcal{L}_\nu = (V, \sigma, E)$ is defined to be the (n, n) -signed colored graph with the following data:

1. $V = \{w \in S_n : w \text{ is the shifted content word of some } \mathbf{T} \in \text{SYT}(\nu)\}$.
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 $1 < i < |\nu|$, where $\tau = \tau(\nu)$.

Example 4.8 Consider $\nu = ((2), (2), (1), (1))$. A portion of the LLT graph \mathcal{L}_ν is presented in Fig. 16. Here, \mathcal{L}_ν is a subgraph of $\mathcal{G}_6^{(\tau)}$ with $\tau = 566666$. The entire connected component of the vertices in Fig. 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 $1 < i < 6$.

Remark 4.9

1. If $\mathbf{v} = (\lambda/\rho)$ is a 1-tuple, then each $\mathbf{T} \in \text{SYT}(\mathbf{v})$ corresponds to a unique $T \in \text{SYT}(\lambda/\rho)$ in the obvious fashion. In particular, the shifted content word of \mathbf{T} is equal to the content reading word of T . By Remark 3.3, $\mathcal{G}_{\lambda/\rho} \cong \mathcal{L}_{\mathbf{v}}$. In the case of straight shapes, $\mathbf{v} = (\lambda)$, Theorem 3.2 implies that $P: S_n \rightarrow \text{SYT}(n)$ induces an isomorphism from $\mathcal{L}_{\mathbf{v}}$ to \mathcal{G}_{λ} .
2. Given some w in the vertex set of $\mathcal{L}_{\mathbf{v}}$, we may readily describe the component of w in a restriction of $\mathcal{L}_{\mathbf{v}}$. Let w be the shifted content word of $\mathbf{T} \in \text{SYT}(\mathbf{v})$, and let \mathcal{C} be the (m, m) -component of w in $\mathcal{L}_{\mathbf{v}}$. The isomorphism type of \mathcal{C} can be found by removing the cells of \mathbf{T} containing values in $\{m + 1, \dots, n\}$, creating some $\mathbf{T}' \in \text{SYT}(\mathbf{v}')$ with shifted content word w' . Then \mathcal{C} is isomorphic to the component of w' in $\mathcal{L}_{\mathbf{v}'}$.

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

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

While LLT graphs do not necessarily satisfy Axiom 4 or Axiom 6, as can be seen in Example 4.8, they do satisfy a subset of the dual equivalence axioms. This is made precise in the following proposition and theorem.

Proposition 4.10 ([3, Prop. 4.6]) *Any LLT graph $\mathcal{L}_{\mathbf{v}}$ obeys Axioms 1, 2, 3, and 5. Furthermore, the $\text{inv}_{\mathbf{v}}$ statistic is constant on each connected component of $\mathcal{L}_{\mathbf{v}}$.*

With Proposition 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 4.11 *The LLT graph $\mathcal{L}_{\mathbf{v}}$ is a dual equivalence graph if and only if $\text{diam}(\mathbf{v}) \leq 3$.*

Proof First suppose that $\text{diam}(\mathbf{v}) \leq 3$. By Proposition 4.10, $\mathcal{L}_{\mathbf{v}}$ obeys Axioms 1, 2, 3, and 5. By Theorem 3.17, we need only show that $\mathcal{L}_{\mathbf{v}}$ satisfies Axiom 4⁺ to prove that $\mathcal{L}_{\mathbf{v}}$ is a dual equivalence graph.

We begin by showing that $\mathcal{L}_{\mathbf{v}}$ obeys Axiom 4. Applying the same logic as Part 1 of Remark 3.18, we may use restrictions to only consider signed colored graphs of type $(5, 5)$. By Part 2 of Remark 4.9, each connected component of any such restriction is a component of $\mathcal{L}_{\boldsymbol{\mu}}$ for some $\boldsymbol{\mu}$ with $\text{diam}(\boldsymbol{\mu}) \leq 3$. We may then reduce to the case where $|\mathbf{v}| = 5$ and $\text{diam}(\mathbf{v}) \leq 3$. It thus suffices to consider $\mathcal{G}_5^{(\tau)}$ for all $\tau = \tau(\mathbf{v})$ with \mathbf{v} satisfying these properties, in which case, it can be checked explicitly that all components of $\mathcal{G}_5^{(\tau)}$ that are contained in some $\mathcal{L}_{\mathbf{v}}$ with $\text{diam}(\mathbf{v}) \leq 3$

are dual equivalence graphs. This fact is verified via computer at Springer site ([Electronic supplementary material](#)), cited as [23].

To demonstrate Axiom 4^+ , we may similarly inspect $\mathcal{G}_6^{(\tau)}$ for all τ that are derived from ν with $|\nu| = 6$ and $\text{diam}(\nu) \leq 3$. It is possible to explicitly check that \mathcal{L}_ν 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 3.20. This fact is also verified via computer at [23]. Thus, \mathcal{L}_ν is a dual equivalence graph whenever $\text{diam}(\nu) \leq 3$.

Now suppose that $\text{diam}(\nu) \geq 4$. By the definition of diameter, there exist four cells of ν 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 ν . Similarly, c_2, c_3 , and c_4 occur in different skew tableaux of ν . There is then some value of i and some standard filling of ν 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}_ν to the edge colored graph $(V, E_i \cup E_{i+1})$ contains a component with at least three distinct edges. Therefore, \mathcal{L}_ν violates Axiom 4 and is not a dual equivalence graph. □

Remark 4.12 By applying [2, Prop. 5.3], it can be proven that \mathcal{L}_ν is a dual equivalence graph whenever ν 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 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 $P: V \rightarrow \text{SYT}(n)$ induces an isomorphism from \mathcal{C} to some standard dual equivalence graph \mathcal{G}_λ .*

Proof By Proposition 4.10, \mathcal{C} satisfies Axioms 1, 2, 3, and 5. By Theorem 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 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 [23].

To demonstrate that P induces a morphism, we show that \mathcal{C} satisfies the hypotheses of Theorem 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 3.12. \square

Lemma 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 4.3. Instead, we will first need to identify specific components of LLT graphs satisfying the hypotheses of Lemma 4.13.

Lemma 4.14 *Let \mathbf{v} be a tuple of skew shapes such that $|\mathbf{v}| = n$ and $\text{diam}(\mathbf{v}) \leq 3$. Let $\mathcal{C} = (V, \sigma, E)$ be a connected component of $\mathcal{L}_{\mathbf{v}}$, let $\tau = \tau(\mathbf{v})$, and let v be any vertex in \mathcal{C} . If $\text{sh}(v)$ 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 $\mathbf{T} = (T^{(0)}, \dots, T^{(k-1)}) \in \text{SYT}(\mathbf{v})$. 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 4.13. Thus $\text{sh}(D_i^{(\tau)}(v)) = \text{sh}(v)$. 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 $\text{sh}(v)$ 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 Sect. 2.2. Suppose that $D_i^{(\tau)}(v) = \tilde{d}_i(v) \neq v$ 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 $\text{diam}(\mathbf{v}) \leq 3$, the cells of \mathbf{v} 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 \mathbf{v} 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 . \square

Corollary 4.15 *Let \mathbf{v} be a tuple of skew shapes such that $|\mathbf{v}| = n$ and $\text{diam}(\mathbf{v}) \leq 3$. Let $\mathcal{C} = (V, \sigma, E)$ be a connected component of $\mathcal{L}_{\mathbf{v}}$, and let w be any vertex in \mathcal{C} . If $\lambda = \text{sh}(w)$ has strictly less than three rows, then $P: V \rightarrow \text{SYT}(n)$ induces an isomorphism from \mathcal{C} to the standard dual equivalence graph \mathcal{G}_{λ} .*

Proof The corollary follows immediately from Lemmas 4.13 and 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 4.3.

Lemma 4.16 *Let μ and \mathbf{v} be tuples of skew shapes such that $\text{diam}(\mu), \text{diam}(\mathbf{v}) \leq 3$ and $|\mu| = |\mathbf{v}| = n$. Let \mathcal{C} and \mathcal{D} be connected components of \mathcal{L}_{μ} and $\mathcal{L}_{\mathbf{v}}$, respectively, and let $\phi: \mathcal{C} \rightarrow \mathcal{D}$ be an isomorphism. If w is a vertex in \mathcal{C} and $\lambda \vdash n$, then $w \in \text{SYam}(\lambda)$ if and only if $\phi(w) \in \text{SYam}(\lambda)$.*

Proof Suppose, for the sake of contradiction, that $v = \phi(w) \in \text{SYam}(\lambda)$ and $w \notin \text{SYam}(\lambda)$. 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 Fig. 17 for an example). Now define

$$p = \sum_{j=1}^{m-2} \lambda_j, \quad q = \lambda_{m-1} + \lambda_m, \quad S_i = p + i, \quad \text{and} \quad S = \{S_1, S_2, \dots, S_q\}. \quad (4.9)$$

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 Sect. 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 $S_1 S_2 \cdots S_{\lambda_{m-1}} S_q$ as an increasing subword (see Fig. 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 $T \neq U_{\lambda}$.

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 4.9. Specifically, consider $\mathbf{T} \in \text{SYT}(\mu)$ 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 μ' with shifted content



Fig. 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 $S_1 S_2 \cdots S_{\lambda_{m-1}} S_q = 4569$

word w' . Then C' is isomorphic to the component of $\mathcal{L}_{\mu'}$ 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 $\mathcal{L}_{\mu'}$, we will apply Corollary 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 4.15, P induces an isomorphism from C' to the standard dual equivalence graph $\mathcal{G}_{\text{sh}(w')}$. The length of the first row in $\text{sh}(w')$ is the length of the longest increasing subword in w' , as noted in Sect. 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 $\text{sh}(w')$ has exactly two rows, it follows that $\text{sh}(w') = (l, q - l)$. Therefore, $C' \cong \mathcal{G}_{(l, q-l)}$.

Alternatively, we may find the isomorphism type of C' by applying the same restrictions to \mathcal{D} . Let D' be the (q, q) -component of v in the p -upward restriction of \mathcal{D} . In particular, $C' \cong D'$. As before, D' is isomorphic to the component of some $\mathcal{L}_{\nu'}$ containing the vertex v' , where v' is obtained by subtracting p from each of the values in v_S . Also, as before, $\text{sh}(v')$ has two rows. Therefore Corollary 4.15 guarantees that $D' \cong \mathcal{G}_{\text{sh}(v')}$. 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, $\text{sh}(v') = (\lambda_{m-1}, \lambda_m)$. Since $l > \lambda_{m-1}$, $C' \cong \mathcal{G}_{(l, q-l)} \not\cong \mathcal{G}_{(\lambda_{m-1}, \lambda_m)} \cong D'$, a contradiction. Therefore, $w \in \text{SYam}(\lambda)$ whenever $v = \phi(w) \in \text{SYam}(\lambda)$.

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

Proof of Theorem 4.3 We begin by reducing Theorem 4.3 to a statement about signed colored graphs. Let V_1, V_2, \dots, V_m be the vertex sets of the connected components of $\mathcal{L}_\nu = (V, \sigma, E)$. Applying (2.8) and Definition 4.7, we have

$$\begin{aligned}
 G_\nu(X; q) &= \sum_{\mathbf{T} \in \text{SYT}(\nu)} q^{\text{inv}_k(\mathbf{T})} F_{\sigma(\mathbf{T})}(X) = \sum_{v \in V} q^{\text{inv}_\nu(v)} F_{\sigma(v)}(X) \\
 &= \sum_{j=1}^m \sum_{v \in V_j} q^{\text{inv}_\nu(v)} F_{\sigma(v)}(X), \tag{4.10}
 \end{aligned}$$

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 ν 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 ν is defined via the signature of its row reading word.

To further simplify (4.10), we turn our attention to the individual components of \mathcal{L}_ν when $\text{diam}(\nu) \leq 3$. Let $\mathcal{C} = (V_j, \sigma, E)$ be a connected component of \mathcal{L}_ν , and choose any fixed $v_j \in V_j$. By Proposition 4.10, the inv_ν statistic is constant on V_j . By Theorems 2.12 and 4.11, \mathcal{C} is isomorphic to \mathcal{G}_λ for some $\lambda \vdash |\nu|$. From (2.21) it follows that

$$\sum_{v \in V_j} q^{\text{inv}_\nu(v)} F_{\sigma(v)}(X) = q^{\text{inv}_\nu(v_j)} s_\lambda. \tag{4.11}$$

To prove Theorem 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 $\text{SYam}(\lambda)$.

We will find such a v_j explicitly by applying Lemma 4.16. Consider the LLT graph \mathcal{L}_λ , where λ is the 1-tuple (λ) . It follows from Part 2 of Remark 4.9 that $\text{diam}(\lambda) \leq 2$. By Part 1 of Remark 4.9, P induces an isomorphism from \mathcal{L}_λ to \mathcal{G}_λ . In particular, the vertex set of \mathcal{L}_λ has exactly one vertex w such that $P(w) = U_\lambda$. Thus, $w \in \text{SYam}(\lambda)$, and w is the only vertex of \mathcal{L}_λ that is a standardized Yamanouchi word. Because \mathcal{C} is isomorphic to \mathcal{G}_λ , and \mathcal{G}_λ is isomorphic to \mathcal{L}_λ , there exists an isomorphism $\phi: \mathcal{L}_\lambda \rightarrow \mathcal{C}$. Applying Lemma 4.16, we get that $v_j = \phi(w)$ is the unique standardized Yamanouchi word in V_j . □

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 [20], while the latter has already seen some progress in [4]. 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 Sect. 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 [14].

Is there a good description of when a component of an LLT graph is a dual equivalence graph? While Theorem 4.11 classifies when an LLT graph is a dual equivalence graph, it makes no claims about specific components of LLT graphs in the

$\text{diam}(\mathbf{v}) \geq 4$ 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 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.

Acknowledgements We would like to thank David Anderson, Sami Assaf, Jonathan Browder, Nantel Bergeron, William McGovern, and Isabella Novik for inspiring conversations on this topic. We would also like to thank the referees for their thorough feedback and thoughtful review. Particular thanks belong to Sara Billey for her invaluable guidance throughout the process of creating this paper.

Partially supported by DMS-1101017 from the NSF.

References

1. Assaf, S.H.: Dual equivalence graphs, ribbon tableaux and Macdonald Polynomials. Ph.D. thesis, University of California at Berkeley (2007)
2. Assaf, S.H.: A generalized major index statistic. *Sémin. Lothar. Comb.* **60**, (2008/09)
3. Assaf, S.H.: Dual equivalence graphs I: A combinatorial proof of LLT and Macdonald polynomials (2011). [arXiv:1005.3759v2](https://arxiv.org/abs/1005.3759v2) [math.CO]
4. Assaf, S.H., Billey, S.: Affine dual equivalence and k -Schur functions. *J. Comb.* **3**(3), 343–399 (2012)
5. Assaf, S.H., Bergeron, N., Sottile, F.: Multiplying Schubert polynomials by Schur polynomials (manuscript) (2012)
6. Carré, C., Leclerc, B.: Splitting the square of a Schur function into its symmetric and antisymmetric parts. *J. Algebr. Comb.* **4**(3) (1995)
7. Fishel, S.: Statistics for special q, t -Kostka polynomials. *Proc. Am. Math. Soc.* **123**(10), 2961–2969 (1995)
8. Fulton, W.: Young Tableaux, with Applications to Representation Theory and Geometry. London Mathematical Society Student Texts, vol. 35. Cambridge University Press, Cambridge (1997)
9. Gessel, I.M.: Multipartite P -partitions and inner products of skew Schur functions. *Contemp. Math.* **34**, 289–317 (1984)
10. Godsil, C., Royle, G.: Algebraic Graph Theory. Graduate Texts in Mathematics, vol. 207. Springer, New York (2004)
11. Haglund, J.: A combinatorial model for the Macdonald polynomials. *Proc. Natl. Acad. Sci. USA* **101**(46), 16127–16131 (2004)
12. Haglund, J.: The q, t -Catalan Numbers and the Space of Diagonal Harmonics. University Lecture Series, vol. 41. Am. Math. Soc., Providence (2008)
13. Haglund, J., Haiman, M., Loehr, N.: A combinatorial formula for Macdonald polynomials. *J. Am. Math. Soc.* **18**(3) (2005)
14. Haglund, J., Haiman, M., Loehr, N., Remmel, J.B., Ulyanov, A.: A combinatorial formula for the character of the diagonal coinvariants. *Duke Math. J.* **126**(2), 195–232 (2005)
15. Haiman, M.: Hilbert schemes, polygraphs and the Macdonald positivity conjecture. *J. Am. Math. Soc.* **14**(4), 941–1006 (2001)
16. Haiman, M.D.: Dual equivalence with applications, including a conjecture of Proctor. *Discrete Math.* **99**(1–3), 79–113 (1992)

17. Knuth, D.E.: Permutations, matrices, and generalized Young tableaux. *Pac. J. Math.* **34**, 709–727 (1970)
18. Lapointe, L., Morse, J.: Tableaux statistics for two part Macdonald polynomials. In: *Algebraic Combinatorics and Quantum Groups*, pp. 61–84. World Scientific, Singapore (2003)
19. Lascoux, A., Leclerc, B., Thibon, J.-Y.: Ribbon tableaux, Hall–Littlewood functions, quantum affine algebras, and unipotent varieties. *J. Math. Phys.* **38**(2), 1041–1068 (1997)
20. Loehr, N.A., Warrington, G.S.: Quasisymmetric expansion of Schur function plethysms. *Proc. Am. Math. Soc.* **140**(4), 1159–1171 (2012)
21. Macdonald, I.G.: A new class of symmetric functions. In: *Actes du 20^e Séminaire Lotharingien*, vol. 372, pp. 131–171 (1988)
22. Macdonald, I.G.: *Symmetric Functions and Hall Polynomials*, 2nd edn. Oxford Mathematical Monographs. Oxford University Press, New York (1995)
23. Roberts, A.: Computer proofs: LLT polynomials and dual equivalence graphs (2012). See Springer site (Electronic supplementary material)
24. Sagan, B.E.: *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, 2nd edn. Graduate Texts in Mathematics, vol. 203. Springer, New York (2001)
25. Stanley, R.P.: *Enumerative Combinatorics*. Vol. 2. Cambridge Studies in Advanced Mathematics, vol. 62. Cambridge University Press, Cambridge (1999)
26. Stanley, R.P.: Catalan Addendum (2012). www-math.mit.edu/rstan/ec/catadd.pdf
27. Stanley, R.P.: *Enumerative Combinatorics*. Vol. 1, 2nd edn. Cambridge Studies in Advanced Mathematics, vol. 49. Cambridge University Press, Cambridge (2012)
28. van Leeuwen, M.A.A.: Some bijective correspondences involving domino tableaux. *Electron. J. Comb.* **7** (2000). Research Paper 35, 25 pp (electronic)
29. Zabrocki, M.: A Macdonald vertex operator and standard tableaux statistics for the two-column (q, t) -Kostka coefficients. *Electron. J. Comb.* **5** (1998). Research Paper 45, 46
30. Zabrocki, M.: Positivity for special cases of (q, t) -Kostka coefficients and standard tableaux statistics. *Electron. J. Comb.* **6**, 36 (1999). Research Paper 41