



DUAL EQUIVALENCE GRAPHS I: A NEW PARADIGM FOR SCHUR POSITIVITY

SAMI H. ASSAF

Department of Mathematics, University of Southern California, Los Angeles,
CA 90089-2532, USA;
email: shassaf@usc.edu

Received 19 December 2014; accepted 11 June 2015

Abstract

We make a systematic study of a new combinatorial construction called a dual equivalence graph. We axiomatize these graphs and prove that their generating functions are symmetric and Schur positive. This provides a universal method for establishing the symmetry and Schur positivity of quasisymmetric functions.

2010 Mathematics Subject Classification: 05E05 (primary); 05A30, 05E10 (secondary)

1. Introduction

Symmetric function theory plays an important role in many areas of mathematics, including combinatorics, representation theory, and algebraic geometry. Multiplicities of irreducible components, dimensions of algebraic varieties, and various other algebraic constructions that require the computation of certain integers may often be translated to the computation of the Schur coefficients of a given function. Thus a quintessential problem in symmetric functions is to prove that a given function has nonnegative integer coefficients when expressed as a sum of Schur functions. In this paper, we introduce a new combinatorial construction, called a *dual equivalence graph*, by which one can establish the symmetry and Schur positivity of a function.

To illustrate the general problem and to demonstrate this new solution, consider the problem of expanding the product of two Schur functions as a sum of Schur

functions, that is,

$$s_\mu s_\nu = \sum_{\lambda} c_{\mu,\nu}^\lambda s_\lambda. \quad (1.1)$$

Since Schur functions are a basis for symmetric functions, this problem is well posed, and since they are an integral basis, these so-called Littlewood–Richardson coefficients $c_{\mu,\nu}^\lambda$ are integers. In fact, they are nonnegative. One way to see this is to realize that the Schur functions are the characters for irreducible representations of the general linear group, and so $c_{\mu,\nu}^\lambda$ counts multiplicities of irreducible representations in tensor products. Another way to understand the nonnegativity is to realize Schur polynomials as Schubert classes for the cohomology of the Grassmannian, and so $c_{\mu,\nu}^\lambda$ counts objects in the intersection of Schubert varieties. The celebrated Littlewood–Richardson rule gives a direct combinatorial interpretation for $c_{\mu,\nu}^\lambda$ without appealing to representation theory or geometry. Briefly, $c_{\mu,\nu}^\lambda$ counts the number of standard Young tableaux of shape μ adjoin ν that rectify to a specified standard Young tableau of shape λ . Dual equivalence graphs abstract this rule to a general tool with universal applicability.

The general set-up is as follows. Begin with a set \mathcal{A} of combinatorial objects together with a notion of a descent set Des sending an object to a subset of positive integers. For Littlewood–Richardson coefficients, \mathcal{A} is the set of standard Young tableaux and Des is the usual notion of descents. Optionally, we may also have a nonnegative, possibly multivariate, integer statistic associated to each object. Define the quasisymmetric generating function for \mathcal{A} with respect to Des by

$$f(X; q) = \sum_{T \in \mathcal{A}} q^{\text{stat}(T)} Q_{\text{Des}(T)}(X),$$

where Q denotes the fundamental basis for quasisymmetric functions [5].

A *dual equivalence* for $(\mathcal{A}, \text{Des})$ is a family of involutions on \mathcal{A} whose local equivalence classes are Schur functions and which commute when their indices are far away. For Littlewood–Richardson coefficients, these are Haiman’s original dual equivalence involutions [7]. A dual equivalence is compatible with a statistic when the involutions preserve the statistic. From this framework, we obtain an explicit set $\text{Dom} \subset \mathcal{A}$ such that

$$f(X; q) = \sum_{\lambda} \left(\sum_{\substack{S \in \text{Dom}(\mathcal{A}) \\ \alpha(S) = \lambda}} q^{\text{stat}(S)} \right) s_{\lambda}(X),$$

where α is an explicit map derived from Des that associates to each element of Dom a partition. This is the generalized notion of implicit rectification. For

example, in the Littlewood–Richardson case, the set Dom is precisely the set of skew tableaux that rectify to a particular standard tableau of straight shape. In particular, giving a dual equivalence for the data $(\mathcal{A}, \text{Des})$ that is compatible with the statistic proves that the generating function $f(X; q)$ is symmetric and Schur positive, and provides an explicit combinatorial formula for the Schur coefficients.

After reviewing symmetric functions and the associated tableaux combinatorics in Section 2, we review the dual equivalence relation on standard tableaux. In Section 3, we use this relation to define an edge-colored graph on standard tableaux. We axiomatize the graph, resulting in a new combinatorial method for establishing the Schur positivity of a quasisymmetric function. In Section 4, we reformulate this machinery in terms of involutions on a set and give a more explicit characterization of the Schur coefficients. As a first application of this theory, Section 5 defines involutions on k -tuples of tableaux that give a dual equivalence when $k = 2$, giving a surprisingly short proof of Schur positivity for the ribbon tableaux generating function introduced by Lascoux *et al.* [9] in the case of dominoes. While these involutions on tuples of tableaux are not, in general, a dual equivalence, we conjecture that the equivalence classes are always Schur positive. These larger equivalence classes are studied further in a forthcoming paper.

2. Preliminaries

2.1. Partitions and tableaux. We represent an integer *partition* λ by the decreasing sequence of its (nonzero) parts

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

We denote the size of λ by $|\lambda| = \sum_i \lambda_i$. If $|\lambda| = n$, we say that λ is a *partition of n* . The *Young diagram* of a partition λ is the set of points (i, j) in the $\mathbb{Z} \times \mathbb{Z}$ lattice such that $1 \leq i \leq \lambda_j$. We draw the diagram so that each point (i, j) is represented by the unit cell southwest of the point; see Figure 1.

For partitions λ, μ , we write $\mu \subset \lambda$ whenever the diagram of μ is contained within the diagram of λ ; equivalently $\mu_i \leq \lambda_i$ for all i . In this case, we define the *skew diagram* λ/μ to be the set-theoretic difference $\lambda - \mu$; for an example, see Figure 1. For our purposes, we depart from the norm by *not* identifying skew shapes that are translates of one another.

A *filling* of a (skew) diagram λ is a map $S : \lambda \rightarrow \mathbb{Z}_+$. A *semistandard Young tableau* is a filling that is weakly increasing along each row and strictly increasing along each column. A semistandard Young tableau is *standard* if it is a bijection



Figure 1. The Young diagram for (5, 4, 4, 1) and the skew diagram for (5, 4, 4, 1)/(3, 2, 2).

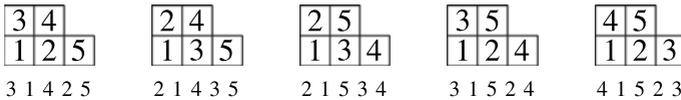


Figure 2. The standard Young tableaux of shape (3, 2) with their content reading words.

from λ to $[n]$, where $[n] = \{1, 2, \dots, n\}$. For λ a diagram of size n , define

$$\begin{aligned} \text{SSYT}(\lambda) &= \{\text{semi-standard tableaux } T : \lambda \rightarrow \mathbb{Z}_+\}, \\ \text{SYT}(\lambda) &= \{\text{standard tableaux } T : \lambda \xrightarrow{\sim} [n]\}. \end{aligned}$$

For $T \in \text{SSYT}(\lambda)$, we say that T has *shape* λ . If T contains entries $1^{\pi_1}, 2^{\pi_2}, \dots$ for some composition π , then we say that T has *weight* π . Thus $T \in \text{SYT}(\lambda)$ if and only if T has weight (1^n) .

The *content* of a cell of a diagram indexes the diagonal on which it occurs; that is, $c(x) = i - j$, when the cell x lies in position $(i, j) \in \mathbb{Z}_+ \times \mathbb{Z}_+$. The *content reading word* of a semistandard tableau is obtained by reading the entries in increasing order of content, going southwest to northeast along each diagonal (on which the content is constant). For examples, see Figure 2.

2.2. Symmetric functions. We have the familiar integral bases for Λ , the ring of symmetric functions, from [11], all indexed by partitions of n : the monomial symmetric functions m_λ , the elementary symmetric functions e_λ , the complete homogeneous symmetric functions h_λ , and, most importantly, the *Schur functions*, s_λ , which may be defined in several ways. For the purposes of this paper, we take the tableau approach:

$$s_\lambda(x) = \sum_{T \in \text{SSYT}(\lambda)} x^T, \tag{2.1}$$

where x^T is the monomial $x_1^{\pi_1} x_2^{\pi_2} \dots$ when T has weight π . This formula also defines the *skew Schur functions*, $s_{\lambda/\mu}$, by taking the sum over semistandard tableaux of shape λ/μ .

As we shall see in Section 3, it will often be useful to express a function in terms of Gessel's fundamental quasisymmetric functions [5] rather than monomials. For $\sigma \in \{\pm 1\}^{n-1}$, the *fundamental quasisymmetric function* $Q_\sigma(x)$ is defined by

$$Q_\sigma(x) = \sum_{\substack{i_1 \leq \dots \leq i_n \\ \sigma_{j=-1} \Rightarrow i_j < i_{j+1}}} x_{i_1} \cdots x_{i_n}. \quad (2.2)$$

We have indexed quasisymmetric functions by sequences of $+1$ and -1 , though, by setting $D(\sigma) = \{i \mid \sigma_i = -1\}$, we may change the indexing to the more familiar one of subsets of $[n - 1]$.

To connect quasisymmetric functions with Schur functions, for T a standard tableau on $[n]$ with content reading word w_T , define the *descent signature* $\sigma(T) \in \{\pm 1\}^{n-1}$ by

$$\sigma(T)_i = \begin{cases} +1 & \text{if } i \text{ appears to the left of } i + 1 \text{ in } w_T, \\ -1 & \text{if } i + 1 \text{ appears to the left of } i \text{ in } w_T. \end{cases} \quad (2.3)$$

For example, the descent signatures for the tableaux in Figure 2 are $+ - ++$, $- + -+$, $- + +-$, $+ - +-$, $++ -+$, from left to right. Note that, if we replace the content reading word with either the row or column reading word, the signature given by (2.3) remains unchanged.

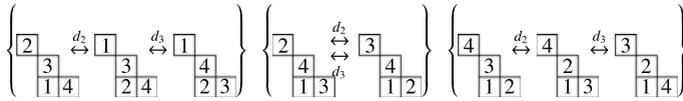
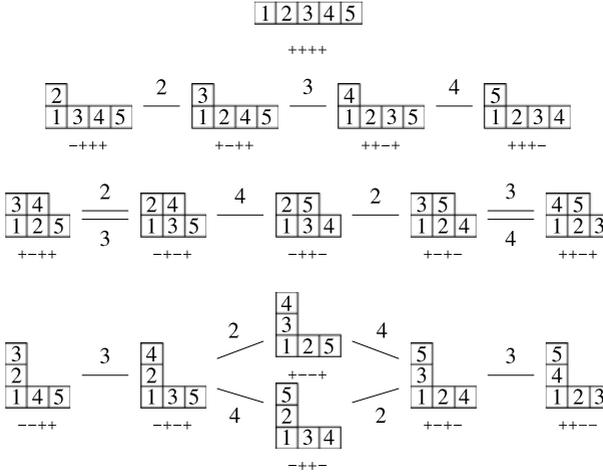
PROPOSITION 2.1 [5]. *The Schur function s_λ is expressed in terms of quasisymmetric functions by*

$$s_\lambda(x) = \sum_{T \in \text{SYT}(\lambda)} Q_{\sigma(T)}(x). \quad (2.4)$$

Comparing (2.1) with (2.4), using quasisymmetric functions instead of monomials allows us to work with standard tableaux rather than semistandard tableaux. One advantage of this formula is that, unlike (2.1), the right-hand side of (2.4) is finite. Continuing with the example in Figure 2,

$$s_{(3,2)}(x) = Q_{+-++}(x) + Q_{-+-+}(x) + Q_{-+--}(x) + Q_{+--+}(x) + Q_{++-+}(x).$$

2.3. Dual equivalence. Dual equivalence was first explicitly defined by Haiman [7] as a relation on tableaux dual to *jeu de taquin* equivalence under the


 Figure 5. Dual equivalence classes for $s_{(1)}s_{(2,1)}$.

 Figure 6. The standard dual equivalence graphs \mathcal{G}_5 , $\mathcal{G}_{4,1}$, $\mathcal{G}_{3,2}$, and $\mathcal{G}_{3,1,1}$.

respectively rectify to the tableaux in Figure 4. Taking generating functions, we see that the dual equivalence classes correspond precisely to the Schur expansion

$$s_{(1)}s_{(2,1)} = s_{(3,1)} + s_{(2,2)} + s_{(2,1,1)}.$$

3. Dual equivalence graphs

3.1. Axiomatization of dual equivalence. Construct a graph whose edges are colored on standard tableaux of partition shape from the dual equivalence relation in the following way. Whenever two standard tableaux T, U have content reading words that differ by an elementary dual equivalence for $i - 1, i, i + 1$, connect T and U with an edge colored by i . Associate to each tableau T the signature $\sigma(T)$ defined by (2.3). For an example, see Figure 6.

The connected components of the graph so constructed are the dual equivalence classes of standard tableaux. Let \mathcal{G}_λ denote the subgraph on tableaux of shape λ . Proposition 2.3 states that the \mathcal{G}_λ exactly give the connected components of the graph.

Define the generating function associated to \mathcal{G}_λ by

$$\sum_{v \in V(\mathcal{G}_\lambda)} Q_{\sigma(v)}(x) = s_\lambda(x). \tag{3.1}$$

By Proposition 2.1, this is Gessel’s quasisymmetric function expansion for a Schur function. In particular, the generating function of any graph whose vertices have signatures and whose connected components are isomorphic to the graphs \mathcal{G}_λ is Schur positive.

In this section, we characterize \mathcal{G}_λ in terms of edges and signatures so that we can readily identify those graphs that are isomorphic to some \mathcal{G}_λ .

DEFINITION 3.1. A signed colored graph of type (n, N) consists of the following data:

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

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

DEFINITION 3.2. A signed colored graph $\mathcal{G} = (V, \sigma, E)$ of type (n, N) is a dual equivalence graph of type (n, N) if $n \leq N$ and the following hold.

- (ax1) For $w \in V$ and $1 < i < n$, $\sigma(w)_{i-1} = -\sigma(w)_i$ if and only if there exists $x \in V$ such that $\{w, x\} \in E_i$. Moreover, x is unique when it exists.
- (ax2) For $\{w, x\} \in E_i$, $\sigma(w)_j = -\sigma(x)_j$ for $j = i - 1, i$, and $\sigma(w)_h = \sigma(x)_h$ for $h < i - 2$ and $h > i + 1$.
- (ax3) For $\{w, x\} \in E_i$, if $\sigma(w)_{i-2} = -\sigma(x)_{i-2}$, then $\sigma(w)_{i-2} = -\sigma(w)_{i-1}$, and if $\sigma(w)_{i+1} = -\sigma(x)_{i+1}$, then $\sigma(w)_{i+1} = -\sigma(w)_i$.
- (ax4) Every connected component of $(V, \sigma, E_{i-1} \cup E_i)$ appears in Figure 7 and every connected component of $(V, \sigma, E_{i-2} \cup E_{i-1} \cup E_i)$ appears in Figure 8.
- (ax5) If $\{w, x\} \in E_i$ and $\{x, y\} \in E_j$ for $|i - j| \geq 3$, then $\{w, v\} \in E_j$ and $\{v, y\} \in E_i$ for some $v \in V$.
- (ax6) Any two vertices of a connected component of $(V, \sigma, E_2 \cup \dots \cup E_i)$ may be connected by a path crossing at most one E_i edge.

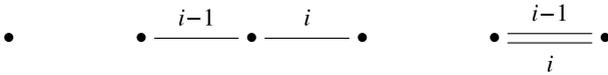


Figure 7. Allowed two-color connected components of a dual equivalence graph.

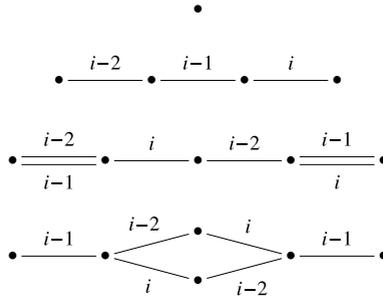


Figure 8. Allowed three-color connected components of a dual equivalence graph.

Note that, if $n > 4$, then the allowed structure for connected components of $(V, \sigma, E_{i-2} \cup E_{i-1} \cup E_i)$ dictates that every connected component of $(V, \sigma, E_{i-1} \cup E_i)$ appears in Figure 7.

Every connected component of a dual equivalence graph of type (n, N) is again a dual equivalence graph of type (n, N) .

It is often useful to consider a restricted set of edges of a signed colored graph. To be precise, for $m \leq n$ and $M \leq N$, the (m, M) -restriction of a signed colored graph \mathcal{G} of type (n, N) consists of the vertex set V , signature function $\sigma : V \rightarrow \{\pm 1\}^{M-1}$ obtained by truncating σ at $M - 1$, and the edge set $E_2 \cup \dots \cup E_{m-1}$. For $m \leq n, M \leq N$, the (m, M) -restriction of a dual equivalence graph of type (n, N) is a dual equivalence graph of type (m, M) .

The graph for $\mathcal{G}_{\lambda'}$ is obtained from \mathcal{G}_{λ} by conjugating each standard tableau and multiplying the signatures coordinatewise by -1 . Therefore the structure of $\mathcal{G}_{(2,1,1,1)}$, $\mathcal{G}_{(2,2,1)}$, and $\mathcal{G}_{(1,1,1,1,1)}$ is also indicated by Figure 6. Comparing this with Figure 8, axiom 4 stipulates that the components of a dual equivalence graph restricted to three consecutive edge colors are exactly the graphs for \mathcal{G}_{λ} when λ is a partition of 5.

PROPOSITION 3.3. *For λ a partition of n , \mathcal{G}_{λ} is a dual equivalence graph of type (n, n) .*

Proof. For $T \in \text{SYT}(\lambda)$, $\sigma(T)_{i-1} = -\sigma(T)_i$ if and only if i does not lie between $i - 1$ and $i + 1$ in the content reading word of T . In this case, there exists $U \in$

$\text{SYT}(\lambda)$ such that T and U differ by an elementary dual equivalence for $i - 1, i, i + 1$. Therefore U is obtained from T by swapping i with $i - 1$ or $i + 1$, whichever lies further away, with the result that $\sigma(T)_j = -\sigma(U)_j$ for $j = i - 1, i$ and also $\sigma(T)_h = \sigma(U)_h$ for $h < i - 2$ and $i + 1 < h$. This verifies axioms 1 and 2.

For axiom 3, if $\sigma(T)_{i-2} = -\sigma(U)_{i-2}$, then i and $i - 1$ have interchanged positions with $i - 2$ lying between, so that T and U also differ by an elementary dual equivalence for $i - 2, i - 1, i$, and similarly for $i + 1$. From this, we obtain an explicit description of double edges as those connecting vertices where $i - 2$ and $i + 1$ lie between $i - 1$ and i . Therefore axiom 4 becomes a straightforward finite check on permutations of 5. If $|i - j| \geq 3$, then $\{i - 1, i, i + 1\} \cap \{j - 1, j, j + 1\} = \emptyset$, so the elementary dual equivalences for $i - 1, i, i + 1$ and for $j - 1, j, j + 1$ commute, thereby demonstrating axiom 5.

Finally, for $T, U \in \text{SYT}(\lambda)$, $|\lambda| = i + 1$, we must show that there exists a path from T to U crossing at most one E_i edge. Let \mathcal{C}_T (respectively \mathcal{C}_U) denote the connected component of the (i, i) -restriction of \mathcal{G}_λ containing T (respectively U). Let μ (respectively ν) be the shape of T (respectively U) with the cell containing $i + 1$ removed. Then $\mathcal{C}_T \cong \mathcal{G}_\mu$ and $\mathcal{C}_U \cong \mathcal{G}_\nu$. If $\mu = \nu$, then, by Proposition 2.3, $\mathcal{C}_T = \mathcal{C}_U$, and axiom 6 holds. Assume, then, that $\mu \neq \nu$. Since $\mu, \nu \subset \lambda$ and $|\mu| = |\nu| = |\lambda| - 1$, both cells λ/μ and λ/ν must be northeastern corners of λ . Therefore there exists $T' \in \text{SYT}(\lambda)$ with i in position λ/ν , $i + 1$ in position λ/μ , and $i - 1$ between i and $i + 1$ in the content reading word of T' . Let U' be the result of swapping i and $i + 1$ in T' , in particular, $\{T', U'\} \in E_i$. By Proposition 2.3, T' is in \mathcal{C}_T and U' is in \mathcal{C}_U ; hence there exists a path from T to T' and a path from U' to U , each crossing only edges $E_h, h < i$. This establishes axiom 6. \square

DEFINITION 3.4. For partitions $\lambda \subset \rho$, with $|\lambda| = n$ and $|\rho| = N$, choose a tableau A of shape ρ/λ with entries $n + 1, \dots, N$. Define the set of standard Young tableaux of shape λ augmented by A , denoted $\text{ASYT}(\lambda, A)$, to be those $T \in \text{SYT}(\rho)$ such that T restricted to ρ/λ is A . Let $\mathcal{G}_{\lambda,A}$ be the signed colored graph of type (n, N) constructed on $\text{ASYT}(\lambda, A)$ with i -edges given by elementary dual equivalences for $i - 1, i, i + 1$ with $i < n$.

Note that $\mathcal{G}_{\lambda,A}$ is a dual equivalence graph of type (n, N) , and the (n, n) -restriction of $\mathcal{G}_{\lambda,A}$ is \mathcal{G}_λ .

Proposition 3.3 is the first step toward justifying Definition 3.2, and it also allows us to refer to \mathcal{G}_λ as the *standard dual equivalence graph corresponding to λ* . In order to show the converse, when two graphs satisfy axiom 1, as all graphs in this paper do, we define an isomorphism between them to be a sign-preserving bijection on vertex sets that respects color adjacency. Precisely, we have the following.

DEFINITION 3.5. A *morphism* between two signed colored graphs of type (n, N) satisfying dual equivalence graph axiom 1, say $\mathcal{G} = (V, \sigma, E)$ and $\mathcal{H} = (W, \tau, F)$, is a map $\phi : V \rightarrow W$ such that, for every $u, v \in V$,

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

A morphism is an *isomorphism* if it is a bijection on vertex sets.

LEMMA 3.6. *If ϕ is a morphism from a signed colored graph \mathcal{G} of type (n, N) satisfying axiom 1 to an augmented standard dual equivalence graph $\mathcal{G}_{\lambda, A}$, then ϕ is surjective.*

Proof. Suppose that $T = \phi(v)$ for some $T \in \text{ASYT}(\lambda, A)$ and some vertex v of \mathcal{G} . Then, for every $1 < i < n$, if $\{T, U\} \in E_i$, then, since $\sigma(v) = \sigma(T)$, by axiom 1 there exists a unique vertex w of \mathcal{G} such that $\{v, w\} \in E_i$ in \mathcal{G} . Since ϕ is a morphism, we must have $\{T, \phi(w)\} \in E_i$ in $\mathcal{G}_{\lambda, A}$. Thus, by the uniqueness condition of axiom 1, $\phi(w) = U$, and so U also lies in the image of ϕ . Therefore the i -neighbor of any vertex in the image of ϕ also lies in the image, since ϕ preserves i -edges. Since $\mathcal{G}_{\lambda, A}$ is connected, ϕ is surjective. \square

The final justification of this axiomatization is the following converse of Proposition 3.3.

THEOREM 3.7. *Every connected component of a dual equivalence graph of type (n, n) is isomorphic to \mathcal{G}_λ for a unique partition λ of n .*

The proof of Theorem 3.7 is the content of Section 3.2. We conclude this section by interpreting Theorem 3.7 in terms of symmetric functions.

COROLLARY 3.8. *Let \mathcal{G} be a dual equivalence graph of type (n, n) such that every vertex is assigned some additional statistic α that is constant on connected components of \mathcal{G} . Then*

$$\sum_{v \in V(\mathcal{G})} q^{\alpha(v)} Q_{\sigma(v)}(X) = \sum_{\lambda} \sum_{\mathcal{C} \cong \mathcal{G}_\lambda} q^{\alpha(\mathcal{C})} s_\lambda(X), \quad (3.2)$$

where the inner sum is over connected components of \mathcal{G} that are isomorphic to \mathcal{G}_λ . In particular, the generating function for \mathcal{G} so defined is symmetric and Schur positive.

3.2. The structure of dual equivalence graphs. We begin the proof of Theorem 3.7 by showing that the standard dual equivalence graphs are nonredundant in the sense that they are mutually nonisomorphic and have no nontrivial automorphisms. Both results stem from the observation that \mathcal{G}_λ contains a unique vertex such that the composition formed by the lengths of the runs of $+1$ in the signature gives a maximal partition.

PROPOSITION 3.9. *If $\phi : \mathcal{G}_\lambda \rightarrow \mathcal{G}_\mu$ is an isomorphism, then $\lambda = \mu$ and $\phi = \text{id}$.*

Proof. Let T_λ be the tableau obtained by filling the numbers 1 through n into the rows of λ from left to right, bottom to top, in which case $\sigma(T_\lambda) = +^{\lambda_1-1}, -, +^{\lambda_2-1}, -, \dots$. For any standard tableau T such that $\sigma(T) = \sigma(T_\lambda)$, the numbers 1 through λ_1 , and also $\lambda_1 + 1$ through $\lambda_1 + \lambda_2$, and so on, must form horizontal strips. In particular, if $\sigma(T) = \sigma(T_\lambda)$ for some T of shape μ , then $\lambda \leq \mu$, with equality if and only if $T = T_\lambda$.

Suppose that $\phi : \mathcal{G}_\lambda \rightarrow \mathcal{G}_\mu$ is an isomorphism. Let T_λ be as above for λ , and let T_μ be the corresponding tableau for μ . Then, since $\sigma(\phi(T_\lambda)) = \sigma(T_\lambda)$, $\lambda \leq \mu$. Conversely, since $\sigma(\phi^{-1}(T_\mu)) = \sigma(T_\mu)$, $\mu \leq \lambda$. Therefore $\lambda = \mu$. Furthermore, $\phi(T_\lambda) = T_\lambda$. For $T \in \text{SYT}(\lambda)$ such that $\{T_\lambda, T\} \in E_i$, we have $\{T_\lambda, \phi(T)\} \in E_i$, so $\phi(T) = T$ by dual equivalence axiom 1. Extending this, every tableau connected to a fixed point by some sequence of edges is also a fixed point for ϕ , and hence $\phi = \text{id}$ on each \mathcal{G}_λ by Proposition 2.3. \square

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

DEFINITION 3.10. Let $\mathcal{G} = (V, \sigma, E)$ be a signed colored graph of type (n, N) satisfying axiom 1. For $1 < i < N$, we say that a vertex $w \in V$ admits an i -neighbor if $\sigma(w)_{i-1} = -\sigma(w)_i$.

For $1 < i < n$, if $\sigma(w)_{i-1} = -\sigma(w)_i$ for some $w \in V$, then axiom 1 implies the existence of $x \in V$ such that $\{w, x\} \in E_i$. That is, if w admits an i -neighbor for some $1 < i < n$, then w has an i -neighbor in \mathcal{G} . For $n \leq i < N$, though i -edges do not exist in \mathcal{G} , if \mathcal{G} were the restriction of a graph of type $(i + 1, N)$ also satisfying axiom 1, then the condition $\sigma(w)_{i-1} = -\sigma(w)_i$ would imply the existence of a vertex x such that $\{w, x\} \in E_i$ in the type $(i + 1, N)$ graph. When convenient, E_i may be regarded as an involution, where $E_i(w) = x$ if w admits an i -neighbor and $\{w, x\} \in E_i$, and otherwise $E_i(w) = w$.

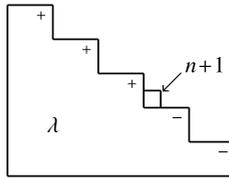


Figure 9. Identifying the unique position for $n + 1$ based on σ_n .

Recall the notion of augmenting a partition λ by a skew tableau A and the resulting dual equivalence graph $\mathcal{G}_{\lambda,A}$ from Definition 3.4.

LEMMA 3.11. *Let $\mathcal{G} = (V, \sigma, E)$ be a connected dual equivalence graph of type (n, N) with $n < N$, and let ϕ be an isomorphism from the (n, n) -restriction of \mathcal{G} to \mathcal{G}_λ for some partition λ of n . Then there exists a semistandard tableau A of shape ρ/λ , $|\rho| = N$, with entries $n + 1, \dots, N$ such that ϕ gives an isomorphism from \mathcal{G} to $\mathcal{G}_{\lambda,A}$. Moreover, the position of the cell of A containing $n + 1$ is unique.*

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

By dual equivalence axiom 2, σ_n is constant on connected components of the $(n - 1, n + 1)$ -restriction of \mathcal{G} . By Proposition 2.3, a connected component of the $(n - 1, n + 1)$ -restriction of \mathcal{G}_λ consists of all standard Young tableaux where n lies in a particular northeastern cell of λ . Therefore, for each connected component of the $(n - 1, n + 1)$ -restriction of \mathcal{G} , we may identify its image under ϕ with \mathcal{G}_μ for some partition $\mu \subset \lambda$, $|\mu| = n - 1$, with n lying in position λ/μ . We will show that σ_n has the monotonicity property on connected components of the $(n - 1, n + 1)$ -restriction of \mathcal{G} depicted in Figure 9; that is, there is an inner corner above which $\sigma_n = +1$ and below which $\sigma_n = -1$.

Let \mathcal{C} and \mathcal{D} be two distinct connected components of the $(n - 1, n + 1)$ -restriction of \mathcal{G} such that there exist vertices u of \mathcal{C} and v of \mathcal{D} with $\{u, v\} \in E_{n-1}$. Let $\phi(\mathcal{C}) \cong \mathcal{G}_\mu$, and let $\phi(\mathcal{D}) \cong \mathcal{G}_\nu$. Since $\{u, v\} \in E_{n-1}$, $\phi(v)$ must have $n - 1$ in position λ/μ with $n - 2$ lying between $n - 1$ and n in the content reading word. Since ϕ preserves E_{n-1} edges, $\phi(u)$ must be the result of an elementary dual equivalence on $\phi(v)$ for $n - 2, n - 1, n$, which will necessarily interchange $n - 1$ and n . Since ϕ preserves signatures, λ/ν lies northwest of the position of λ/μ if and

only if $\sigma(v)_{n-2,n-1} = +-$ and $\sigma(u)_{n-2,n-1} = -+$. If λ/ν lies northwest of the position of λ/μ and $\sigma(v)_n = -1$, then $\sigma(v)_n = \sigma(v)_{n-1}$. Thus, by axiom 3, $\sigma(u)_n = \sigma(v)_n = -1$. Similarly, if λ/ν lies northwest of the position of λ/μ and $\sigma(u)_n = +1$, then $\sigma(u)_n = \sigma(u)_{n-1}$. Thus, by axiom 3, $\sigma(v)_n = \sigma(u)_n = +1$.

Abusing notation and terminology, we identify $\phi(\mathcal{C})$ with the cell of λ in position λ/μ , where $\phi(\mathcal{C}) \cong \mathcal{G}_\mu$. With this convention, we have shown that, if $\sigma_n(\mathcal{C}) = +1$ and \mathcal{D} is any component connected to \mathcal{C} by an $n - 1$ -edge such that $\phi(\mathcal{D})$ lies northwest of $\phi(\mathcal{C})$, then $\sigma_n(\mathcal{D}) = +1$ as well. Similarly, if $\sigma_n(\mathcal{C}) = -1$ and \mathcal{D} is any component connected to \mathcal{C} by an $n - 1$ -edge such that $\phi(\mathcal{D})$ lies southeast of $\phi(\mathcal{C})$, then $\sigma_n(\mathcal{D}) = -1$ as well. By dual equivalence graph axiom 6, for any two distinct connected components \mathcal{C} and \mathcal{D} of the $(n - 1, n + 1)$ -restriction of \mathcal{G} and any pair of vertices w on \mathcal{C} and x on \mathcal{D} , there is a path from w to x crossing at most one, and hence exactly one, $n - 1$ edge. Therefore, for any \mathcal{C} and \mathcal{D} , there exist vertices v of \mathcal{C} and u of \mathcal{D} such that $\{v, u\} \in E_{n-1}$. Hence every two connected components of the $(n - 1, n + 1)$ -restriction of \mathcal{G} are connected by an $n - 1$ -edge, thus establishing the monotonicity depicted in Figure 9.

This established, it follows that there exists a unique row such that $\sigma(\mathcal{C})_n = -1$ whenever the $\phi(\mathcal{C})$ has n south of this row and $\sigma(\mathcal{C})_n = +1$ whenever the $\phi(\mathcal{C})$ has n north of this row. In this case, the cell containing $n + 1$ must be placed at the eastern end of this pivotal row, and doing so extends ϕ to an isomorphism between $(n, n + 1)$ graphs. \square

Once Theorem 3.7 has been proved, Lemma 3.11 may be used to obtain the following generalization of Theorem 3.7 for dual equivalence graphs of type (n, N) .

COROLLARY 3.12. *Every connected component of a dual equivalence graph of type (n, N) is isomorphic to $\mathcal{G}_{\lambda,A}$ for a unique partition λ and some skew tableau A of shape ρ/λ , $|\rho| = N$, with entries $n + 1, \dots, N$.*

Finally we have all of the ingredients necessary to prove the main result of this section.

THEOREM 3.13. *Let \mathcal{G} be a connected signed colored graph of type $(n + 1, n + 1)$ satisfying axioms 1 through 5 such that each connected component of the (n, n) -restriction of \mathcal{G} is isomorphic to a standard dual equivalence graph. Then there exists a morphism ϕ from \mathcal{G} to \mathcal{G}_λ for some unique partition λ of $n + 1$.*

Proof. When $n + 1 = 2$ or, more generally, when \mathcal{G} has no n -edges, the result follows immediately from Lemma 3.11. Therefore we proceed by induction,

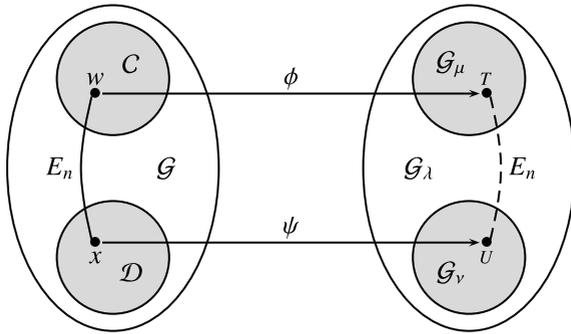


Figure 10. An illustration of the gluing process.

assuming that \mathcal{G} has at least one n -edge, and assuming the result for graphs of type (n, n) .

By induction, for every connected component \mathcal{C} of the $(n, n + 1)$ -restriction of \mathcal{G} , we have an isomorphism from the (n, n) -restriction of \mathcal{C} to \mathcal{G}_μ for a unique partition μ of n . By Lemma 3.11, this isomorphism extends to an isomorphism from \mathcal{C} to $\mathcal{G}_{\mu, A}$ for a unique augmenting tableau A , say with shape λ/μ . We will show that for any \mathcal{C} the shape of μ augmented with A is the same, and that we may glue these isomorphisms together to obtain a morphism from \mathcal{G} to \mathcal{G}_λ .

Suppose that $\{w, x\} \in E_n$. Let \mathcal{C} (respectively \mathcal{D}) denote the connected component of the $(n, n + 1)$ -restriction of \mathcal{G} containing w (respectively x). Let ϕ (respectively ψ) be the isomorphism from \mathcal{C} (respectively \mathcal{D}) to $\mathcal{G}_{\mu, A}$ (respectively $\mathcal{G}_{\nu, B}$), and set $T = \phi(w)$; see Figure 10. We will show that $\psi(x) = E_n(T)$, and hence, if μ, A has shape λ , then so does ν, B , and the maps ϕ and ψ glue together to give an morphism from $\mathcal{C} \cup \mathcal{D}$ to \mathcal{G}_λ that preserves n -edges. There are two cases to consider, based on the relative positions of $n - 1, n, n + 1$ in T , regarded as a tableau of shape λ .

First suppose that $n + 1$ lies between n and $n - 1$ in the reading word of T . We will show that, in this case, $\mathcal{C} = \mathcal{D}$. Since $n + 1$ lies between n and $n - 1$ in the reading word of T , both $n - 1$ and n must be northeastern corners of μ , and so there is a cell with entry less than $n - 1$ that also lies between them. By Proposition 2.3, there exists a tableau T' with $n - 1, n, n + 1$ in the same positions as in T , but now with $n - 2$ lying between n and $n - 1$ in the reading word of T' . Furthermore, since both T and T' lie on the $(n - 2, n + 1)$ -restriction of $\mathcal{G}_{\mu, A}$, there is a path from T to T' in $\mathcal{G}_{\mu, A}$ using only edges E_h with $h \leq n - 3$. Let $U' = E_n(T')$. Since $n - 2$ lies between n and $n - 1$ in U' , we have $U' = E_{n-1}(T')$ as well. By axiom 5, all edges in the path from T to T' commute with E_n , and so the same path takes $U = E_n(T)$ to U' , and each pair of corresponding tableaux

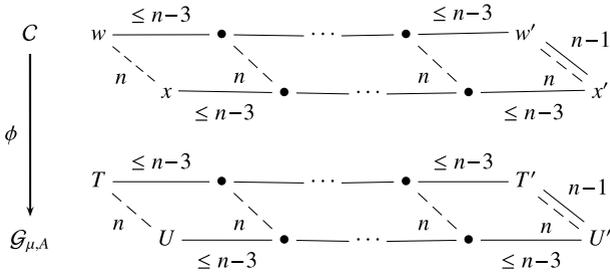


Figure 11. An illustration of the path from T to U in $\mathcal{G}_{\mu,A}$ and its lift in \mathcal{C} .

on the two paths is connected by an E_n edge; see Figure 11. We now embark on a diagram chase around Figure 11, from w to x and from T to U , to show that $\mathcal{C} = \mathcal{D}$.

Since the path from T to T' to U' to U uses only edges from $\mathcal{G}_{\mu,A}$, this path lifts via the isomorphism ϕ to a path in \mathcal{C} . Let $w' = \phi^{-1}(T')$ and $x' = \phi^{-1}(U')$. We will show that $x = \phi^{-1}(U)$, and so lies on \mathcal{C} . Since ϕ preserves signatures, both w' and x' must admit an n -edge in \mathcal{G} . As summarized in Figure 7, axioms 3 and 4 dictate that the only way for two vertices connected by an $n - 1$ -edge both to admit an n -edge is for $\{w', x'\} \in E_n$ in \mathcal{G} . By axioms 2 and 5, the path from w' to w gives an identical path from x' to $\phi^{-1}(U)$. Since each corresponding pair along the two paths must be paired by an n -edge, we have $\phi^{-1}(U) = E_n(w) = x$, as desired. Therefore x lies on \mathcal{C} , and ϕ respects the n -edge between w and x . In this case $\mathcal{C} = \mathcal{D}$ and, by Proposition 3.9, $\psi = \phi$.

For the second case, suppose that $n - 1$ lies between n and $n + 1$ in T . Consider the subset of tableaux in $\mathcal{G}_{\mu,A}$ with n and $n + 1$ fixed in the same position as in T and $n - 1$ lying anywhere between them in the reading word. We claim that this set uniquely determines μ and A . In terms of the graph structure, these are all tableaux reachable from T using edges E_h with $h \leq n - 3$ and a certain subset of the E_{n-2} edges. We will return soon to the question of which E_{n-2} edges these are. For now, let \mathcal{T} denote the union of the graphs $\mathcal{G}_{\rho,R}$, where ρ is a partition of $n - 2$ with augmenting tableau R consisting of a single cell containing $n - 1$ such that ρ, R is the shape of T with n and $n + 1$ removed and the augmented cell of R lies strictly between the positions of n and $n + 1$ in T . Clearly the set of ρ, R uniquely determines the cells containing n and $n + 1$, and so uniquely determines λ . Furthermore, which of $n, n + 1$ occupies which cell is determined by σ_n , and this is constant on this subset by axioms 2 and 3. Lifting \mathcal{T} to \mathcal{C} using ϕ^{-1} gives rise to an induced subgraph of \mathcal{C} that completely determines λ as well as the positions of n and $n + 1$ in the image of this subgraph under ϕ . We will

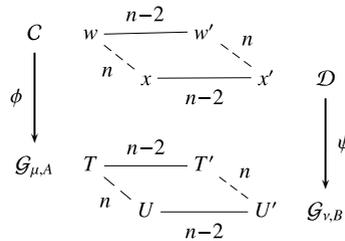


Figure 12. An illustration of E_{n-2} edges on $\mathcal{T} \cup \mathcal{U}$ and their lifts in $\mathcal{C} \cup \mathcal{D}$.

show that the corresponding induced subgraph for \mathcal{D} is isomorphic but with the opposite sign for σ_n .

To prove the assertion, we return to the question of which E_{n-2} edges are allowed in generating \mathcal{T} . Any E_{n-2} edge that keeps $n - 1$ between n and $n + 1$ clearly does not change σ_{n-1} or σ_n . Therefore such E_{n-2} edges must pair vertices both of which admit an n -neighbor. Further, neither of these vertices may have E_n as a double edge with E_{n-1} since $n - 1$ lies between n and $n + 1$. By axiom 4, the E_{n-2} edges that meet these conditions are precisely those in the lower component of Figure 8. In particular, these E_{n-2} edges commute with E_n edges as depicted in Figure 12. By axiom 5, E_h also commutes with E_n for $h \leq n - 3$. Therefore all edges on the induced subgraph of \mathcal{C} containing $\phi^{-1}(\mathcal{T})$ commute with E_n . Therefore E_n may be regarded as an isomorphism from this subgraph to $\mathcal{X} = E_n(\phi^{-1}(\mathcal{T}))$. Since $\{w, x\} \in E_n$ and $w \in \phi^{-1}(\mathcal{T})$, we have $x \in \mathcal{X}$. Since all edges of the induced subgraph have color at most $n - 2$, it follows that $\mathcal{X} \subset \mathcal{D}$.

Let $U = \psi(x)$, and let $\mathcal{U} = \psi(\mathcal{X})$. Since ϕ , ψ and E_n are isomorphisms of the $(n - 2, n - 2)$ -restrictions, \mathcal{U} together with its induced edges is isomorphic to \mathcal{T} together with its induced edges, though, by axiom 1, the signs for σ_n and σ_{n+1} are reversed. By the definition of \mathcal{T} and the fact that it uniquely determines μ and A , this implies that the tableaux in \mathcal{U} have shape λ , with the cells containing n and $n + 1$ reversed from that in \mathcal{T} . In particular, $\mathcal{U} = E_n(\mathcal{T})$; that is to say, ϕ and ψ glue to give a morphism from $\mathcal{C} \cup \mathcal{D} \subset \mathcal{G}$ to $\mathcal{G}_{\mu,A} \cup \mathcal{G}_{\nu,B} \subset \mathcal{G}_\lambda$ that respects E_n edges of the induced subgraphs.

Since T admits an n -neighbor, n cannot lie between $n - 1$ and $n + 1$, so these two are the only cases. Thus we now have a well-defined morphism from the $(n, n + 1)$ -restriction of \mathcal{G} to the $(n, n + 1)$ -restriction of \mathcal{G}_λ that respects n -edges. As such, this map lifts to a morphism from \mathcal{G} to \mathcal{G}_λ . \square

By Lemma 3.6, the morphism of Theorem 3.13 is necessarily surjective, though in general it need not be injective. The smallest example where injectivity fails

was first observed by Gregg Musiker in a graph of type $(6, 6)$ with generating function $2s_{(3,2,1)}(X)$.

For any \mathcal{G} satisfying the hypotheses of Theorem 3.13, the fiber over each vertex of \mathcal{G}_λ in the morphism from \mathcal{G} to \mathcal{G}_λ has the same cardinality. Letting ϕ be the morphism from \mathcal{G} to \mathcal{G}_λ , there is a bijective correspondence between connected components of $\phi^{-1}(\mathcal{G}_\mu)$ and connected components of $\phi^{-1}(\mathcal{G}_\nu)$ that follows from the following result.

COROLLARY 3.14. *Let \mathcal{G} satisfy the hypotheses of Theorem 3.13, and let ϕ be the morphism from \mathcal{G} to \mathcal{G}_λ . For any connected component \mathcal{C} of the (n, n) -restriction of \mathcal{G} , say with $\phi(\mathcal{C}) = \mathcal{G}_\mu$, and any partition $\nu \subset \lambda$ of size n , there is a unique connected component \mathcal{D} of the (n, n) -restriction of \mathcal{G} with $\phi(\mathcal{D}) = \mathcal{G}_\nu$ that can be reached from \mathcal{C} by crossing at most one E_n edge.*

Proof. To prove existence, if $\nu \neq \mu$, let T be a tableau of shape λ with $n + 1$ in position λ/μ , n in position λ/ν , and $n - 1$ lying between in the reading word. Otherwise let T be a tableau with $n + 1$ in position λ/μ and n and $n - 1$ lying on opposite sides in the reading word. Let w be the unique element in $\phi^{-1}(T) \cap \mathcal{C}$. Then w admits an n -neighbor, and, since ϕ is a morphism, $\phi(E_n(w)) = E_n(\phi(w)) \in \mathcal{G}_\nu$.

To prove uniqueness, let $\{w, x\} \in E_n$ with $w \in \mathcal{C} \cong \mathcal{G}_\mu$ and $x \in \mathcal{D} \cong \mathcal{G}_\nu$. If $n + 1$ lies between n and $n - 1$ in $\phi(w)$, then $\mu = \nu$, and, just as in the proof of Theorem 3.13, we conclude that $\mathcal{D} = \mathcal{C}$, as desired. Alternately, assume that $n - 1$ lies between n and $n + 1$ in $\phi(w)$, and suppose that $\{w', x'\} \in E_{n-1}$ with $w' \in \mathcal{C}$ and $x' \in \mathcal{D}' \cong \mathcal{G}_\nu$. Since $\phi(w)$ and $\phi(w')$ have the same shape, and $E_n(\phi(w)) = \phi(E_n(w)) = \phi(x)$ and $E_n(\phi(w')) = \phi(E_n(w')) = \phi(x')$ have the same shape, just as in the proof of Theorem 3.13, there must be a path from $\phi(w)$ to $\phi(w')$ in \mathcal{G}_ν using only edges E_h with $h \leq n - 3$ and those E_{n-2} that commute with E_n . Therefore this path gives rise to the same path from $\phi(x)$ to $\phi(x')$ in \mathcal{G}_μ . The former path lifts to a path from w to w' in \mathcal{C} , and so the latter lifts to a path from $E_n(w) = x$ to $E_n(w') = x'$ in $\mathcal{D} = \mathcal{D}'$, which is as desired. \square

In order to ensure that the morphism in the conclusion of Theorem 3.13 is an isomorphism, and thereby complete the proof of Theorem 3.7, we need only invoke the heretofore uninvoked axiom 6.

Proof of Theorem 3.7. Let \mathcal{G} be a dual equivalence graph of type $(n + 1, n + 1)$. We aim to show that \mathcal{G} is isomorphic to \mathcal{G}_λ for a unique partition λ of $n + 1$. We proceed by induction on $n + 1$, noting that the result is trivial for $n + 1 = 2$. Every connected component of the (n, n) -restriction of \mathcal{G} is a dual equivalence graph, and so, by induction, is isomorphic to a standard dual equivalence graph.

Thus, by Theorem 3.13, there exists a morphism, say ϕ , from \mathcal{G} to \mathcal{G}_λ for a unique partition λ of $n + 1$. By Corollary 3.14, for any connected component \mathcal{C} of the (n, n) -restriction of \mathcal{G} and any partition $\nu \subset \lambda$ of size n , there is a unique connected component \mathcal{D} of the (n, n) -restriction of \mathcal{G} that can be reached from \mathcal{C} by crossing at most one E_n edge such that $\phi(\mathcal{D}) = \mathcal{G}_\nu$. By dual equivalence axiom 6, any two connected components of the (n, n) -restriction of \mathcal{G} can be connected by a path using at most one E_n edge. Therefore the connected components of the (n, n) -restriction of \mathcal{G} are pairwise nonisomorphic. Hence the morphism from \mathcal{G} to \mathcal{G}_λ is injective on the $(n, n + 1)$ -restrictions, and so it is injective on all of \mathcal{G} . Surjectivity follows from Lemma 3.6; thus ϕ is an isomorphism. \square

4. Abstract dual equivalence

4.1. A recharacterization of dual equivalence. The inspiration for dual equivalence graphs comes from the elementary dual equivalence involutions. In this section, we reformulate the machinery of dual equivalence graphs back in terms of involutions. Given a set \mathcal{A} of combinatorial objects together with a notion of a descent set Des sending an object to a subset of positive integers, the goal is to define equivalence relations on objects in \mathcal{A} so that the sum over objects in any single equivalence class is a single Schur function.

Given $(\mathcal{A}, \text{Des})$ and involutions $\varphi_2, \dots, \varphi_{n-1}$ on \mathcal{A} , for $1 < h < i < n$ we consider the *restricted dual equivalence class* $[T]_{(h-1, i+1)}$ generated by $\varphi_h, \dots, \varphi_i$. In addition, we consider the *restricted and shifted descent set* $\text{Des}_{(h-1, i+1)}(T)$ obtained by intersecting $\text{Des}(T)$ with $\{h - 1, \dots, i\}$ and subtracting $h - 2$ from each element so that $\text{Des}_{(h-1, i+1)}(T) \subseteq [i - h + 2]$.

DEFINITION 4.1. Let \mathcal{A} be a finite set, and let Des be a map on \mathcal{A} such that $\text{Des}(T) \subseteq [n - 1]$ for all $T \in \mathcal{A}$. A *dual equivalence* for $(\mathcal{A}, \text{Des})$ is a family of involutions $\{\varphi_i\}_{1 < i < n}$ on \mathcal{A} such that the following hold.

- (i) For all $i - h \leq 3$ and all $T \in \mathcal{A}$, there exists a partition λ of $i - h + 3$ such that

$$\sum_{U \in [T]_{(h, i)}} \mathcal{Q}_{\text{Des}_{(h, i)}(U)}(X) = s_\lambda(X).$$

- (ii) For all $|i - j| \geq 3$ and all $T \in \mathcal{A}$, we have

$$\varphi_j \varphi_i(T) = \varphi_i \varphi_j(T).$$

On the surface, Definition 4.1 appears to have only two conditions. However, condition (i) includes four cases, for $i - h = 0, 1, 2, 3$. These cases correspond

to the dual equivalence graph axioms as follows: $i - h = 0$ means there is one edge color, and this corresponds to axioms 1 and 2; $i - h = 1$ means there are two edge colors, and this corresponds to axiom 3 and the first half of axiom 4; $i - h = 2$ means there are three edge colors, and this corresponds to the second half of axiom 4; $i - h = 3$ means there are four edge colors, and this corresponds to the axiom 6.

THEOREM 4.2. *For $\{\varphi_i\}$ a family of involutions on \mathcal{A} , let $\mathcal{G} = (\mathcal{A}, \sigma, \Phi)$ be the corresponding signed colored graph with descent signature and i -colored edges given by*

$$\sigma(T)_i = -1 \Leftrightarrow i \in \text{Des}(T) \quad \Phi_i = \{ \{T, \varphi_i(T)\} \mid T \notin \mathcal{A}^{\varphi_i} \}. \quad (4.1)$$

Then $\{\varphi_i\}$ is a dual equivalence for $(\mathcal{A}, \text{Des})$ if and only if \mathcal{G} is a dual equivalence graph.

Proof. We may assume that \mathcal{G} is connected. If \mathcal{G} is a dual equivalence graph, then, by Theorem 3.7, we may assume that $\mathcal{G} = \mathcal{G}_\lambda$ for some partition λ , and the result follows from the observation that d_i is a dual equivalence for SYT(λ) [7].

Assume then that $\{\varphi_i\}$ is a dual equivalence for $(\mathcal{A}, \text{Des})$. If $\sigma(w)_{i-1} = \sigma(w)_i$, then the degree-3 generating function of $[w]_i$ is $s_{(3)}$ or $s_{(1^3)}$, so w must be in an i -equivalence class of its own; that is, $\varphi_i(w) = w$. On the other hand, if $\sigma(w)_{i-1} = -\sigma(w)_i$, then the equivalence class of w must have generating function $s_{(2,1)}$, so there exists $x \neq w$ such that $\varphi_i(w) = x$. Moreover, since φ_i is an involution, x is unique. This establishes axiom 1. This also shows that, up to interchanging w and x , $\sigma(w)_{i-1,i} = +-$ and $\sigma(x)_{i-1,i} = -+$, satisfying the first half of axiom 2.

Assume that $x = \varphi_i(w) \neq w$, $\sigma(w)_{i-2} = -\sigma(x)_{i-2}$ and, contrary to axiom 3, $\sigma(w)_{i-2} = \sigma(w)_{i-1}$. By axiom 1 and the first half of axiom 2, we may assume that $\sigma(w)_{i-1,i} = +-$ and $\sigma(x)_{i-1,i} = -+$. Therefore $\sigma(w)_{i-2,i-1,i} = + + -$ and $\sigma(x)_{i-2,i-1,i} = - - +$. By restricting to the $(i-1, i)$ -equivalence class, condition (i) implies that w must lie in a class with generating function $s_{(3,1)}$, and x must lie in a class with generating function $s_{(2,1,1)}$, contradicting the fact that w, x are in the same restricted class with generating function a single Schur function. Therefore axiom 3 must hold.

Now consider nontrivial connected components of $\Phi_{i-1} \cup \Phi_i$. Up to conjugating, the restricted generating function is either $s_{(3,1)}$ or $s_{(2,2)}$. In the latter case, there are two vertices, and neither can be a fixed point for φ_{i-1} or φ_i , by axiom 1. Therefore the structure indeed matches the right graph of Figure 7. In the former case, there are three vertices, say $\sigma(w)_{i-2,i-1,i} = + + -$, $\sigma(x)_{i-2,i-1,i} = + - +$, and $\sigma(u)_{i-2,i-1,i} = - + +$. By axiom 1 and the first half

of axiom 2, we must have $\varphi_i(w) = x$, $\varphi_{i-1}(u) = x$, with w fixed for φ_{i-1} and u fixed for φ_i . Again, the structure indeed matches the middle graph of Figure 7.

Next we consider nontrivial connected components of $\Phi_{i-1} \cup \Phi_i \cup \Phi_{i+1}$. Up to conjugating, the restricted generating function is either $s_{(4,1)}$, $s_{(3,2)}$, or $s_{(3,1,1)}$. In the former case, by axiom 1, exactly two vertices are not fixed points for φ_j for each $j = i - 1, i, i + 1$, and making the forced pairing results in the structure of $\mathcal{G}_{(4,1)}$. In the other cases, there are four vertices that are not fixed points for φ_j for each $j = i - 1, i, i + 1$, so, taking the first half of axiom 2 into account, there are two possible j -edge pairings for each $j = i - 1, i, i + 1$. For generating function $s_{(3,1,1)}$, taking axiom 3 into account uniquely forces each edge pairing so that the resulting graph has the structure of $\mathcal{G}_{(3,1,1)}$. For generating function $s_{(3,2)}$, given that two color components must appear in Figure 7, there are two possibilities: the structure of $\mathcal{G}_{(3,2)}$ or a vertex w with $\varphi_{i-1}(w) = \varphi_i(w) = \varphi_{i+1}(w)$, in which case the restricted equivalence class has only two elements and not the required five. Therefore axiom 4 holds.

Axiom 5 is precisely condition (ii) for dual equivalence.

For the second half of axiom 2, assume that $|i - j| \geq 3$. If w is a fixed point for φ_i , then, by axiom 5, $\varphi_i(\varphi_j(w)) = \varphi_j(\varphi_i(w)) = \varphi_j(w)$, so $\varphi_j(w)$ is a fixed point for φ_i as well. In particular, if $x = \varphi_i(w) \neq w$ and $\sigma(w)_j \neq \sigma(x)_j$ for some $j < i - 2$, then, by this observation and axiom 1, $\sigma(w)_j \neq \sigma(x)_j$ for every $j < i - 2$. Similarly, if $x = \varphi_i(w) \neq w$ and $\sigma(w)_j \neq \sigma(x)_j$ for some $j > i + 1$, then $\sigma(w)_j \neq \sigma(x)_j$ for every $j > i + 1$. Therefore it suffices to show that $\sigma(w)_j \neq \sigma(x)_j$ for $j = i - 3$ and $j = i + 2$. This follows by axioms 1 and 4 and the first half of axiom 2.

Finally, since axiom 6 follows from axiom 4 when $n \leq 5$, we prove axiom 6 by induction, and assume that $i = n - 1$. Since \mathcal{G} has been shown to satisfy axioms 1 through 5 and the $(n - 1, n - 1)$ -restriction satisfies axiom 6 by induction, Theorem 3.7 ensures that the hypotheses of Theorem 3.13 and Corollary 3.14 are met. Therefore the generating function of \mathcal{G} is ks_λ for some positive integer k . When $n = 6$, condition (i) ensures that $k = 1$ and the map to \mathcal{G}_λ is an isomorphism. In particular, we may assume that the restriction of \mathcal{G} to edges $\Phi_{n-4}, \dots, \Phi_{n-1}$ satisfies dual equivalence graph axiom 6.

Suppose that $T, U, V, W \in \mathcal{A}$, and that $\{T, U\}, \{V, W\} \in \Phi_{n-1}$ and U and V lie on the same connected component of $\Phi_2 \cup \dots \cup \Phi_{n-2}$. We will show that there exist T', W' lying on the same connected component of $\Phi_2 \cup \dots \cup \Phi_{n-2}$ as T, W , respectively, such that there is a path from T' to W' crossing at most one Φ_{n-1} edge; see Figure 13. This is enough to establish axiom 6.

By Theorem 3.7, every connected dual equivalence graph is isomorphic to the graph on tableaux of a given shape. By the inductive hypothesis and Theorem 3.13, we may identify T, U, V, W with tableaux of shape κ , $|\kappa| = n$,

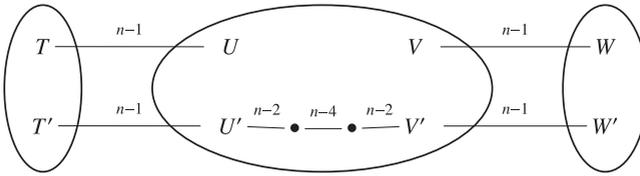


Figure 13. Establishing axiom 6 from local conditions.

and, when restricted to entries less than n , T, U, V, W have shapes $\mu, \lambda, \lambda, \nu$, respectively, with μ, λ, ν distinct partitions contained in κ by Corollary 3.14. Then $\kappa \setminus \xi$ must be a corner (end of row, top of column) for $\xi = \mu, \lambda, \nu$. There are six cases to consider, based on the relative positions of $\kappa \setminus \mu, \kappa \setminus \nu$, and $\kappa \setminus \lambda$ in κ . We treat one case in full detail, noting that the others can be resolved in a completely analogous way.

Assume that these cells appear with $\kappa \setminus \mu$ northwest of $\kappa \setminus \nu$ northwest of $\kappa \setminus \lambda$. Let T' be any tableau of shape κ with n in position $\kappa \setminus \mu$, $n - 1$ in position $\kappa \setminus \lambda$, $n - 2$ in position $\kappa \setminus \nu$, $n - 3$ between n and $n - 2$ in the reading order, $n - 4$ between $n - 1$ and $n - 2$ in the reading order, and $n - 5$ between $n - 3$ and $n - 4$ in the reading order. Set $U' = \varphi_{n-1}(T')$. Since the shape of U' restricted to entries less than n is λ , U' and $U = \varphi_{n-1}(T)$ must lie on the same connected component of $\Phi_2 \cup \dots \cup \Phi_{n-2}$, as must $V' = \varphi_{n-2}\varphi_{n-4}\varphi_{n-2}(U')$. Set $W' = \varphi_{n-1}(V')$. Since the shape of W' restricted to entries less than n is ν , W' and $W = \varphi_{n-1}(V)$ must lie on the same connected component of $\Phi_2 \cup \dots \cup \Phi_{n-2}$. Moreover, by unraveling the definitions of U', V' , and W' , we have $W' = \varphi_{n-1}\varphi_{n-2}\varphi_{n-4}\varphi_{n-2}\varphi_{n-1}(T')$. Therefore T' and W' lie on the same restricted graph of size 6, so by induction there exists a path from T' to W' using at most one edge in Φ_{n-1} . \square

REMARK 4.3. In practice, the characterization of dual equivalence being local makes it far better than the axioms for a dual equivalence graph for establishing Schur positivity. The equivalence of axiom 6 to a local condition was first observed and proved by Roberts [12].

COROLLARY 4.4. *If there exists a dual equivalence for $(\mathcal{A}, \text{Des})$, then*

$$\sum_{T \in \mathcal{A}} Q_{\text{Des}(T)}(X)$$

is symmetric and Schur positive.

4.2. Formulas from dual equivalence. Dual equivalence may be used to prove that a function is symmetric and Schur positive, and it gives a combinatorial interpretation of the Schur coefficients as the number of equivalence classes or connected components of a certain type. A classical example of a similar formula is the standard tableau version of the Littlewood–Richardson rule.

THEOREM 4.5 (Littlewood–Richardson rule). *Define integers $c_{\mu,v}^\lambda$ by*

$$s_\mu s_\nu = \sum_{\lambda} c_{\mu,\nu}^\lambda s_\lambda. \quad (4.2)$$

Then $c_{\mu,v}^\lambda$ is the number of standard tableaux of shape μ appended to ν that rectify by jeu de taquin to a chosen standard Young tableau of shape λ .

We use dual equivalence to give a simple proof of this rule as a corollary (at $q = 1$) to Theorem 5.4.

Dual equivalence for \mathcal{A} can be regarded as implicitly giving a rectification rule via the unique isomorphism, say θ , from $(\mathcal{A}, \sigma, \Phi)$ to \mathcal{G}_λ . That is, say that $T \in \mathcal{A}$ rectifies to $\theta(T) \in \text{SYT}(\lambda)$. Then choosing only those tableaux that rectify to a fixed $T \in \text{SYT}(\lambda)$ is equivalent to choosing a distinguished equivalence class member and taking all objects that map to that member under θ .

In this section, we present two good choices for distinguished elements that avoid the need for explicitly constructing the isomorphism θ .

For $T \in \mathcal{A}$, let $\alpha(T)$ be the composition of n corresponding to $\text{Des}(T)$. Recall the dominance order on partitions of n , which we extend to compositions of n by

$$\alpha \geq \beta \Leftrightarrow \alpha_1 + \cdots + \alpha_k \geq \beta_1 + \cdots + \beta_k \quad \forall k. \quad (4.3)$$

We can now define the set of distinguished elements.

DEFINITION 4.6. Let $\{\varphi_i\}$ be a dual equivalence for $(\mathcal{A}, \text{Des})$. Then $T \in \mathcal{A}$ is called *dominant* if $\alpha(T) \geq \alpha(S)$ for every S in the dual equivalence class of T .

Since dominance order is a partial order, it is not immediately obvious that dominant objects exist. Not only do they exist, but each dual equivalence class contains a unique dominant element, and $\alpha(T)$ is a partition for T dominant.

THEOREM 4.7. *Let $\{\varphi_i\}$ be a dual equivalence for $(\mathcal{A}, \text{Des})$ preserving stat . Then*

$$f(X; q) = \sum_{T \in \text{Dom}(\mathcal{A})} q^{\text{stat}(T)} s_{\alpha(T)}(X) = \sum_{\lambda} \left(\sum_{\substack{T \in \text{Dom}(\mathcal{A}) \\ \alpha(T) = \lambda}} q^{\text{stat}(T)} \right) s_{\lambda}(X), \quad (4.4)$$

where $\text{Dom}(\mathcal{A})$ is the set of dominant objects of \mathcal{A} with respect to $\{\varphi_i\}$.



Figure 14. The superstandard (dominant) and substandard (subordinate) tableaux of shape $(4, 3, 2)$.

Proof. Given λ , let $T_\lambda \in \text{SYT}(\lambda)$ denote the *superstandard tableau of shape* λ obtained by filling the first row with $1, 2, \dots, \lambda_1$, the second row with $\lambda_1 + 1, \dots, \lambda_1 + \lambda_2$, and so on. For an example, see Figure 14. For any $T \in \text{SYT}(\lambda)$, we have $\alpha(T) \leq \lambda$, with equality if and only if $T = T_\lambda$. Since the dual equivalence classes on tableaux include all tableaux of a given shape, each dual equivalence class contains a unique dominant element, and the map α gives the corresponding Schur function for the class. The result for arbitrary \mathcal{A} now follows from Corollary 4.4 since the descent sets must be the same for the elements of a dual equivalence class and the set of tableaux of shape λ for some partition λ . \square

REMARK 4.8. Theorem 4.7 makes use of the implicit bijection between \mathcal{A} and tableaux that exists whenever there is a dual equivalence for \mathcal{A} . This bijection can be realized by identifying each $T \in \text{Dom}(\mathcal{A})$ with the superstandard tableau $T_{\alpha(T)}$ and then applying the same sequence of dual equivalence involutions to both.

REMARK 4.9. There is another distinguished element that can be chosen from each equivalence class which is almost as natural as the dominant element. Say that $T \in \mathcal{A}$ is *subordinate* if $\alpha(T) \leq \alpha(S)$ for every S in the dual equivalence class of T . For example, the right tableau in Figure 14 is the subordinate tableau of shape $(4, 3, 2)$. Then each dual equivalence class contains a unique subordinate object. Define a map β on subordinate objects by sending T to $\alpha([n - 1] \setminus \text{Des}(T))'$. That is, complement the set and conjugate the shape. Then, in Theorem 4.7, the set Dom of dominant objects may be replaced with the set Sub of subordinate objects when α is replaced with β .

5. A graph for LLT polynomials

5.1. LLT polynomials. The LLT polynomial $\tilde{G}_\mu^{(k)}(x; q)$, originally defined by Lascoux *et al.* [9] in 1997, is the q -generating function of k -ribbon tableaux of shape μ weighted by a statistic called *cospin*. By the Stanton–White correspondence [13], k -ribbon tableaux are in bijection with certain k -tuples of tableaux, from which it follows that LLT polynomials are q -analogs of products of Schur functions. An alternative definition of $\tilde{G}_\mu(x; q)$ as the q -generating function of tuples of semistandard tableaux of shapes $\boldsymbol{\mu} = (\mu^{(0)}, \dots, \mu^{(k-1)})$

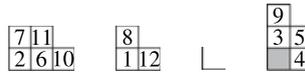


Figure 15. A standard 4-tuple of shape $((3, 2), (2, 1), \emptyset, (2, 2, 1)/(1))$.

weighted by a statistic called diagonal inversions is given in [6]. For a detailed account of the equivalence of these definitions (actually $q^a \tilde{G}_\mu^{(k)}(x; q) = \tilde{G}_\mu(x; q)$ for a constant $a \geq 0$ depending on μ), see [2, 6].

Extending prior notation, for $\lambda = (\lambda^{(0)}, \dots, \lambda^{(k-1)})$, define

$$\text{SSYT}(\lambda) = \{\text{semi-standard tuples of tableaux of shapes } (\lambda^{(0)}, \dots, \lambda^{(k-1)})\},$$

$$\text{SYT}(\lambda) = \{\text{standard tuples of tableaux of shapes } (\lambda^{(0)}, \dots, \lambda^{(k-1)})\}.$$

As with tableaux, if $\mathbf{T} = (T^{(0)}, \dots, T^{(k-1)}) \in \text{SSYT}(\lambda)$ has entries $1^{\pi_1}, 2^{\pi_2}, \dots$, then we say that \mathbf{T} has *shape* λ and *weight* π . Note that a standard tuple of tableaux has weight (1^n) (for an example, see Figure 15), and this is not the same as a k -tuple of standard tableaux.

For a k -tuple of (skew) shapes $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$, define the *shifted content* of a cell x by

$$\tilde{c}(x) = k \cdot c(x) + i \tag{5.1}$$

when x is a cell of $\lambda^{(i)}$, where $c(x)$ is the usual content of x regarded as a cell of $\lambda^{(i)}$. For $\mathbf{T} \in \text{SSYT}(\lambda)$, let $\mathbf{T}(x)$ denote the entry of the cell x in \mathbf{T} . Define the *set of diagonal inversions of \mathbf{T}* by

$$\text{dInv}(\mathbf{T}) = \{(x, y) \mid k > \tilde{c}(y) - \tilde{c}(x) > 0 \text{ and } \mathbf{T}(x) > \mathbf{T}(y)\}. \tag{5.2}$$

Then the *diagonal inversion number of \mathbf{T}* is given by

$$\text{dinv}(\mathbf{T}) = |\text{dInv}(\mathbf{T})|. \tag{5.3}$$

It will also be convenient to track the *set of diagonal descents of \mathbf{T}* , defined by

$$\text{dDes}(\mathbf{T}) = \{(x, y) \mid \tilde{c}(y) - \tilde{c}(x) = k \text{ and } \mathbf{T}(x) > \mathbf{T}(y)\}. \tag{5.4}$$

Of course, these descents are determined by the *shape* of \mathbf{T} by the increasing rows and columns condition for tableaux.

For example, suppose that \mathbf{T} is the 4-tuple of tableaux in Figure 15. Since \mathbf{T} is standard, let us abuse notation by representing a cell of \mathbf{T} by the entry it contains. Then the set of diagonal inversions is

$$\text{dInv}(\mathbf{T}) = \left\{ (9, 7), (9, 8), (7, 3), (8, 3), (8, 2), (3, 2), (3, 1), (2, 1), (11, 1), (11, 5), (6, 4), (12, 4), (12, 10) \right\},$$

and so $\text{div}(\mathbf{T}) = 13$. The diagonal descents describe the shape of \mathbf{T} . In this case,

$$\text{dDes}(\mathbf{T}) = \{(7, 2), (11, 6), (8, 1), (9, 3), (5, 4)\}.$$

By [6], the LLT polynomial $\tilde{G}_\mu(x; q)$ is given by

$$\tilde{G}_\mu(x; q) = \sum_{\mathbf{T} \in \text{SSYT}(\mu)} q^{\text{div}(\mathbf{T})} x^{\mathbf{T}}, \quad (5.5)$$

where $x^{\mathbf{T}}$ is the monomial $x_1^{\pi_1} x_2^{\pi_2} \cdots$ when \mathbf{T} has weight π .

Notice that when $q = 1$, (5.5) reduces to a product of Schur functions:

$$\sum_{\mathbf{T} \in \text{SSYT}(\lambda)} x^{\mathbf{T}} = \prod_{i=0}^{k-1} \sum_{T^{(i)} \in \text{SSYT}(\lambda^{(i)})} x^{T^{(i)}} = \prod_{i=0}^{k-1} s_{\lambda^{(i)}}(x). \quad (5.6)$$

Define the *content reading word* of a tuple of tableaux to be the word obtained by reading entries in increasing order of shifted content and reading diagonals southwest to northeast. For the example in Figure 15, the content reading word is (9, 7, 8, 3, 2, 11, 1, 5, 6, 12, 4, 10).

For \mathbf{T} a standard tuple of tableaux, define $\sigma(\mathbf{T})$ analogously to (2.3) using the content reading word. Expressed in terms of quasisymmetric functions, (5.5) becomes

$$\tilde{G}_\mu(x; q) = \sum_{\mathbf{T} \in \text{SYT}(\mu)} q^{\text{div}(\mathbf{T})} Q_{\sigma(\mathbf{T})}(x). \quad (5.7)$$

Using Fock space representations of quantum affine Lie algebras constructed by Kashiwara *et al.* [8], Lascoux *et al.* [9] proved that $\tilde{G}_\mu(x; q)$ is a symmetric function. Thus we may define the Schur coefficients, $\tilde{K}_{\lambda, \mu}(q)$, by

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

Using Kazhdan–Lusztig theory, Leclerc and Thibon [10] proved that $\tilde{K}_{\lambda, \mu}(q) \in \mathbb{N}[q]$ for μ a tuple of straight shapes. The original motivation for dual equivalence graphs is to understand these Schur coefficients combinatorially and for arbitrary skew shapes.

5.2. Dual equivalence for tuples of tableaux. Haiman's dual equivalence involutions do not always preserve the increasing row and column conditions for a standard tuple of shapes. Therefore something more is needed.

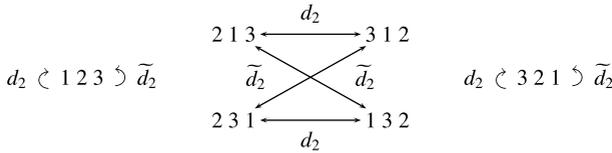


Figure 16. The two dual equivalences for \mathfrak{S}_3 .

DEFINITION 5.1. Define the *elementary twisted dual equivalence involution* \tilde{d}_i , $1 < i < n$, on permutations w as follows. If i lies between $i - 1$ and $i + 1$ in w , then $\tilde{d}_i(w) = w$. Otherwise, \tilde{d}_i cyclically rotates $i - 1, i, i + 1$ so that i lies on the other side of $i - 1$ and $i + 1$.

Note that Haiman's dual equivalence involutions and the twisted dual equivalence involutions are the only two possible dual equivalences for \mathfrak{S}_3 ; see Figure 16. For the general case, the number of dual equivalences on \mathfrak{S}_n is determined by the multiplicities of permutations of n with each possible inverse descent set.

Define the distance between two entries i and j of $\mathbf{T} \in \text{SYT}(\mu)$ to be the difference of the shifted contents of their cells, with the extension $\text{dist}(a_1, \dots, a_l) = \max_{i,j} \{\text{dist}(a_i, a_j)\}$. Note that none of $i - 1, i, i + 1$ may occur with the same content. For fixed k , combine these two dual equivalences into an involution D_i on k -tuples of tableaux by

$$D_i(w) = \begin{cases} d_i(w) & \text{if } \text{dist}(i - 1, i, i + 1) > k, \\ \tilde{d}_i(w) & \text{if } \text{dist}(i - 1, i, i + 1) \leq k, \end{cases} \quad (5.8)$$

where w is the content reading word of \mathbf{T} .

PROPOSITION 5.2. For μ a tuple of shapes and $\mathbf{T} \in \text{SYT}(\mu)$, \mathbf{T} and $D_i(\mathbf{T})$ have the same diagonal descent set and the same diagonal inversion number. In particular, D_i is a well-defined involution on $\text{SYT}(\mu)$ that preserves the number of diagonal inversions.

Proof. If i lies between $i - 1$ and $i + 1$ in $w(\mathbf{T})$, then the assertion is trivial. Assume then that i does not lie between $i - 1$ and $i + 1$ in $w(\mathbf{T})$. If $\text{dist}(i - 1, i, i + 1) > k$ in \mathbf{T} , then the relative order of the middle letter with the two outer letters remains unchanged, and the outer letters do not form a potential descent or a potential inversion. Therefore $\text{dDes}(\mathbf{T}) = \text{dDes}(d_i(\mathbf{T}))$ and $\text{dInv}(\mathbf{T}) = \text{dInv}(d_i(\mathbf{T}))$.

If $\text{dist}(i - 1, i, i + 1) \leq k$ in \mathbf{T} , then the relative order of the outer letters in \mathbf{T} is the same as the relative order of the outer letters in $\tilde{d}_i(\mathbf{T})$. If these positions

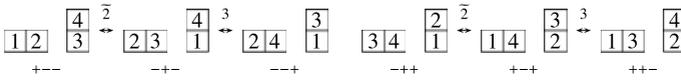


Figure 17. The graph on domino tableaux of shape $((2), (1, 1))$.

form a potential descent, which happens if and only if $\text{dist}(i - 1, i, i + 1) = k$, the descent or lack thereof between the outer two letters is preserved. Therefore $\text{dDes}(\mathbf{T}) = \text{dDes}(\tilde{d}_i(\mathbf{T}))$. In the case when $\text{dist}(i - 1, i, i + 1) < k$, these outer letters form a potential inversion, which similarly is preserved. The map \tilde{d}_i toggles the middle letter between the smallest, $i - 1$, and the largest, $i + 1$. Therefore exactly one of the potential inversions involving the middle letter is an inversion, and this holds for \mathbf{T} as well as for $\tilde{d}_i(\mathbf{T})$. Therefore $\text{dInv}(\mathbf{T}) = \text{dInv}(\tilde{d}_i(\mathbf{T}))$, though $\text{dInv}(\mathbf{T}) \neq \text{dInv}(\tilde{d}_i(\mathbf{T}))$.

Since D_i preserves the diagonal descent set, it is a well-defined involution on $\text{SYT}(\mu)$. □

Let \mathcal{G}_μ be the signed colored graph on $\text{SYT}(\mu)$ determined by the involutions D_i . An example of \mathcal{G}_μ is given in Figure 17.

Since the graph in Figure 17 is a dual equivalence graph, we have

$$\tilde{G}_{(2),(1,1)}(x; q) = qs_{3,1}(x) + q^2s_{2,1,1}(x).$$

In general, \mathcal{G}_μ does not satisfy dual equivalence axioms 4 or 6. Instead of having restricted equivalence classes be single Schur functions, the restricted equivalence classes for D_i are conjecturally Schur positive.

DEFINITION 5.3. Let \mathcal{A} be a finite set, and let Des be a descent set map on \mathcal{A} such that $\text{Des}(T) \subseteq [n - 1]$ for all $T \in \mathcal{A}$. A *weak dual equivalence* for $(\mathcal{A}, \text{Des})$ is a family of involutions $\{\varphi_i\}_{1 < i < n}$ on \mathcal{A} such that the following hold.

- (i) For all $|i - j| \leq 3$ and all $T \in \mathcal{A}$,

$$\sum_{U \in [T]_{(j,i)}} Q_{\text{Des}_{(j,i)}(U)}(X)$$

is Schur positive, and is a single Schur function if $i = j$.

- (ii) For all $|i - j| \geq 3$ and all $T \in \mathcal{A}$, we have

$$\varphi_j \varphi_i(T) = \varphi_i \varphi_j(T).$$

We refer to condition (i) of Definition 5.3 as *local Schur positivity*.

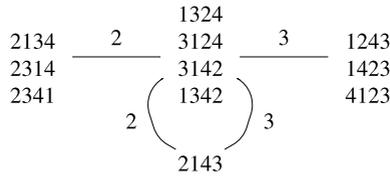


Figure 18. Possible vertices and edges for graph using d_i or \tilde{d}_i for edges colored $i = 2, 3$.

THEOREM 5.4. *For μ a tuple of shapes, the involutions $\{D_i\}$ give a weak dual equivalence for $\text{SYT}(\mu)$, and when μ consists of at most two shapes, this is a strong dual equivalence.*

Proof. Condition (i) for $i = j$ follows immediately from the fact that both d_i and \tilde{d}_i are dual equivalences for \mathfrak{S}_3 . Condition (ii) follows from the fact that, if $|i - j| \geq 3$, then $\{i - 1, i, i + 1\}$ and $\{j - 1, j, j + 1\}$ are disjoint.

For $|i - j| = 1$, consider first a component of $E_2 \cup E_3$ containing a vertex with signature $\sigma_{1,2,3} = -++$. The only possible permutations with this descent pattern are the three depicted on the left side of Figure 18. Applying d_2 and \tilde{d}_2 to each of these gives one of the top four permutations in the middle of Figure 18. Applying d_3 and \tilde{d}_3 to each of these gives either one of the three permutations on the right of Figure 18, or the bottom permutation in the middle of Figure 18. Finally, applying d_2 and \tilde{d}_2 to 2143 gives either 1342 or 3142, both of which appear in the middle of Figure 18. Thus traversing the graph in Figure 18 by starting on the left and alternating between 2-edges and 3-edges must eventually end on the right. Taking signatures into account, the possible generating functions are of the form $s_{(3,1)} + ms_{(2,2)}$ for some $m \in \mathbb{N}$ (in fact, a more detailed analysis shows that $m = 0$ or 1). In particular, a component containing a vertex with signature $\sigma_{1,2,3} = -++$ is Schur positive. The same figure applies when working with a vertex with signature $\sigma_{1,2,3} = ++-$. Reversing the permutations and multiplying the signatures componentwise by -1 proves the result for components with a vertex with signature $\sigma_{1,2,3} = +- -$ or $- - +$. A component with a vertex with signature $\sigma_{1,2,3} = +++$ or $---$ is a single vertex and has generating function $s_{(4)}$ or $s_{(1,1,1,1)}$, respectively. The only remaining case is an alternating loop of vertices with signatures $\sigma_{1,2,3} = +- +$ and $- + -$. As before, the top four permutations in the middle of Figure 18 must connect by a 2-edge and a 3-edge to 2143, and thus the component has two vertices and generating function $s_{(2,2)}$. Similarly, the reverse of the top four permutations in the middle of Figure 18 must connect by a 2-edge and a 3-edge to 3412, and thus the component has two vertices and generating function $s_{(2,2)}$. This covers all cases, so local Schur

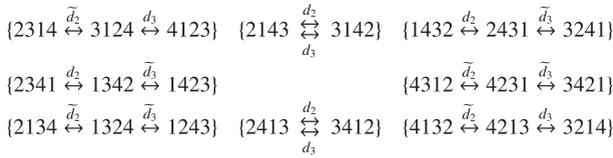


Figure 19. The nontrivial dual equivalence graphs on \mathfrak{S}_4 using \tilde{d}_i only when $i - 1, i, i + 1$ are adjacent.

positivity holds. Moreover, the case when μ has two shapes can be seen to be a single Schur function directly from Figure 19.

For $|i - j| = 2, 3$, we enumerate all cases and check. Notice that one may regard the shape μ as specifying attacking positions in a permutation. That is, for a permutation w , say that w_p attacks w_q if $p < q$ and the difference in shifted contents between p and q is at most k . Therefore the structure of \mathcal{G}_μ is given by the graph on permutations where the edges are given by \tilde{d}_i if i attacks the rightmost of $i \pm 1$ or if the leftmost of $i \pm 1$ attacks i , and by d_i otherwise. Since attacking positions are determined by distance, if w_p attacks w_r with $p < q < r$, then w_p attacks w_q as well. Therefore the graph on permutations of size n is determined by (a_1, \dots, a_{n-1}) , where a_j is the rightmost position that w_j attacks. Since none of $i - 1, i, i + 1$ can have the same shifted content, we may assume that each position attacks its right neighbor, and so $j + 1 \leq a_j \leq n$. Moreover, if w_p attacks w_r with $p < q < r$, then w_q attacks w_r as well. Therefore $a_p \leq a_{p+1}$. Hence the number of attacking vectors to consider for permutations of length n is the $n - 1$ th Catalan number. In particular, there are 14 graph structures on permutations of 5 and 42 on permutations of 6. These cases can be checked by hand or by computer. \square

An alternative and completely elementary proof of Theorem 5.4 for μ consisting of at most two shapes follows immediately from [1, Theorem 5.3].

Theorems 4.2 and 5.4 together with Proposition 5.2 prove the following, which can also be found in [1, Corollary 5.4].

COROLLARY 5.5. For $\mu = (\mu^{(0)}, \mu^{(1)})$, the LLT polynomial $\tilde{G}_\mu(x; q)$ is Schur positive.

In 1995, Carré and Leclerc [3] gave a combinatorial interpretation of $\tilde{K}_{\lambda, \mu}(q)$ when μ has two shapes in their study of 2-ribbon tableaux, though a complete proof of their result was not found until 2005 by van Leeuwen [14] using the theory of crystal graphs. That proof is quite long and involved, as compared to the one-page proof above using dual equivalence.

$$\begin{array}{l}
 \{2314 \xleftrightarrow{\tilde{d}_2} 3124 \xleftrightarrow{\tilde{d}_3} 2143 \xleftrightarrow{\tilde{d}_2} 1342 \xleftrightarrow{\tilde{d}_3} 1423\} \quad \{1432 \xleftrightarrow{\tilde{d}_2} 2413 \xleftrightarrow{\tilde{d}_3} 3214\} \\
 \{2341 \xleftrightarrow{\tilde{d}_2} 3142 \xleftrightarrow{\tilde{d}_3} 4123\} \quad \{4312 \xleftrightarrow{\tilde{d}_2} 4231 \xleftrightarrow{\tilde{d}_3} 3421\} \\
 \{2134 \xleftrightarrow{\tilde{d}_2} 1324 \xleftrightarrow{\tilde{d}_3} 1243\} \quad \{4132 \xleftrightarrow{\tilde{d}_2} 4213 \xleftrightarrow{\tilde{d}_3} 3412 \xleftrightarrow{\tilde{d}_2} 2431 \xleftrightarrow{\tilde{d}_3} 3241\}
 \end{array}$$

Figure 20. The twisted dual equivalence classes of \mathfrak{S}_4 .

In general, D_i is not a dual equivalence. For instance, if w has the pattern 2431 with $\text{dist}(1, 2, 3) \leq k$ (which forces $k \geq 3$), then $D_2(w)$ contains the pattern 3412. However, $D_3(w)$ contains the pattern 3241, which is not the same. For example, see Figure 20. Therefore the restricted generating function is not a single Schur function (though it is Schur positive).

CONJECTURE 5.6. *For μ a tuple of shapes, each equivalence class under D_i is symmetric and Schur positive.*

Conjecture 5.6 has been verified for $n \leq 12$. Additionally, the extreme case when each shape in μ is a single (nonskewed) box has connected components with particularly nice Schur expansions that can be proved by more elementary means.

In this case, $D_i = \tilde{d}_i$ for all i , and so there are no double edges in \mathcal{G}_μ . For the standard dual equivalence graphs, \mathcal{G}_λ has no double edges if and only if λ is a hook, that is, $\lambda = (m, 1^{n-m})$ for some $m \leq n$. Therefore the generating function for a dual equivalence graph with no double edges is a sum of Schur functions indexed by hooks. The analog of this fact for \mathcal{G}_μ is that the generating function is a sum of skew Schur functions indexed by ribbons.

Let ν be a ribbon of size n . Label the cells of ν from 1 to n in increasing order of content. Define the *descent set* of ν , denoted $\text{Des}(\nu)$, to be the set of indices i such that the cell labeled $i + 1$ lies south of the cell labeled i . Define the *major index of a ribbon* by

$$\text{maj}(\nu) = \sum_{i \in \text{Des}(\nu)} i. \tag{5.9}$$

Any connected component of \mathcal{G}_μ such that $D_i = \tilde{d}_i$ on the entire component not only has constant diagonal inversion number, but the relative ordering of the first and last letters of each vertex is constant as well. This is because \tilde{d}_i does not change the relative order of the outer two letters among $i - 1, i, i + 1$. Therefore, $w_1 > w_n$ for some $w \in [u]$ if and only if $w_1 > w_n$ for all $w \in [u]$. When $w_1 > w_n$, say that $(1, n)$ is an *inversion* in $[u]$.

THEOREM 5.7. *Let $[U]$ be an equivalence class for $\text{SYT}(\mu)$ under D_i for which $D_i(T) = \tilde{d}_i(T)$ for all $T \in [U]$. Then*

$$\sum_{T \in [U]} Q_{\sigma(T)}(x) = \sum_{\substack{v \text{ an } n\text{-ribbon} \\ \text{maj}(v) = \text{dinv}(U) \\ n-1 \in \text{Des}(v) \Leftrightarrow (1, n) \in \text{Inv}(U)}} s_v. \tag{5.10}$$

Proof. We may assume that D_i is acting by \tilde{d}_i on the set of permutations of $[n]$, in which case diagonal inversions are the usual inversions for permutations. By earlier remarks, for $S, T \in [U]$, $\text{inv}(S) = \text{inv}(T)$ and $(1, n) \in \text{Inv}(S)$ if and only if $(1, n) \in \text{Inv}(T)$. In fact, it is an easy exercise to show by induction that this necessary condition for two vertices to coexist in $[U]$ is also sufficient. That is to say, $[U]$ is the set of words S with $\text{inv}(S) = \text{inv}(U)$ and $(1, n) \in \text{Inv}(S)$ if and only if $(1, n)$ is an inversion of U .

Recall Foata’s bijection on words [4]. For w a word and x a letter, ϕ is built recursively using an inner function γ_x by $\phi(wx) = \gamma_x(\phi(w))x$. From this structure it follows that the last letter of w is the same as the last letter of $\phi(w)$. Furthermore, γ_x is defined so that the last letter of w is greater than x if and only if the first letter of $\gamma_x(w)$ is greater than x , and ϕ preserves the descent set of the inverse permutation; that is, $\sigma(w) = \sigma(\phi(w))$. Finally, the bijection satisfies $\text{maj}(w) = \text{inv}(\phi(w))$. Summarizing these properties, ϕ is a σ -preserving bijection between the following sets:

$$\{w \mid \text{inv}(w) = j \text{ and } (1, n) \in \text{Inv}(w)\} \xleftrightarrow{\sim} \{w \mid \text{maj}(w) = j \text{ and } n - 1 \in \text{Des}(w)\},$$

$$\{w \mid \text{inv}(w) = j \text{ and } (1, n) \notin \text{Inv}(w)\} \xleftrightarrow{\sim} \{w \mid \text{maj}(w) = j \text{ and } n - 1 \notin \text{Des}(w)\}.$$

A standard filling of a ribbon v is just a permutation w such that $\text{Des}(w) = \text{Des}(v)$. Therefore, by (2.4), the Schur function s_v may be expressed as

$$s_v(x) = \sum_{\text{Des}(w) = \text{Des}(v)} Q_{\sigma(w)}(x). \tag{5.11}$$

Applying ϕ to this formula yields (5.10). □

Acknowledgements

The author is grateful to Mark Haiman for inspiring and helping to develop many of the ideas contained in this paper and in its precursor [2]. The author is indebted to N. Bergeron, S. Billey, A. Garsia, M. Haiman, J. Haglund, G. Musiker, and F. Sottile for carefully reading earlier drafts and providing feedback that greatly improved the exposition.

References

- [1] S. Assaf, 'A generalized major index statistic', *Sém. Lothar. Combin.* **60** (2008), Art. B50c; 13pp. (electronic).
- [2] S. H. Assaf, 'Dual equivalence graphs, ribbon tableaux and Macdonald polynomials', PhD Thesis, University of California Berkeley, 2007.
- [3] C. Carré and B. Leclerc, 'Splitting the square of a Schur function into its symmetric and antisymmetric parts', *J. Algebraic Combin.* **4**(3) (1995), 201–231.
- [4] D. Foata, 'On the Netto inversion number of a sequence', *Proc. Amer. Math. Soc.* **19** (1968), 236–240.
- [5] I. M. Gessel, 'Multipartite P -partitions and inner products of skew Schur functions', in: *Combinatorics and Algebra (Boulder, CO, 1983)*, Contemporary Mathematics, 34 (American Mathematical Society, Providence, RI, 1984), 289–317.
- [6] J. Haglund, M. Haiman, N. Loehr, J. B. Remmel and A. Ulyanov, 'A combinatorial formula for the character of the diagonal coinvariants', *Duke Math. J.* **126** (2005), 195–232.
- [7] M. D. Haiman, 'Dual equivalence with applications, including a conjecture of Proctor', *Discrete Math.* **99**(1–3) (1992), 79–113.
- [8] M. Kashiwara, T. Miwa and E. Stern, 'Decomposition of q -deformed Fock spaces', *Selecta Math. (N.S.)* **1**(4) (1995), 787–805.
- [9] A. Lascoux, B. Leclerc and J.-Y. Thibon, 'Ribbon tableaux, Hall–Littlewood functions, quantum affine algebras, and unipotent varieties', *J. Math. Phys.* **38**(2) (1997), 1041–1068.
- [10] B. Leclerc and J.-Y. Thibon, 'Littlewood–Richardson coefficients and Kazhdan–Lusztig polynomials', in: *Combinatorial Methods in Representation Theory (Kyoto, 1998)*, Advanced Studies in Pure Mathematics, 28 (Kinokuniya, Tokyo, 2000), 155–220.
- [11] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, 2nd edn, Oxford Mathematical Monographs (The Clarendon Press, Oxford University Press, New York, 1995). With contributions by A. Zelevinsky, Oxford Science Publications.
- [12] A. Roberts, 'Dual equivalence graphs revisited and the explicit Schur expansion of a family of LLT polynomials', *J. Algebraic Combin.* **39**(2) (2014), 389–428.
- [13] D. W. Stanton and D. E. White, 'A Schensted algorithm for rim hook tableaux', *J. Combin. Theory Ser. A* **40**(2) (1985), 211–247.
- [14] M. A. A. van Leeuwen, 'Spin-preserving Knuth correspondences for ribbon tableaux', *Electron. J. Combin.* **12** (2005), Research Paper 10; 65pp. (electronic).