Discrete Mathematics 99 (1992) 79–113
North-Holland

Dual equivalence with applications, including a conjecture of Proctor

Mark D. Haiman*
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02 139, USA

Received 6 June 1989
Revised 5 February 1991

Abstract

Haiman, M.D., Dual equivalence with applications, including a conjecture of Proctor, Discrete Mathematics 99 (1992) 79–113.

We make a systematic study of a new concept in the theory of *jeu-de-taquin,* which we call *dual equivalence.* Using this, we prove a conjecture of Proctor establishing a bijection between standard tableaux of 'shifted staircase' shape and reduced expressions for the longest element in the Coxeter group B_l. We also get a new and more illuminating proof of the analogous theorem, due to Greene and Edelman, for the Coxeter group A_l, and arrive at yet one more theorem of a similar type. We explain some symmetric functions associated to reduced expressions by Stanley and prove his conjecture that one of these for B_l is the Schur function s_λ for λ an l-by-l square. We classify shifted and unshifted shapes for which the *total promotion* operator has special properties; in one case this proves another conjecture of Stanley. We determine the previously unknown 'dual Knuth relations' for the shifted Schensted correspondence.

1. Introduction; notational conventions

The immediate purpose of this article is to prove a conjecture of Proctor establishing a bijection between reduced expresions for the longest element w_{B_l} of the Coxeter group of type B_l and standard tableaux of 'shifted staircase' shape $(2l-1, 2l-3, \ldots, 1)$. This generalizes a theorem of Greene and Edelman for the Coxeter group of type A_l and the unshifted staircase shape $(l, l-1, \ldots, 1)$. As it turns out, we get a new and simpler proof of the Greene–Edelman theorem and yet another theorem of the same type which had not even been conjectured before. We show how these bijections explain the mysterious symmetric functions used by Stanley to study the case A_l and prove a conjecture of his concerning one of these functions for B_l.

* Research supported in part by N.S.F. grant number DMS-8717795.

Our real purpose, however, is not only to obtain the above results but to develop a heretofore missing but crucial tool in the theory of *jeu-de-taquin* and the Schensted correspondence. This tool is a relation on tableaux we call *dual equivalence* because it is precisely dual in the sense of the Schensted correspondence to the relation of *jeu-de-taquin* equivalence. In particular, for permutations, it captures the property of having the same 'recording' tableau. However, it is defined intrinsically for any tableaux, and has a marvelous number of equivalent alternative characterizations.

Once dual equivalence is understood, all manner of things in addition to the Proctor conjecture and Greene–Edelman theorem follow quite naturally, of which we discuss a number here. They include the following.

(1) The fundamental theorems of *jeu-de-taquin,* which can now be given purely *jeu-de-taquin* proofs which are unified for the shifted and unshifted cases and do not require the use of any extraneous considerations such as the Greene invariant or detailed analysis of the Schensted 'bumping' process.

(2) Dual 'Knuth relations'. These can be determined immediately, again without the Greene invariant. In the shifted case, they were not known before.

(3) Shapes for which total promotion has special properties. All the known cases are covered by the theory, as well as a case that had only been conjectured.

We deal in this paper only with shifted and unshifted standard tableaux. Thus for our purposes a *tableau* T is an order-preserving bijective function from a shape $\operatorname{sh} T = \lambda$ to the set $\{1, \ldots, n = |\lambda|\}$, where a *shape* is a segment in the partial order $\mathbf{Z} \times \mathbf{Z}$ (the *plane*) or $\{(i, j) \in \mathbf{Z} \times \mathbf{Z} \mid i \leq j\}$ (the *shifted plane*). We view the planes as ordered matrix-style with i increasing downward and j to the right, so the shifted plane lies on and above the diagonal in the plane. All figures and all such terminology as 'above', 'below', 'left', 'right' use this matrix convention. The elements of the planes are called *cells.*

A *normal* shape is one with a unique upper-left corner, which in the shifted case is required to lie on the diagonal. A normal shape can be specified by giving its row lengths as a *partition,* or in the shifted case, a *strict partition,* e.g., $(3, 2, 1)$ indicates a triangular shape with six cells.

One operates on tableaux by *jeu-de-taquin slides* [9]. Say that a shape μ *extends* λ if $\lambda \cup \mu$ is a shape containing λ and μ as complementary initial and final segments. If $\{c\}$ extends $\operatorname{sh} T$, one performs a *forward slide on T into the cell c* by moving the greater (or only) of the entries of T to c's left and above c into the cell c, then continuing in like fashion with the cell just vacated, until one reaches an upper-left-most cell of $T \cup \{c\}$ where one stops. This last cell is said to be *vacated* by the slide. If $\operatorname{sh} T$ extends $\{c\}$, one performs a *reverse slide on T into c* by the opposite process.

Tableaux that can be obtained each from the other by slides are *jeu-de-taquin equivalent.* Equivalence can be tested using the *Schensted correspondence*: let the *reading word* $w = w(T)$ be the permutation consisting of the entries of T read left-to-right from bottom to top, one row at a time. We denote the Schensted insertion and 'recording' tableaux for w by $\Box \leftarrow w$ and $R{:}\Box \leftarrow w$. In the shifted case there is a *shifted Schensted correspondence* due to Worley and Sagan [7, 12] whose insertion and recording tableaux we denote $\boxbslash \leftarrow w$ and $R{:}\boxbslash \leftarrow w$. Then $\Box \leftarrow w$ (or $\boxbslash \leftarrow w$) is the unique tableau of normal shape equivalent to T.

A sequence of cells $(c_1, \ldots, c_l)$ describes a *sequence of slides* for T if it is meaningful to form $T = T_0, T_1, \ldots, T_l$ in which each T_j is the result of a slide on T_{j-1} into the cell c_j. In particular, if X is a tableau such that $\operatorname{sh} X$ extends $\operatorname{sh} T$, then the cells of X taken in the order of the entries of X form a sequence of forward slides for T. The result of applying this sequence we denote $j_X(T)$. Likewise, if $\operatorname{sh} T$ extends $\operatorname{sh} X$ we form $j^X(T)$ by reverse slides into the cells of X, taken in reverse order. The sequence of cells vacated as we form $j_X(T)$ describes a tableau which we denote $V{:}j_X(T)$; similarly we obtain $V{:}j^X(T)$. Evidently $j^{V:j_X(T)}(j_X(T)) = T = j_{V:j^Y(T)}(j^Y(T))$.

When, as above, we have tableaux such that $\mu = \operatorname{sh} X$ extends $\lambda = \operatorname{sh} T$ we will sometimes write $T \cup X$ to mean the tableau of shape $\lambda \cup \mu$ whose least $|T|$ entries form the tableau T and whose greatest $|X|$ entries form X (up to a constant added to all entries). Similarly we have $Y \cup T \cup X$ and like expressions.

Both the plane and shifted plane have an order-reversing reflection sending (i, j) to $(-j, -i)$. Operations, properties, or statements that correspond under this reflection will be referred to as 'anti'. Thus a forward slide is the *anti-operation* of a reverse slide, and the two identities $j^{V:j_X(T)}(j_X(T)) = T = j_{V:j^Y(T)}(j^Y(T))$ are one another's *anti-statements*.

2. Definition of dual equivalence and alternative formulations

We now proceed to our key concept, an equivalence relation among tableaux of a given shape λ. The definition we give here is just one of many possible characterizations of this relation. Other such characterizations are established in the propositions following the definition.

Definition. Let S and T be tableaux. Suppose that every sequence $(c_1, \ldots, c_k)$ which is a sequence of slides for both S and T yields two taleaux of the same shape, when applied to S and T respectively. Then S and T are said to be *dual equivalent*, written $S \approx T$.

In the definition, $(c_1, \ldots, c_k)$ is allowed to be the empty sequence; thus $S \approx T$ entails that $\operatorname{sh} S = \operatorname{sh} T$. To see that dual equivalence is a genuine equivalence relation, note first that it is obviously reflexive and symmetric. If $S \approx T$ then it is easy to see, by induction on the length of the sequence, that any sequence of slides for S is also a sequence of slides for T. Transitivity follows directly from this observation.

The sequence $(c_1, \ldots, c_k)$ is allowed to involve both forward and reverse slides. Thus it is implicit in the definition that anti-dual equivalence is the same as dual equivalence.

A warning to the reader is in order here about the shifted and unshifted versions of dual equivalence: unshifted tableaux S and T may be dual equivalent, but not be when regarded as shifted tableaux. This is because slides are possible in the shifted plane which cannot be simulated in the unshifted one.

Crucial to our study of dual equivalence in this section and to our applications in later sections will be the reduction of a general dual equivalence to a chain of especially simple 'elementary' ones. The reduction is based on the following simple observation.

Lemma 2.1. *Let λ, μ, and ν be shapes such that μ extends λ and ν extends $\lambda \cup \mu$. Let X, Y, S, and T be tableaux with* $\operatorname{sh} X = \lambda$, $\operatorname{sh} S = \operatorname{sh} T = \mu$, *and* $\operatorname{sh} Y = \nu$. *If $S \approx T$, then $X \cup S \cup Y \approx X \cup T \cup Y$.*

Proof. Consider the action of a forward slide, say into a cell c, on $X \cup S \cup Y$. Its effect is to first slide Y into the cell c, then slide S into the cell vacated by the slide on Y, and finally slide X into the cell vacated by this slide on S. This results in a tableaux $X' \cup S' \cup Y'$. If $S \approx T$, then any slide on T vacates the same cell as the corresponding slide on S. Therefore sliding $X \cup T \cup Y$ into c yields a tableaux $X' \cup T' \cup Y'$, with the same X', Y' as before. Moreover, we will have $S' \approx T'$, since they result from slides of S and T into the same cell. All we have just said applies, mutatis mutandis, to reverse slides as well. Thus applying a slide to $X \cup S \cup Y$ and $X \cup T \cup Y$, we obtain two tableaux of the same shape which again conform to the hypotheses of the lemma, and by induction the same is true for any number of slides, which proves $X \cup S \cup Y \approx X \cup T \cup Y$. □

Let us now introduce a convention which will save on verbiage later.

Definition. A *miniature* shape or tableau is one having m cells, where $m = 3$ in the context of the unshifted theory, and $m = 4$ in the shifted context.

Proposition 2.2. *Each dual equivalence class of miniature tableaux consists of either one or two elements. A miniature tableau T belongs to a two-element class if and only if its reading word belongs to list A below, for the unshifted case, or list B for the shifted case. In either case, the other tableau T' dual equivalent to T differs from T by exchange of the entries x and y. In the lists, 1, 2, 3, 4 stand for the entries of the tableau in increasing order.*

List A: $x1y$ $x3y$

List B: $1x2y$ $x12y$ $1x4y$ $x14y$ $4x1y$ $x41y$ $4x3y$ $x43y$

Proof. For each entry on list A, write down all pairs of tableaux of common shape whose reading words match the entry. Then verify that any slide on such a pair yields another such pair—a trivial but tedious exercise which we leave for the skeptical reader. Likewise for list B, using shifted tableaux. This proves that the members of each such pair are dual equivalent, as claimed.

Next suppose S and T are distinct dual equivalent miniature tableaux, say of shape λ. Choose any tableau X such that $\operatorname{sh} X$ is normal and λ extends $\operatorname{sh} X$. Then $S' = j^X(S)$ and $T' = j^X(T)$ result from the same slide sequence applied to S and T respectively, hence have the same shape μ, which is normal. Furthermore, $S \approx T$ implies $\mathrm{V}: j^X(S) = \mathrm{V}: j^X(T) = Y$, say: Thus $S = j_Y(S')$ and $T = j_Y(T')$. In particular $S' \neq T'$.

Now S' and T', as distinct miniature tableaux of the same normal shape, can only be

$$\left\{ \begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 & \multicolumn{1}{c}{} \\ \cline{1-1} \end{array}\,, \quad \begin{array}{|c|c|}\hline 1 & 3 \\ \hline 2 & \multicolumn{1}{c}{} \\ \cline{1-1} \end{array} \right\}$$

in the unshifted case, or

$$\left\{ \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline \multicolumn{1}{c|}{} & 4 & \multicolumn{1}{c}{} \\ \cline{2-2} \end{array}\,, \quad \begin{array}{|c|c|c|}\hline 1 & 2 & 4 \\ \hline \multicolumn{1}{c|}{} & 3 & \multicolumn{1}{c}{} \\ \cline{2-2} \end{array} \right\}$$

in the shifted case. Thus they form a pair whose reading words match an entry in list A or B, in fact the entry $x1y$ from A or $x12y$ from B. Since we have already seen that any slide carries such a pair into another such pair, S and T are such a pair, which proves the proposition. □

Definition. An *elementary dual equivalence* is one of the form $X \cup S \cup Y \approx X \cup T \cup Y$ as in Lemma 2.1, where $S \approx T$ is a dual equivalence of distinct miniature tableaux.

As an aid to understanding we pause to summarize in informal language the content of this definition, taken together with Proposition 2.2. An unshifted elementary dual equivalence amounts to an exchange in a tableau of consecutive entries x and y that are separated in the cross-order (the order ascending upward and to the right) by an entry w consecutive with $\{x, y\}$. A shifted elementary dual equivalence also amounts to such an exchange, but where additionally w is preceded in the cross-order by an entry v consecutive with $\{w, x, y\}$.

We are now prepared to reduce all dual equivalences to elementary ones. We begin with another simple lemma, then use it to proceed from the special case of normal-shape tableaux to the general case.

Lemma 2.3. *Let $U \approx V$ be an elementary dual equivalence. Then applying any slide to U and V respectively yields U' and V' such that $U' \approx V'$ is elementary.*

Proof. This follows immediately from the description given in the proof of Lemma 2.1 of the action of a slide on the tableaux $U = X \cup S \cup Y$ and $V = X \cup T \cup Y$. □

Proposition 2.4. *Let S and T be any two tableaux of the same normal shape λ. Then S and T are connected by a chain of elementary dual equivalences.*

Proof. We proceed by induction on $|\lambda|$. The case $\lambda = \emptyset$ is trivial. If $|\lambda| = n > 0$, we can assume by induction that for any lower-right corner cell c in λ, there are elementary dual equivalences connecting all tableaux having entry n in cell c. In particular, if there is only one such c, we are done.

Otherwise, let c and c' be two lower-right corner cells. It suffices to prove that there exist tableaux S and T of shape λ having entry n in cells c and c' respectively, such that $S \approx T$ is an elementary dual equivalence.

In the unshifted case, choose from among all cells lying strictly between c and c' in the cross-order a lower-right-most one c''. Such a cell clearly exists in a normal shape. Put $n-2$ in c'', n and $n-1$ in c and c' in either order, and $1, \ldots, n-3$ in the remaining cells arbitrarily to form S and T. The exchange of $n-1$ and n is then an elementary dual equivalence $S \approx T$.

In the shifted case, proceed just as above, but choose also a lower-right-most cell c''' preceding c'' in the cross-order. Again the fact that λ is a normal shape—here a normal *shifted* shape—guarantees this cell's existence. Filling in the cells as above, with now also $n-3$ in c''', we get a shifted elementary dual equivalence $S \approx T$. □

Corollary 2.5. *All tableaux of any given normal shape λ are dual equivalent.*

Theorem 2.6. *Tableaux S and T are dual equivalent if and only if they are connected by a chain of elementary dual equivalences.*

Proof. The 'if' part is obvious.

For 'only if', assume $S \approx T$. Then $\operatorname{sh} S = \operatorname{sh} T = \lambda$, say; let X be a tableau of normal shape such that λ extends $\operatorname{sh} X$. As in the proof of Proposition 2.2, let $S' = j^X(S)$, $T' = j^X(T)$, and $Y = \mathrm{V}:j^X(S) = \mathrm{V}:j^X(T)$.

By Proposition 2.4, there is a chain $S' = R_0' \approx R_1' \approx \cdots \approx R_k' = T'$ of elementary dual equivalences connecting S' and T'. By Lemma 2.3, $S = R_0 \approx R_1 \approx \cdots \approx R_k = T$ is a chain of elementary dual equivalences, where $R_i = j_Y(R_i')$ for each i.
□

In the remainder of this section, we develop several alternative properties characterizing the relation of dual equivalence. Readers wishing to move quickly to applications may stop after Corollary 2.9, which is the last result used in subsequent sections, except Section 3, which uses Theorem 2.12. The remaining material is included in order to explicate the relationship between dual equivalence and the foundations of *jeu-de-taquin* theory. This material also contains the reason for the name 'dual equivalence.'

Lemma 2.7. *Let S and T be tableaux such that* $\operatorname{sh} T$ *extends* $\operatorname{sh} S$. *Then* $\mathrm{V}{:}\jmath_T(S) = \jmath^S(T)$ *and* $\mathrm{V}{:}\jmath^S(T) = \jmath_T(S)$.

Proof. Let P stand for the poset $\operatorname{sh} S \cup \operatorname{sh} T$; thus $S \cup T$ describes an order-preserving bijective function $\phi : P \to \{1, \ldots, n\}$, where $n = |P|$. If now θ is any such bijection, define $r_{i,i+1}\theta$ to be either θ, in case $\theta^{-1}(i) < \theta^{-1}(i+1)$, or else $(i, i+1) \circ \theta$, otherwise. It is plain that $r_{i,i+1}\theta$ is an order-preserving bijection from P to $\{1, \ldots, n\}$.

Let $k = |\operatorname{sh} S|$, $l = |\operatorname{sh} T|$. By interpreting the definition of *jeu-de-taquin* slides appropriately in terms of the operations $r_{i,i+1}$, we see that

$$\begin{aligned}(\mathrm{V}{:}\jmath_T(S)) \cup \jmath_T(S) = r_{l,l+1} r_{l+1,l+2} \cdots r_{n-1,n} \\ \cdots r_{2,3} r_{3,4} \cdots r_{k+1,k+2} \\ \cdots r_{1,2} r_{2,3} \cdots r_{k,k+1}\phi.\end{aligned}$$

Likewise,

$$\begin{aligned}\jmath^S(T) \cup (\mathrm{V}{:}\jmath^S(T)) = r_{l,l+1} r_{l-1,l} \cdots r_{1,2} \\ \cdots r_{n-2,n-1} r_{n-3,n-2} \cdots r_{k-1,k} \\ \cdots r_{n-1,n} r_{n-2,n-1} \cdots r_{k,k+1}\phi.\end{aligned}$$

Now $r_{i,i+1}$ and $r_{j,j+1}$ commute if $|i-j| > 1$. These commutation relations imply that the preceding two expressions are equal, proving the lemma. □

It should be noted that the concepts of shapes, tableaux, and slides have natural interpretations with any poset in place of the plane or the shifted plane. Lemma 2.7 remains valid in this more general setting, as is clear from the proof just given.

Corollary 2.8. *Let S, T and X be tableaux such that $S \approx T$ and* $\operatorname{sh} S = \operatorname{sh} T$ *extends* $\operatorname{sh} X$. *Then $\jmath_S(X) = \jmath_T(X)$. In other words, dual equivalent slide sequences have the same effect on any tableaux X. Note that this applies to both forward and reverse slide sequences by the anti-statement of this corollary.*

Proof. By Lemma 2.7, $\jmath_S(X) = \mathrm{V}{:}\jmath^X(S)$ and $\jmath_T(X) = \mathrm{V}{:}\jmath^X(T)$. It follows directly from the definition of dual equivalence that $S \approx T$ implies $\mathrm{V}{:}\jmath^X(S) = \mathrm{V}{:}\jmath^X(T)$. □

Corollary 2.9. *With the hypotheses of Corollary* 2.8, $\mathrm{V}:\jmath_S(X) \approx \mathrm{V}:\jmath_T(X)$.

Proof. By Lemma 2.7, the conclusion is the same as $\jmath^X(S) \approx \jmath^X(T)$, and this is clear from the definition of dual equivalence. □

Theorem 2.10. *Let S and T be tableaux and suppose* $\operatorname{sh} S = \operatorname{sh} T = \lambda$. *Then the following conditions are equivalent*:

(1) $S \approx T$.

(2) $\mathrm{V}:\jmath^X(S) = \mathrm{V}:\jmath^X(T)$ *for all X such that λ extends* $\operatorname{sh} X$.

(3) $\mathrm{V}:\jmath^X(S) = \mathrm{V}:\jmath^X(T)$ *for some X such that λ extends* $\operatorname{sh} X$ *and* $\operatorname{sh} X$ *is normal.*

(4) $\mathrm{V}:\jmath_X(S) = \mathrm{V}:\jmath_X(T)$ *for all X such that* $\operatorname{sh} X$ *extends λ.*

(5) $\mathrm{V}:\jmath_X(S) = \mathrm{V}:\jmath_X(T)$ *for some X such that* $\operatorname{sh} X$ *extends λ and* $\operatorname{sh} X$ *is anti-normal.*

(6) $\jmath_S(X) = \jmath_T(X)$ *for all X such that λ extends* $\operatorname{sh} X$.

(7) $\jmath_S(X) = \jmath_X(X)$ *for some X such that λ extends* $\operatorname{sh} X$ *and* $\operatorname{sh} X$ *is normal.*

(8) $\jmath^S(X) = \jmath^T(X)$ *for all X such that* $\operatorname{sh} X$ *extends λ.*

(9) $\jmath^S(X) = \jmath^T(X)$ *for some X such that* $\operatorname{sh} X$ *extends λ and* $\operatorname{sh} X$ *is anti-normal.*

(10) *There exist tableaux S' and T' of the same normal shape μ and a tableau Y of shape extending μ such that $S = \jmath_Y(S')$ and $T = \jmath_Y(T')$.*

(11) *There exist tableaux S' and T' of the same anti-normal shape μ and a tableau Y for which μ extends* $\operatorname{sh} Y$ *such that $S = \jmath^Y(S')$ and $T = \jmath^Y(T')$.*

Proof. (1)⇒(2) follows directly from the definition of dual equivlaence. (2)⇒(3) is trivial. (3)⇒(10) follows by taking $S' = \jmath^X(S)$, $T' = \jmath^X(T)$, and $Y = \mathrm{V}:\jmath^X(S) = \mathrm{V}:\jmath^X(T)$. (10)⇒(1) because of Corollary 2.5. (6) and (7) are equivalent to (2) and (3) by Lemma 2.7.

Finally, (4), (5), (8), (9), and (11) are the anti-statements of (2), (3), (6), (7), and (10), so they are all equivalent to (1) as (1) is its own anti-statement. □

There is one more, rather deeper, characterization of dual equivalence, this time in terms of the reading words of the tableaux. To explain this, it is convenient to identify *words* with *permutation tableaux* in the manner introduced by Schützenberger. Namely, a permutation tableau is one whose shape is an anti-chain. Any permutation can be the reading word of such a tableau and so we use the tableau to represent the word.

Given a shape λ, it is easy to describe a sequence of slides that carries any tableau T of shape λ to its reading word (as a permutation tableau), by first separating the rows of T, then spreading each row into individual cells. The action of such a slide sequence does not depend on the actual contents of T; thus the slides vacate cells in a fixed sequence that is the same for all T. The reverse of this last sequence is then a sequence of slides that carries the reading words back to the corresponding tableaux. Now if two tableaux are dual equivalent, then so are the results of any slide sequence applied to the two tableaux. In the present context this proves the following.

Lemma 2.11. *Tableaux S and T of the same shape λ are dual equivalent if and only if their reading words (regarded as permutation tableaux) are.*

Now we shall determine when the two reading words are dual equivalent. The result, incidentally, explains how to compute lists A and B of Proposition 2.2, which were introduced there by fiat.

Theorem 2.12. *Permutation tableaux S and T are dual equivalent if and only if $\mathrm{R}:\square\leftarrow S = \mathrm{R}:\square\leftarrow T$ (or $\mathrm{R}:\text{◹}\leftarrow S = \mathrm{R}:\text{◹}\leftarrow T$ in the shifted theory).*

Proof. We prove 'only if' first. For the unshifted case, suppose $S \approx T$. $\mathrm{R}:\square\leftarrow S$ records the shape sequence $\lambda_1 \subset \cdots \subset \lambda_n$, where $\lambda_k = \mathrm{sh}\,\square\leftarrow(s_1 \cdots s_k)$ (here we identify S with the word $s_1 \cdots s_n$). Since $S \approx T$, a sequence of slides that brings $s_1 \cdots s_k$ to normal shape and leaves the other $n-k$ cells put will bring $t_1 \cdots t_k$ to the same normal shape λ_k. Hence $\mathrm{R}:\square\leftarrow T = \mathrm{R}:\square\leftarrow S$.

For the shifted case, the same argument shows that $\mathrm{R}:\text{◹}\leftarrow S$ and $\mathrm{R}:\text{◹}\leftarrow T$ have the entries $1, \ldots, n$ in the same cells. However, in addition, those off-diagonal entries representing shifted Schensted insertions whose 'bump paths' contain a diagonal cell are circled. We must verify that k is circled in $\mathrm{R}:\text{◹}\leftarrow S$ if and only if k is circled in $\mathrm{R}:\text{◹}\leftarrow T$. At this point we simply state the relevant facts, presuming that the reader who is familiar with the simulation of shifted insertion by slides can check them easily.

Put $S_k = \text{◹}\leftarrow(s_1 \ldots s_k)$ and consider the insertion of s_{k+1} into S_k, i.e., the formation of $\text{◹}\leftarrow(S_k \oplus s_{k+1})$ by slides. s_{k+1} is initially placed at the end of an empty row directly above the top row of S_k. If we number the empty cells in this row $1, \ldots, m$, then slides are performed into these cells in the order $m, m-1, \ldots, 1$. Let d be the number of diagonal cells in S_k. The diagonal cell in the last row of S_k will be:

(1) vacated by the slide into cell $d+1$ of the empty row, in case the bump path for insertion of s_{k+1} does not touch the diagonal, or

(2) vacated by the slide into cell d of the empty row, in case the bump path touches the diagonal and continues to the right through at least one more column, or

(3) not vacated at all, in case the bump path touches the diagonal and ends there.

If $S \approx T$, then the slides forming $\text{◹}\leftarrow(S_k \oplus s_{k+1})$ will vacate the same cells as those forming $\text{◹}\leftarrow(T_k \oplus t_{k+1})$, so we will have the same case (1), (2) or (3) for each tableau. But the circling of entry $k+1$ is determined by the case (circled in case (2), otherwise not), and this means $\mathrm{R}:\text{◹}\leftarrow S$ and $\mathrm{R}:\text{◹}\leftarrow T$ have the same entries circled.

Finally we prove the 'if' direction. Let S be a permutation and let $\lambda = \operatorname{sh} \square \leftarrow S$ (or sh ◹$\leftarrow S$). Fix a tableau X describing a slide sequence carrying S into the shape λ, i.e., any tableau of normal shape such that $\operatorname{sh} S$ extends $\operatorname{sh} X$. Put $Y = \mathrm{V}:\jmath^X(S)$. We will have the same λ and Y for any permutation S' such that $S' \approx S$, and since all tableaux of the normal shape λ are dual equivalent, we see that $\jmath^X$ and $\jmath_Y$ define inverse bijections between the dual equivalence class of S and the set of tableaux of shape λ. By the Schensted correspondence, $\jmath^X$ (which computes $\square \leftarrow S$ or ◹$\leftarrow S$) also defines a bijection from the set $R(S)$ of S' with $\mathrm{R}:\square \leftarrow S' = \mathrm{R}:\square \leftarrow S$ (or $\mathrm{R}:$◹$\leftarrow S' = \mathrm{R}:$◹$\leftarrow S$) to the set of tableaux of shape λ. By the 'only if' part of the theorem, $R(S)$ contains the dual equivalence class of S, and so we see that the two sets must be equal. □

It is this last theorem that justifies the term 'dual equivalence'. The equivalence to which dual equivalence is 'dual' is *jeu-de-taquin* equivalence, that is, the relation of being connected by a sequence of slides. Two tableaux are *jeu-de-taquin* equivalent if and only if their reading words have the same Schensted insertion tableau. Our theorem says that, 'dually,' dual equivalence reflects the same concept for the recording tableaux.

Abstractly, the Schensted insertion tableau for a permutation w is merely a representative of its *jeu-de-taquin* equivalence class, and the recording tableau a symbol for its dual equivalence class. We define two permutations to be *shape equivalent* if they are *jeu-de-taquin* equivalent to tableaux of the same normal shape. Then the Schensted correspondence amounts to this: each permutation is determined by its equivalence class and its dual equivalence class; and an equivalence class and dual equivalence class both contain some permutation if and only if they are contained in the same shape class.

The advantage of this viewpoint is that one can describe a 'Schensted correspondence' not only for permutations, but for tableaux of any given shape λ, using dual equivalence classes in place of recording tableaux.

Theorem 2.13. *Let λ be a shape and let T_λ be the set of tableaux of shape λ. Then within each shape class, each jeu-de-taquin equivalence class meets each dual equivalence class in a unique $T \in T_\lambda$.*

Proof. Fix a tableau X such that $\operatorname{sh} X$ is normal and λ extends $\operatorname{sh} X$. We are to show that if $S, U \in T_\lambda$ belong to the same shape class then there is a unique $T \in T_\lambda$ such that $\jmath^X(T) = \jmath^X(S)$ and $T \approx U$. Let $Y = \mathrm{V}:\jmath^X(U)$. We have inverse bijections $\jmath^X$ and $\jmath_Y$ between the dual equivalence class of U and the set of tableau of shape $\mu = \operatorname{sh} \jmath^X(U) = \operatorname{sh} \jmath^X(S)$. Hence $T = \jmath_Y(\jmath^X(S))$ is the unique tableaux with the properties we require. □

We close this section with some remarks on the relevance of dual equivalence to the foundations of *jeu-de-taquin* theory. There are two 'fundamental theorems' of *jeu-de-taquin,* due to Schützenber for the unshifted theory and to Worley and Sagan for the shifted version.

The first fundamental theorem is that the result of a sequence of slides carrying a tableau T to normal shape depends only on T and not on the choice of sequence. The resulting tableau is what we call $\square \leftarrow T$ or ◹ $\leftarrow T$. In view of Corollary 2.8, which depends only on the very general Lemma 2.7, this first fundamental theorem is nothing but the fact (Corollary 2.5) that all tableaux of a given normal shape are dual equivalent, which we proved by direct elementary methods. Thus we have given here a new proof of the first fundamental theorem, which may be in some ways more illuminating than previous ones.

The second fundamental theorem of *jeu-de-taquin* is that the number of tableaux of shape λ going over under *jeu-de-taquin* to a given tableau T of normal shape μ depends only on λ and μ and not on T. The applications of *jeu-de-taquin* to the theory of symmetric functions are based on this theorem, which is normally proved using an operation known as *evacuation.* In our context, the second fundamental theorem is merely a corollary to Theorem 2.13, and we get the additional information that the number in question equals the number of dual equivalence classes contained in the shape class μ in T_λ.

The following proposition characterizes those shapes which can replace normal shapes in the first fundamental theorem. Although it is not hard to prove, we omit the proof as it is somewhat off our central purpose.

Proposition 2.14. *An unshifted shape λ has the property that all tableaux of shape λ are dual equivalent if and only if λ is either normal or anti-normal. A shifted shape λ has this property if and only if λ is either normal, anti-normal, or any shifted shape having a unique upper-left and a unique lower-right corner.*

3. Dual Knuth relations for shifted insertion

Knuth relations are certain elementary transformations on a permutation that preserve the corresponding Schensted insertion or recording tableau. In the cases of the unshifted insertion and recording tableaux and the shifted insertion tableau, there are theorems, due to Knuth [4] in the unshifted theory and to Worley [12] and Sagan [7] in the shifted theory, saying that two permutations have the same insertion or recording tableau if and only if they are connected by a chain of the appropriate Knuth relations.

We now use dual equivalence to deduce the Knuth relations (called *dual* Knuth relations) for preservation of the recording tableau. In the shifted theory, this result is new.

Theorem 3.1. *Let v and w be permutations. Interpret entries from lists A and B of Proposition 2.2 as dual Knuth relations as follows: if the list entry occurs as a* *subsequence* (*not necessarily adjacent*) *of* v, *where* 1, 2, 3, 4 *stand for any consecutive numbers* $m+1$, $m+2$, $m+3$, $m+4$, *the corresponding dual Knuth relation consists of exchanging* x *and* y. *Then* $\mathrm{R}:\square\leftarrow v=\mathrm{R}:\square\leftarrow w$ *if and only if* v *and* w *are connected by dual Knuth relations from list* A. $\mathrm{R}:\triangledown\leftarrow v=\mathrm{R}:\triangledown\leftarrow w$ *if and only if* v *and* w *are connected by dual Knuth relations from list* B.

Proof. Theorem 2.12 shows that v and w have the same recording tableau if and only if they are dual equivalent (regarded as permutation tableaux). This happens if and only if they are connected by elementary dual equivalences, and for permutation tableaux, these amount to the above dual Knuth relations. $\square$

In [2] it is shown that $\mathrm{R}:\triangledown\leftarrow w$ is the insertion tableau $\triangledown\overset{\mathrm{m}}{\leftarrow}w^{-1}$ for a process of *shifted mixed insertion* defined there. Reinterpreting the dual Knuth relations for their effect on the inverse of a permutation, we have the following.

Corollary 3.2. *Let* v *and* w *be permutations. Then* $\triangledown\overset{\mathrm{m}}{\leftarrow}v=\triangledown\overset{\mathrm{m}}{\leftarrow}w$ *if and only if* v *and* w *are connected by Knuth relations of the following forms*:

$$acbd\sim cabd,\qquad bacd\sim bcad,\qquad dacb\sim dcab,\qquad dbac\sim dbca,$$

where a, b, c, d *form a subsequence of adjacent numbers*, $a<b<c$, *and* $d<b$.

4. Promotion and generalized staircases

In this section we apply dual equivalence to the study of the *promotion* operation on tableaux. Especially, we are interested in shapes for which tableau promotion 'commutes' with elementary dual equivalences in an appropriate sense. For these shapes, properties of promotion which hold for one tableau can be carried over to dual equivalent ones.

The shapes to which our reasoning applies can be readily classified. We call them *generalized staircases*. They turn out to be precisely those shapes about which it is known or conjectured that the operation of *total promotion* is the identity (or in one case, the transpose). Consequently we arrive at a new unified proof of this property for the known shapes and for the conjectured ones.

The results of this section also form the basis of our study of the Green–Edelman correspondence and Proctor's conjecture in Section 5.

Definition (Schützenberger [8]). Let T be a tableau and let $n=|\mathrm{sh}\,T|$. The *promotion step* operation p is defined on T as follows: delete the entry n from its cell in T; perform a slide into that cell on the rest of T; and fill the cell vacated by the slide with a new least entry 0. Then add 1 to all entries to regain the standard $1,\ldots,n$. The result is $p(T)$. Note that p is an invertible operation and that p^{-1} is the anti-operation of p. The operation p^n is called *total promotion.*

We will also need the following related operation.

Definition. Let T be as before. The *exhaustion* operation e is defined on T as follows: Delete entry n and perform a slide into its cell. Then do the same to entries $n-1$, $n-2, \ldots, 1$ in succession, so that every cell in sh T is ultimately vacated. Then $e(T)$ is the tableau of shape sh T whose entries record the order in which cells were vacated by the above slides.

Exhaustion e and its anti-operation e^* are involutions (see [8]). Like Lemma 2.7, this is a general statement true of posets that are not necessarily tableau shapes. In fact, it follows from Lemma 2.7, applied to the poset $P = T \cup S$ where S is a chain of $|T|$ elements all greater than every element of T.

It is not hard to deduce from the definitions that the total promotion operator p^n is equal to $e^{*-1} \circ e = e^* \circ e$. For a given shape λ, then, the properties $p^n = \mathrm{id}$ and $e = e^*$ are equivalent. For a normal shape, e^* reduces to the standard *evacuation* operation from the Schützenberger and Worley theories of *jeu-de-taquin.*

Definition. A shape λ is a *brick* if for all tableaux X, Y of shape λ, $X \approx Y \Rightarrow e(X) \approx e(Y)$. Otherwise, λ is a *stone.*

Proposition 4.1. *The unshifted miniature bricks are these:*

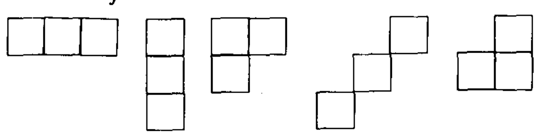

The other unshifted miniature shapes are stones. The shifted miniature bricks are these:

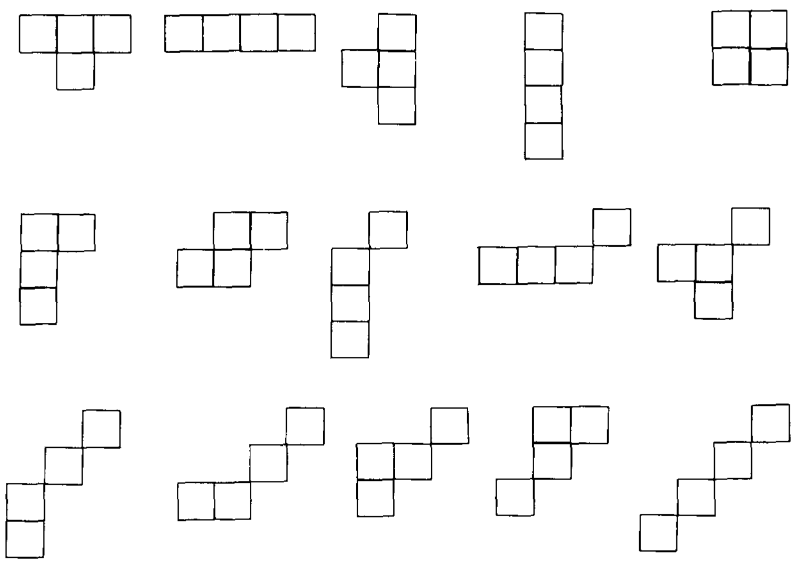

The other shifted miniature shapes are stones. Here the diagrams represent shapes up to isomorphism as posets, so that for instance the last shape in the shifted table stands for any permutation shape, even though the four cells might not be diagonally adjacent.

Proof. For each shape on the lists it is routine to verify that e preserves dual equivalence. For each shape not on the lists, one easily finds counterexamples. The work can be simplified by noting that only two-element dual equivalence classes need be checked.

For example,

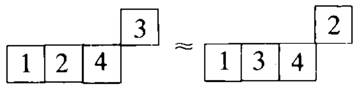

is the only such dual equivalence class of its shape, and e exchanges these two tableaux, so this shape is a brick.

On the other hand,

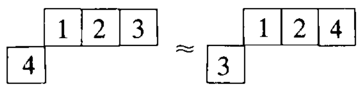

but e carries these to nondual equivalent tableaux, so this shape is a stone.

Note incidentally that normal shapes and permutation shapes are necessarily bricks. □

We can now describe a class of tableaux for which promotion and dual equivalence interact nicely.

Definition. A *generalized staircase* is a rookwise-connected non-empty shape in which every miniature final segment is a brick and every miniature initial segment is an anti-brick.

Note that although an anti-brick may be a stone (since $e \neq e^*$ in general), the preceding definition as a whole is anti-invariant, so an anti-generalized staircase is the same as a generalized staircase.

Proposition 4.2. *The unshifted generalized staircases are the following.*

(1) *The staircase shape $A_l = (l, l-1, \ldots, 1)$ and its anti-shape A_l^*, for $l \geq 1$. We have*

$$|A_l| = \binom{l+1}{2}.$$

(2) *The rectangle shape $R_{l,m} = (m, \ldots, m)$ with l rows, for $l, m \geq 1$. We have $|R_{l,m}| = lm$.*

The shifted generalized staircase are listed next.

(1) *The 'double staircase' shape*

$$D_{l,m} = (l, l-1, \ldots, 1) + (m, m-1, \ldots, 1) = (l+m, l+m-2, \ldots)$$

and its anti-shape $D^*_{l,m}$, *for* $l \geq 1$, $0 \leq m \leq l$. *We have*

$$|D_{l,m}| = \binom{l+1}{2} + \binom{m+1}{2}.$$

(2) *The trapezoid shape* $T_{l,m} = (m+l-1, m+l-3, \ldots, m-l+1)$ *and its anti-shape* $T^*_{l,m}$, *for* $m \geq l \geq 1$. *We have* $|T_{l,m}| = lm$.

(3) *The shape* $U_3 = A_2$, *regarded as a shifted shape, and* $U^*_3 = A^*_2$.

(4) *The shape* $U_4 = R_{2,2}$, *regarded as a shifted shape.*

(5) *The shifted shape* $U_5 = (3, 2)$ *and its anti-shape* U^*_5.

Proof. First we check that these are generalized staircases. The U_j are trivial. The rest are either normal or anti-normal and it suffices to check the normal ones by anti-invariance. Their initial segments are normal, hence anti-bricks, so only final segments need be examined. For A_l these are either permutations or are normal. For $R_{l,m}$ they are anti-normal. Final segments of $T_{l,m}$ can be the first, second, seventh, ninth, or any of the last four entries in the table of shifted bricks in Proposition 4.1, and no others. Final segments of $D_{l,m}$ can be any of the first, third, fourth, sixth, eighth, tenth, eleventh, or any of the last three entries, and no others.

Now we establish that these are the only possibilities. Let λ be a generalized staircase. Let c be the upper-right-most cell in λ. We can assume λ is at least miniature, since all smaller shapes are on the list.

Suppose the cell to the left of c is in λ. Then λ has a unique upper-left corner cell. Otherwise it would have

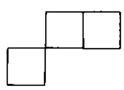

as an initial segment, which is an anti-stone, and thus forbidden, in the unshifted case. Since it is connected, λ would also have one of

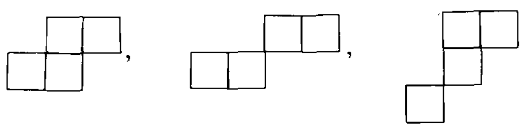

as an initial segment, which are shifted anti-stones.

In the unshifted case, λ is now normal. In the shifted case,

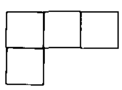

is an anti-stone, so λ is either normal or has only two cells in the top row. In the latter case, since

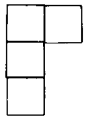

is an anti-stone, λ is U_4 or U_5^*.

By anti-invariance, if the cell below c is in λ, then λ is anti-normal or is U_4 or U_5 in the shifted case. But either the cell to the left of c or the cell below c must be in λ by connectedness, so λ is normal or anti-normal or U_4. We can assume without loss of generality that λ is normal.

In the unshifted case, if the cell below c is in λ then λ is anti-normal and is an $R_{l,m}$. Otherwise the lower-right corners of λ form a staircase because all disconnected unions of a single-cell and a two-cell shape are stones. Thus λ is an A_l.

In the shifted case, let r be the first row which is not followed by a row exactly two cells shorter (i.e., stepping back one cell at the end). Let x be the last cell in row r. If the next row is one cell shorter than r, then the cell below x is in λ. This implies, by the same argument used to show λ is normal or anti-normal, that either all cells below x extending down to the diagonal are in λ, and λ is a $D_{l,m}$, or else r is the next-to-last row, so that x, the cell below x, and the cell to the left of that form a final segment

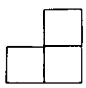

But then λ is U_5 because the final segments

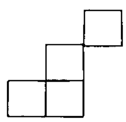 and 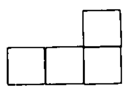

are stones.

Otherwise, neither the cell below x nor the cell to the left of that cell is in λ. But the only brick having two cells at the end of the top row and neither of the cells below them is

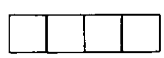,

which shows that r is the last row and λ is a $T_{l,m}$. $\square$

The shapes of type T and type D overlap in two cases. These will be of special interest in Section 5, so we will honor them with their own names.

Definition. For $T_{l,l} = D_{l,l-1} = (2l-1, 2l-3, \ldots, 1)$ write B_l. We have $|B_l| = l^2$. For $D_{l,l} = T_{l,l+1} = (2l, 2l-2, \ldots, 2)$ write C_l. We have $|C_l| = l(l+1)$.

After all these definitions and classifications, we finally come to their purpose, which is the following theorem.

Theorem 4.3. *Let λ be a generalized staircase and let $n = |\lambda|$. We write $(t \bmod n)$ for the element of $\{1, \ldots, n\}$ congruent to $t \pmod n$. Let $S \approx T$ be an elementary dual equivalence of tableaux of shape λ, and let $\{j, j+1, \ldots, k\}$ be the entries of the miniature segment involved in the elementary dual equivalence. Then for $(-t \bmod n) \notin \{j, \ldots, k-1\}$, $p^t(S) \approx p^t(T)$ is an elementary dual equivalence involving the segment $\{(j+t \bmod n), \ldots, (k+t \bmod n)\}$ (the restriction on t makes this a genuine segment).*

Proof. For $t = 1$ and $k < n$, and for $t = -1$ and $j > 1$, the result follows from the analysis of the action of slides used to prove Lemma 2.1 and 2.3. Iterating this implies the result in turn for any sequence of consecutive allowed values of t containing a value t_0, once it is proved for t_0.

All the allowed values of t fall into segments separated by gaps of length $m-1$, where $m = k - j + 1$ is the size of a miniature shape. To carry the result across the gaps, we must show that if $\{j, \ldots, k\} = \{n-m+1, \ldots, n\}$ then $p^m(S) \approx p^m(T)$ is an elementary dual equivalence involving $\{1, \ldots, m\}$. For negative t we need the anti-statement of this, but that follows by anti-invariance.

To prove the last assertion, we compute $p^m(S)$ and $p^m(T)$ by the following steps.

(1) Exhaust the final segments Y_S and Y_T containing $\{n-m+1, \ldots, n\}$ in S and T.

(2) Slide the remaining part X of S and T into the cells of sh Y_S = sh Y_T in the order given by the entries of $e(Y_S)$ and $e(Y_T)$.

(3) Record the order in which cells are vacated by the m slides in Step (2), for S and T.

(4) Perform forward slides into these cells in the order recorded in Step (3), filling the vacated cell with $0, -1, \ldots, -(m-1)$ in succession.

(5) Add m to all entries to regain the standard $1, \ldots, n$.

Now $e(Y_S) \approx e(Y_T)$ because their shape, a final segment of λ, is a brick. In Step (2) we are applying $j_{e(Y_S)}$ and $j_{e(Y_T)}$ to X, which is common to both S and T. By Corollary 2.8, both yield the same result. By Corollary 2.9, the orders recorded for S and T in Step (3) describe dual equivalent tableaux $Z_S \approx Z_T$. In Step (4) we are computing $e^{*-1}(Z_s) = e^*(Z_S)$ and $e^{*-1}(Z_T) = e^*(Z_T)$ (with the entries reduced by m) and they are dual equivalent because initial segments of λ are anti-bricks. Step (5) merely renormalizes the entries, and so we have shown $p^m(S) \approx p^m(T)$. □

As an immediate application we get the following.

Theorem 4.4. *For tableaux T of shape A_l or A_l^* we have $p^{|A_l|}(T) = T^t$, the transpose of T. For tableaux of any other generalized staircase shape λ except U_5 or U_5^* we have $p^{|\lambda|}(T) = T$, i.e., total promotion is the identity.*

Proof. Treating U_3 and U_4 as unshifted shapes, λ is either normal or anti-normal. Without loss of generality assume λ is normal. All tableaux of shape λ are connected by elementary dual equivalences, so in view of Theorem 4.3 it suffices to exhibit one tableau T_0 of shape λ such that $p^{|\lambda|}(T_0) = T_0$, or T_0^t if $\lambda = A_l$. For most of the generalized staircases it is a simple matter to verify that a T_0 which works is the tableau whose entries increase from left to right one row at a time.

The case $\lambda = D_{l,m}$ is harder than the others, so we consider it in detail. In this case, it is easier to show $e \circ e^*(T_0) = T_0$ than to compute $p^{|\lambda|}(T_0)$ directly. First we describe $e^*(T_0)$. Consider the nested normal shapes $\emptyset = \lambda_0 \subset \lambda_1 \subset \cdots \subset \lambda_l = D_{l,m}$, where for each i, $\theta_i = \lambda_i / \lambda_{i-1}$ is the rookwise connected path extending along the lower-right edge of λ_i from the diagonal to the end of the first row. Let S_i be the tableaux of shape θ_i in which the entries $1, 2, \ldots$ lie in descending rows up to some entry j, and the remaining entries $j, j+1, \ldots$ lie in ascending columns. Equivalently, S_i is the unique tableau of shape θ_i such that $\nabla \leftarrow S_i$ has just one row. From standard rules for shifted evacuation [2] we find that $e^*(T_0) = S_1 \cup \cdots \cup S_l$.

Now we compute $e(e^*(T_0))$. As we eliminate entries from θ_l, the segment $S_1 \cup \cdots \cup S_{l-1}$ is subjected to forward slides. The tableau S_l has the property that the sequence of these slides is described by S_l itself, so they yield $j_{S_l}(S_1 \cup \cdots S_{l-1}) = \mathrm{V} : j^{S_1 \cup \cdots \cup S_{l-1}}(S_l)$. Since sh $\nabla \leftarrow S_l$ is the single row $(l+m)$, the entries $1, \ldots, l+m$ of $e(e^*(T_0))$ occupy the first row. Furthermore, after eliminating the entries from θ_l, the remaining entries have been moved by slides into the normal subshape consisting of $D_{l,m}$ minus its first row and hence they form a tableau identical to $S_1 \cup \cdots \cup S_{l-1}$. It follows by induction that $e(e^*(T_0)) = T_0$. □

For $D_{l,m}$ this result is new; it had been conjectured by Stanley [11]. For the other shapes, the result was known from ad hoc proofs adapted to each shape.

Since the original manuscript of this paper was written, Kim and I [3] have proven the following.

Theorem 4.5. *If λ is a connected shape having the property described in Theorem 4.3, then λ is a generalized staircase.*

Theorem 4.6. *The only connected shapes λ for which $p^{|\lambda|} =$ identity are the generalized staircases $R_{l,m}$, $D_{l,m}$, $D_{l,m}^*$, $T_{l,m}$, $T_{l,m}^*$. The only symmetric shapes λ for which $p^{|\lambda|} =$ transpose are the generalized staircases A_l and A_l^*.*

5. The Greene–Edelman correspondence and the Proctor conjecture

In the paper [1], Greene and Edelman gave a remarkable bijection between standard tableaux of shape A_l and reduced expressions for the longest element w_{A_l} of the Coxeter group of type A_l. The fact that these two entities were equinumerous had been conjectured by Stanley and also proved by him, nonbijectively, through the intervention of certain symmetric functions he had concocted expressly for this purpose [10], see Section 6.

Stanley had also conjectured that the number of reduced expressions for the longest element w_{B_l} in the Coxeter group of type B_l was equal to the number of standard tableaux of square shape $R_{l,l}$. This last number is also (by hook formulas, or see [2] for a bijection) the number of standard tableaux of trapezoidal shape $B_l = T_{l,l}$. It was conjectured by Proctor [10] that for this case, the direct analogue of the Greene–Edelman correspondence would be a bijection between standard tableaux of shape B_l and reduced expresssions for w_{B_l}.

These tableau-to-reduced-expression correspondences involve the *promotion sequence* of the tableaux. In this section we apply the results of Section 4 to the study of promotion sequences in order to arrive in a unified manner at a new proof of the Greene–Edelman theorem and a proof of the Proctor conjecture. As a bonus, we discover that the methods used here apply not only to the shapes A_l and B_l, but also to C_l, for which we give an entirely new correspondence of the Greene–Edelman type.

The machinery common to all three cases is developed in the results leading up to Corollary 5.6. Proposition 5.8 shows that there are no other cases to consider. The remainder of the section is divided into subsections dealing with the peculiarities of each case.

Definition. Let T be a tableau of shape λ, with $|\lambda| = n$. Suppose that the lower-right corner cells of λ have been assigned fixed labels (conventionally, consecutive integers increasing in cross-order, i.e., upward and to the right). Then the *promotion sequence* of T is the doubly infinite sequence $\vec{p}(T) = (\ldots, r_{-1}, r_0, r_1, \ldots)$, where r_k is the label of the cell occupied by the greatest entry in the tableau $p^{n-k}(T)$. Also we write $\hat{p}(T)$ for the *short promotion sequence* $(r_1, \ldots, r_n)$. □

The seemingly curious indexing of $\vec{p}(T)$ can be explained this way: imagine repeated promotion and inverse promotion steps being performed on T, without renormalizing the entries to the standard $1, \ldots, n$ at each step. Then every integer k appears at some stage in a unique corner cell of λ, and r_k is the label of that corner cell. Note in particular that for $1 \leq k \leq n$, $(r_k, \ldots, r_n)$ is determined entirely by the final segment of T containing the entries $\{k, \ldots, n\}$, and is the short promotion sequence of that segment.

The importance of $\hat{p}(T)$ is that for λ a generalized staircase (ignoring U_5), it determines all of $\vec{p}(T)$. Specifically, by Theorem 4.4, $\vec{p}(T)$ is periodic with period n, except for $\lambda = A_l$, where $\vec{p}(T)$ is transpose-periodic; alternate periods have labels corresponding under transpose.

We now isolate those generalized staircase shapes for which the correspondence $T \mapsto \hat{p}(T)$ is amenable to detailed analysis.

Definition. A generalized staircase $\lambda \notin \{U_3, U_3^*, U_4\}$ is a *perfect stairwase* if it has the following property: if $\mu, \nu \subseteq \lambda$ are miniature final segments and $U \approx V$ is an elementary dual equivalence of tableaux of shape μ, then every tableau U' of shape ν having $\hat{p}(U') = \hat{p}(U)$ belongs to an elementary dual equivalence $U' \approx V'$ such that $\hat{p}(V') = \hat{p}(V)$. □

Theorem 5.1. *The members of the families A_l, B_l, and C_l are perfect staircases.*

Proof. When the final segment tableau U mentioned in the definition of a perfect staircase is determined uniquely by its short promotion sequence $\hat{p}(U)$, the condition in the definition is trivially satisfied. This observation pertains to all U of certain shapes, including every miniature final segment in A_l and most of those in B_l and C_l. The exceptions are as follows.

First, in B_l or C_l, if $a < b < c$ are consecutive, then the sequence $bcab$ occurs as $\hat{p}(U)$, for U either of two tableaux:

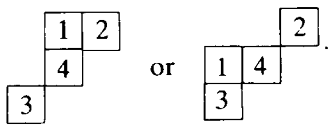

However, the sequence $bacb$ occurs for two tableaux which are dual equivalent to the above two (by switching the entries 2 and 3), so the condition in the definition is satisfied.

Second, in C_l only, the following pairs U, U' have $\hat{p}(U) = \hat{p}(U')$:

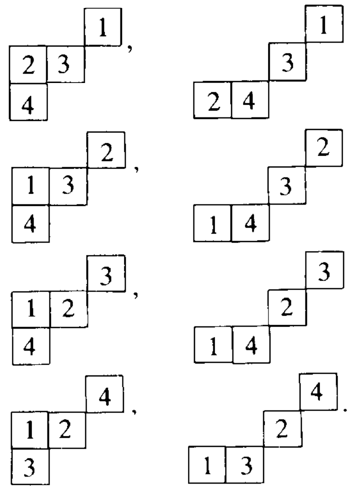

Fortunately, the first pair of tableaux are dual equivalent to the second pair and the third pair to the fourth, so the condition in the definition remains satisfied. □

Lemma 5.2. *Let λ be a perfect staircase. If $S \approx T$ is an elementary dual equivalence of tableaux of shape λ involving the segment with entries $\{j, \ldots, k\}$, then $\hat{p}(S)$ and $\hat{p}(T)$ differ only in the subsequence $r_j, \ldots, r_k$. Moreover, this subsequence in $\hat{p}(S)$ determines the corresponding subsequence in $\hat{p}(T)$.*

Proof. The part before 'moreover' follows directly from Theorem 4.3, even if λ is not perfect.

For the 'moreover', consider $S' = p^{n-k}(S) \approx p^{n-k}(T) = T'$. By Theorem 4.3, this is an elementary dual equivalence involving final segments U in S' and V in T'. By perfection, $\hat{p}(U)$ determines $\hat{p}(V)$, and these are the subsequences in question. □

Lemma 5.3. *Let λ be a perfect staircase, μ a final segment, U and V tableaux of shape μ with $\hat{p}(U) = X$, $\hat{p}(V) = Y$. Suppose there exists a tableau T such that* $\operatorname{sh} T \subseteq \lambda$, *$\mu$ extends* $\operatorname{sh} T$, *$T \cup U \approx T \cup V$, $\hat{p}(T \cup U) = WX$ and $\hat{p}(T \cup V) = WY$, for some sequence of labels W. Then if S is any tableau such that* $\operatorname{sh} S \subseteq \lambda$ *and μ extends* $\operatorname{sh} S$, *we have $\hat{p}(S \cup U) = AX$ and $\hat{p}(S \cup V) = AY$, for some sequence A.*

Proof. We can extend T any way we wish to the whole shape $\lambda \backslash \mu$ and preserve the hypotheses about it. Here the hypothesis that $\hat{p}(T \cup U)$ and $\hat{p}(T \cup V)$ agree in their initial subsequence remains true for the extension because of the assumption $T \cup U \approx T \cup V$.

Similarly we can extend S; this only strengthens the conclusion. Thus we assume

$$\operatorname{sh} T \cup \mu = \operatorname{sh} S \cup \mu = \lambda.$$

We can also assume λ is normal, since otherwise the result is trivial. Hence also $\operatorname{sh} S = \operatorname{sh} T$ is normal.

Now Proposition 2.4 provides us with a chain of elementary dual equivalences transforming T into S. Applying these same elementary dual equivalences to $T \cup U$ and $T \cup V$ transforms their short promotion sequences into those of $S \cup U$ and $S \cup V$. By Lemma 5.2 this transformation changes only the initial subsequence W, into a subsequence A that is determined by W alone. This proves the lemma. □

Proposition 5.4. *Let λ be a perfect staircase and let X and Y be two sequences of corner labels of the same length. Let $\hat{p}^{-1}(X)$ denote the set of tableaux U with* $\operatorname{sh} U$ *a final segment of λ and $\hat{p}(U) = X$. Suppose $\phi : \hat{p}^{-1}(X) \to \hat{p}^{-1}(Y)$ is a bijection such that for each $U \in \hat{p}^{-1}(X)$,* $\operatorname{sh} \phi(U) = \operatorname{sh} U$. *Suppose further that for all such U*

there exists a tableau T with shape contained in λ such that $\operatorname{sh} U$ *extends* $\operatorname{sh} T$, $T \cup U \approx T \cup \phi(U)$, *and*

$$\hat{p}(T \cup U) = WX, \qquad \hat{p}(T \cup \phi(U)) = WY$$

for some sequence W. Then whenever AXB, AYB are sequences of $|\lambda|$ labels, there is a bijection $\psi : \hat{p}^{-1}(AXB) \to \hat{p}^{-1}(AYB)$, where $\psi(S)$ is computed as follows: apply $|B|$ promotion steps; change the final segment U of size $|X|$ to $\phi(U)$; apply $|B|$ inverse promotion steps.

Proof. It is obvious that if the map ψ is well-defined then it is a bijection whose inverse is the analogous map corresponding to ϕ^{-1}.

To see that $\psi(S)$ is well-defined, we must first check that $p^{|B|}(S)$ has a final segment U to which ϕ applies. This is so because $\hat{p}(p^{|B|}(S)) = B'AX$ ends with X.

Secondly, we must check that $\hat{p}(\psi(S)) = AYB$. We observed that $\hat{p}(p^{|B|}(S)) = B'AX$, where by Theorem 4.4, B' is either B or else corresponds to B under transpose in case $\lambda = A_l$. Lemma 5.3 applied to U and $V = \phi(U)$ shows that changing U to $\phi(U)$ changes the short promotion sequence to $\hat{p}(p^{|B|}(\psi(S))) = B'AY$. By Theorem 4.4 again, $\hat{p}(\psi(S)) = AYB$. □

Definition. A pair $\{X, Y\}$ satisfying the hypotheses of Proposition 5.4 for a perfect staircase λ is called a *λ-relation*. The promotion sequences $\hat{p}(T)$ belong to tableaux T of shape λ are called *λ-words*. A set R of λ-relations is *full* if all λ-words are connected by transformations of the form $AXB \mapsto AYB$ for λ-relations $\{X, Y\} \in R$.

Proposition 5.5. *Let λ be a perfect staircase, μ a miniature final segment. If $U \approx V$ is an elementary dual equivalence of tableaux of shape μ then $\{\hat{p}(U), \hat{p}(V)\}$ is a λ-relation. All such λ-relations taken together form a full set.*

Proof. Let $X = \hat{p}(U)$, $Y = \hat{p}(V)$. If $U' \in \hat{p}^{-1}(X)$ then by perfection U' belongs to an elementary dual equivalence $U' \approx V'$ (necessarily unique) with $V' \in \hat{p}^{-1}(Y)$. Then $\phi(U') = V'$ defines the bijection required by Proposition 5.4; ϕ is bijective by perfection. The tableau T required by Proposition 5.4 for each U', V' can be taken to be empty since $U' \approx V'$. This shows $\{X, Y\}$ is a λ-relation.

If $S \approx T$ is an elementary dual equivalence of tableaux of shape λ, then Lemma 5.2 shows that $\hat{p}(S)$ and $\hat{p}(T)$ differ by a λ-relation of the kind just described. Since all tableaux of shape λ are connected by elementary dual equivalences, the set of these λ-relations is full. □

Corollary 5.6. *For tableaux T of a fixed shape $\lambda \in \{A_l, B_l, C_l\}$ the correspondence $T \mapsto \hat{p}(T)$ is injective. Thus it is bijective onto the set of λ-words.*

Proof. By Propositions 5.4 and 5.5, every λ-word belongs to the same number of tableaux. We need only show that this number is 1 for some specially chosen λ-word. We shall illustrate with the argument for $\lambda = B_l$ and leave the other, entirely similar, cases to the reader.

For B_l, take the corner labels to be $0, \ldots, l-1$ and consder the word

$$w = (l-1)(l-2)\cdots 21012\cdots(l-2)(l-1)(l-2)(l-3)$$
$$\cdots 21012\cdots(l-3)(l-2)\cdots 210121010.$$

It is easy to see that $w = \hat{p}(T_0)$ for the tableau T_0 whose entries increase left to right, one row at a time.

Suppose now $\hat{P}(T) = w$ for another tableau T of shape B_l. For the moment let us agree not to renormalize when promoting, so that entries retain their identity across promotion steps. In $p^{(l-1)^2}(T)$, the entries $1, \ldots, 2l-1$ form a final segment U with

$$\hat{p}(U) = (l-1)(l-2)\cdots 101\cdots(l-2)(l-1).$$

This can only happen if U has 1 and $2l-1$ in the first row, 2 and $2l-2$ in the second, and so on, with $l-1$ and $l+1$ in the $(l-1)$st row and l in the sole cell of the last row. For example, with $l=4$, U is

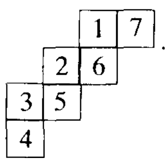

From the definition of promotion we see that U is *jeu-de-taquin* equivalent to the initial segment of size $2l-1$ in T, so this segment is $\triangledown \leftarrow U$, which is a single row of $2l-1$ cells, i.e., the first row of T. Now that we know the entries $1, \ldots, 2l-1$ occupy the first row, we conclude that the rest occupy a subshape of shape B_{l-1}, and thus $T = T_0$ by induction. $\square$

At this point we have fully developed the general machinery for studying the perfect staircases. In each case, the correspondence $T \mapsto \hat{p}(T)$ is a bijection from standard tableaux of a given shape λ to a certain set of words of length $|\lambda|$, and we have a description in principle of this set, in terms of a full set of λ-relations.

For our purposes we need to use full sets of λ-relations more concise than the ones provided by Proposition 5.5.

Definition. Let $\{X, Y\}$ be a λ-relation and suppose $\{WXZ, WYZ\}$ is another. We say the first of these λ-relations *implies* the second.

Lemma 5.7. *A set of λ-relations whose members imply all the members of a full set of λ-relations is itself full.*

Proof. Obvious. □

Before examining each of the families A_l, B_l, and C_l in detail, we justify our focus on them by showing that they are the only cases of interest.

Proposition 5.8. *Disregarding shapes with just one lower-right corner (for which all promotion sequences are trival), the only perfect staircases belong to the families A_l, B_l and C_l.*

Proof. We are to eliminate as not perfect the shapes $T_{l,m}$ with $m-1>l>1$, and $D_{l,m}$ with $l-1>m>0$. For the $D_{l,m}$ cases, both

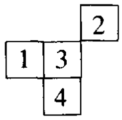

and

either 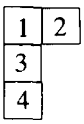

(when $m=l-2$) or 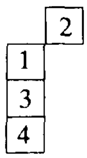

(when $m<l-2$)

have promotion sequence 1211 (labelling corner cells $1,2,\ldots$) but the first is dual equivalent to a tableau with $\hat{p}(V)=2111$, whereas the second is dual equivalent to a tableau with $\hat{p}(V')=1121$. For the $T_{l,m}$ cases, both

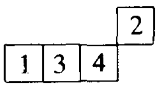

and 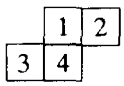

have promotion sequence 1211 but are dual equivalent to tableaux with promotion sequences 1121 and 2121 respectively. □

There are nonperfect generalized staircases λ, as well as shapes which are not generalized staircases, for which $\hat{p}$ is injective on tableaux of shape λ. Examples include $T_{2,4}$ and $D_{3,1}$. Among generalized staircases, one can show that $\hat{p}$ is not injective for $D_{l,m}$ when $l>m+2$ and for $T_{l,m}$ when $m>l+3$. Other cases remain open questions.

Problem 5.9. Is the correspondence $T\mapsto\hat{p}(T)$ injective on tableaux of shape λ, where $\lambda=D_{l,l-2}$, $T_{l,l+2}$, or $T_{l,l+3}$? If so, is there a meaningful combinatorial interpretation of the set of short promotion sequences that occur?

Table 1
The full set of Proposition 5.5. A_l-relations

	Dual equivalence	λ-relation
(1)	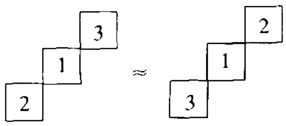	$bac,\ bca\quad (a<b<c)$
(2)	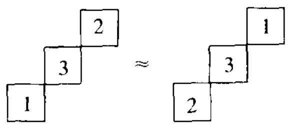	$acb,\ cab\quad (a<b<c)$
(3)	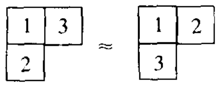	$bab,\ aba\quad (\lvert b-a\rvert=1)$

Table 2
A full set of A_l-relations

$ca,\ ac$	$(\lvert c-a\rvert>1)$
$bab,\ aba$	$(\lvert b-a\rvert=1)$

5.1. The case $\lambda = A_l$

In this subsection, we fix l and the unshifted shape $\lambda = A_l$. As corner labels we use the integers $1, \ldots, l$. Abusing notation we also let A_l denote the Coxeter group of type A_l and use $1, \ldots, l$ to denote the simple reflections. Then reduced expressions of elements of A_l are certain sequences of these integers, as are promotion sequences.

Representing A_l as the symmetric group S_{l+1} acting on $\{1, \ldots, l+1\}$ with the simple reflection i corresponding to the adjacent transposition $(i, i+1)$, the longest element w_{A_l} is the permutation $w_{A_l}(j) = l+2-j$.

There are three classes of elementary dual equivalences of shape a final segment in A_l. By Proposition 5.5 these lead to a full set of λ-relations. They are summarized in Table 1, in which a, b, c stand for integers from $\{1, \ldots, l\}$.

If $|c-a|>1$ then Proposition 5.4 applies to show $\{ac, ca\}$ is a λ-relation. In this application, the tableaux U, $V=\phi(U)$ with $\hat{p}(U)=ac$, $\hat{p}(V)=ca$ are unique, and the required dual equivalence $T\cup U \approx T\cup V$ is the one in entry (1) of Table 1. The λ-relation $\{ac, ca\}$ $(|c-a|>1)$ implies those in entries (1) and (2) of Table 1, so we arrive at a more concise full set by Lemma 5.7, see Table 2.

Of course, the Table 2 relations are nothing but the Coxeter relations connecting reduced expressions for elements of A_l. The tableau T_0 whose entries increase left to right one row at a time has promotion sequence

$$\hat{p}(T_0) = 12\cdots l12\cdots(l-1)\cdots 123121,$$

and this is a reduced expression for the longest element w_{A_l}. Immediately we obtain from Corollary 5.6 two theorems.

Theorem 5.10 (Greene–Edelman). *The operation $T \mapsto \hat{p}(T)$ is a bijection from standard tableaux of shape A_l to reduced expressions for the longest element w_{A_l} in the Coxeter group A_l.*

Theorem 5.11 (Greene–Edelman). *All reduced expressions for the longest element w_{A_l} of the Coxeter group A_l are connected by certain restricted Coxeter relations, which are the A_l-relations from Table* 1.

5.2. The case $\lambda = B_l$

In this subsection, we fix l and the shifted shape $\lambda = B_l$. As corner cell labels we use the integers $0, \ldots, l-1$. As before, we use these symbols also for the Coxeter group of type B_l and its simple reflections. Here 0 is to be the 'special' simple reflection at the end of the Dynkin diagram with the double link.

Representing B_l as the group of signed permutations of $\{1, \ldots, l\}$, the simple reflection 0 corresponds to the sign change $1 \mapsto -1$; the others i to adjacent transpositions $(i, i+1)$. The longest element w_{B_l} is the one that changes all signs, $w_{B_l}(j) = -j$.

The analogue of Table 1 is shown in Table 3. Note that $\hat{p}(T)$ for a permutation tableau is essentially T^{-1}, regarding T as the permutation that is its reading word. The corresponding λ-relations are thus the relations from Proposition 2.2, list B, reinterpreted for their effect on the inverse of a permutation as was done in Corollary 3.2.

From Proposition 5.4 we get $\{ca, ac\}$ ($|c-a|>1$) as a B_l-relation, using for instance entry (4) of Table 3 for the dual equivalence $T \cup U \approx T \cup \phi(U)$. This λ-relation implies those in entries (1) through (5). We can also apply Proposition 5.4 to get $\{bab, aba\}$ ($|b-a|=1$, $a \neq 0$, $b \neq 0$) from entry (6) of Table 3. This λ-relation implies those in entries (6) and (7). In both these applications of Proposition 5.4, the tableaux U and $\phi(U)$ are unique.

Using Lemma 5.7 we arrive at Table 4.

These are exactly the Coxeter relations connecting reduced expressions in B_l. In the proof of Corollary 5.6 we computed $\hat{p}(T_0)$ for the tableau whose entries increase left to right a row at a time; it is a reduced expression for w_{B_l}. Immediately we get two new theorems.

Theorem 5.12 (Proctor's conjecture). *The operation $T \mapsto \hat{p}(T)$ is a bijection from standard tableaux of shape B_l to reduced expressions for the longest element w_{B_l} in the Coxeter group B_l.*

Table 3
The full set of Proposition 5.5 B_l-relations

	Dual equivalence	λ-relation
(1)	Permutation tableaux	*acbd*, *cabd* *bacd*, *bcad* *dacb*, *dcab* *dbac*, *dbca* $(a < b < c,\ a \neq d < b)$
(2)	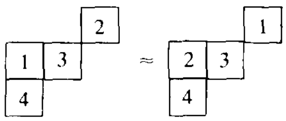	*acba*, *caba* $(a + 1 = b < c)$
(3)	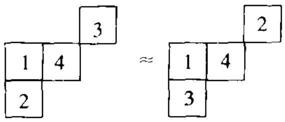	*bacb*, *bcab* $(a + 1 = b < c)$
(4)	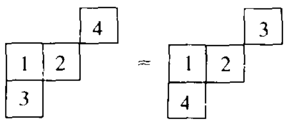	*abac*, *abca* $(a + 1 = b < c)$
(5)	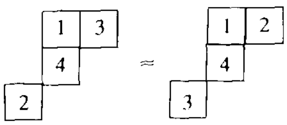	*bacb*, *bcab* $(a < b = c - 1)$
(6)	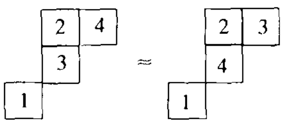	*dbab*, *daba* $(d < a,\ d < b,\ \lvert b - a \rvert = 1)$
(7)	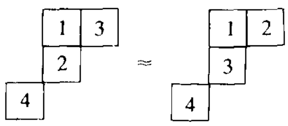	*babd*, *abad* $(d < a,\ d < b,\ \lvert b - a \rvert = 1)$
(8)	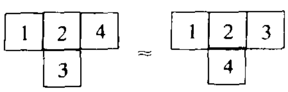	0101, 1010

Table 4
A full set of B_l-relations

ca, *ac*	$(\lvert c - a \rvert > 1)$
bab, *aba*	$(\lvert b - a \rvert = 1,\ a \neq 0,\ b \neq 0)$
0101, 1010	

Theorem 5.13. *All reduced expressions for the longest element w_{B_l} of the Coxeter group B_l are connected by certain restricted Coxeter relations, which are the B_l-relations from Table 3. These can be summarized as:*

$$\left.\begin{array}{l} acbd,\ cabd \\ bacd,\ bcad \\ dacb,\ dcab \\ dbac,\ dbac \end{array}\right\} \quad (a < b < c,\ d \leqslant b,\ \text{if } d = b \text{ then } b = a + 1 \text{ or } b = c - 1)$$

$$\left.\begin{array}{l} daba,\ dbab \\ abad,\ babd \end{array}\right\} \quad (d < a,\ d < b,\ |b - a| = 1)$$

$$0101,\ 1010$$

Kraskiewicz [5] has independently constructed a bijection from reduced expressions for w_{B_l} to tableaux of shape B_l. It can be shown that his bijection is $\hat{p}^{-1}$.

5.3. *The case $\lambda = C_l$*

In this subsection we fix l and the shifted shape $\lambda = C_l$, with corner cells labelled $1, \ldots, l$. We will be concerned here with the Coxeter group B_{l+1} and we use $0, 1, \ldots, l$ for its simple reflections as in Section 5.2.

For C_l, Table 5 is the analogue to Tables 1 and 3. Here we carry over all but entry (8) from Table 3, and add entries for final segments involving both cells in the last row of C_l.

Exactly as for B_l, we get from Proposition 5.4 the λ-relation $\{ca, ac\}$ ($|c - a| > 1$), as well as $\{bab, aba\}$ ($|b - a| = 1,\ b \neq 1,\ a \neq 1$). These λ-relations imply those in entries (1) through (5) of Table 5.

It might appear that Proposition 5.4 would also give $\{212, 121\}$ as a C_l-relation, using the last entry in Table 5. This is false, however, because there is not a unique final segment tableau V with $\hat{p}(V) = 121$. Using Lemma 5.7, Table 6 is the best we can get.

We now define the sequences which will turn out to be the C_l-words.

Definition. A *winnowed expression of order l* is a sequence that can be obtained from a reduced expression for the longest element $w_{B_{l+1}}$ in the Coxeter group B_{l+1} by deleting all instances of the simple reflection 0.

Observe that any reduced expression for $w_{B_{l+1}}$ contains 0 exactly $l + 1$ times, since in its representation as a signed permutation every element of $\{1, \ldots, l + 1\}$ must change sign once. Therefore every winnowed expression has length $l(l + 1) = |C_l|$.

In order to show that the winnowed expressions are the C_l-words, we first establish that the set of them is closed.

Table 5
The full set of Proposition 5.5. C_l-relations

	Dual equivalence	λ-relation
(1)	From Table 3	Entries (1) through (7)
(2)	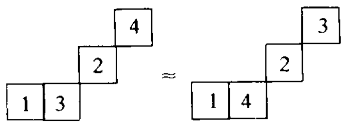	$1b1c$, $1bc1$ $(1<b<c)$
(3)	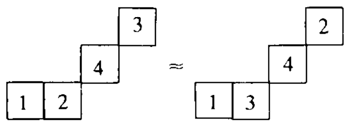	$11cb$ $1c1b$ $(1<b<c)$
(4)	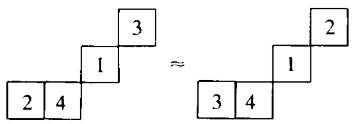	$b1c1$, $bc11$ $(1<b<c)$
(5)	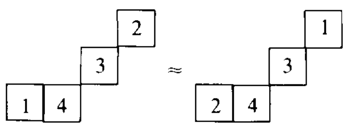	$1cb1$, $c1b1$ $(1<b<c)$
(6)	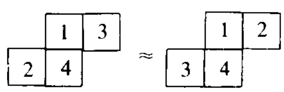	2121, 1211
(7)	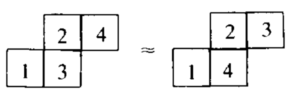	1212, 1121

Table 6
A full set of C_l-relations

ca, *ac*	$(\|c-a\|>1)$
bab, *aba*	$(\|b-a\|=1,\ a\neq 1,\ b\neq 1)$
1212, 1121	
2121, 1211	

Lemma 5.14. *The set of winnowed expressions of order l is closed under substitutions reflecting the C_l-relations of Table 6.*

Proof. Let W be a winnowed expression obtained from the B_{l+1}-word X.

If ca occurs in W, with $|c-a|>1$, it comes from either ca or $c0a$ in X. In B_{l+1}, c and a commute and at least one of them commutes with 0, so $ca=ac$ and $c0a=0ac$ or $ac0$. Hence we can substitute ac for ca and get another winnowed expression.

If bab occurs in W, with $|b-a|=1$ and $a, b\neq 1$, it comes from either bab, $b0ab$, or $ba0b$ in X ($b0a0b$ would not be reduced). Since they are not 1, a and b commute with 0, so $bab = aba$ and $b0ab = ba0b = 0aba$. Hence we can substitute aba for bab and get a winnowed expression.

If 1121 occurs in W it comes from 10121, 101201, or 101021 in X (adjacent 1's are not reduced, nor is 1010201). We have in B_{l+1} that $10121=10212$ and $101201=101021=010121=010212$. Winnowing, we get 1212 to substitute for 1121. Conversely, if 1212 occurs in W it comes from 12012, 10212, 120102, or 102102 in X (without one 0 between the 1's it is not reduced). In B_{l+1}, $12012=12012=10121$ and $120102=102102=102120=101210$. Winnowing, we get 1121 to substitute for 1212. The substitution $\{2121, 1211\}$ follows by symmetry. □

Proposition 5.15. *The set of winnowed expressions of order l is equal to the set of C_l-words.*

Proof. It is readily verified that for the tableau T_0 of shape C_l whose entries increase left to right a row at a time, $\hat{p}(T_0)$ is a winnowed expression. In fact it is the winnowing of the promotion sequence for the similar tableau of shape B_{l+1}. It follows from Lemma 5.14 that every C_l-word is a winnowed expression.

Now the set of all winnowed expressions is connected under substitutions of the forms

$$ca, ac \quad (|c-a|>1),$$
$$bab, aba \quad (|b-a|=1).$$

This is true because the B_{l+1}-words are connected by Coxeter relations from Table 4, which reduce to these substitutions, except for $\{0101, 1010\}$ which reduces to no change at all after winnowing.

If the set of C_l-words were closed under the above substitutions, it would show that every winnowed expression is a C_l-word, completing the proof. Unfortunately, the substitution $\{121, 212\}$ is not valid for C_l-words. However, it suffices to show something weaker, namely that if $A121B$ and $A212B$ are both winnowed expressions, and one is a C_l-word, then so is the other.

Both the set of C_l-words and the set of B_{l+1}-words, hence also the set of winnowed expressions, are closed under cyclic permutations (by Theorem 4.4) so we can assume B is empty and our winnowed expressions are $A212$ and $A121$. Using Lemma 5.14, we can modify A by any C_l-relation. Consequently, if T is a tableau such that $\hat{p}(T)$ is one of our winnowed expressions, we can modify its initial segment of size $|C_l| - 3$ by any elementary dual equivalence, and since this segment is normal we can replace it by any tableau of the same shape.

Now we examine three possible cases. In the diagrams below, $m_1 < m_2 < m_3 < m_4$ stand for the four greatest entries of T.

(I) $\hat{p}(T) = A212$, so T has a final segment

$$\begin{array}{cc} \boxed{m_2} & \boxed{m_4} \\ \boxed{m_3} & \end{array}.$$

Then we can assume this is part of a final segment

$$\begin{array}{ccc} & \boxed{m_2} & \boxed{m_4} \\ \boxed{m_1} & \boxed{m_3} & \end{array}$$

so that $A212 = C1212$ and $A121 = C1121$ is a C_l-word.

(II) $\hat{p}(T) = A121$ and T has a final segment

$$\begin{array}{cc} \boxed{m_2} & \boxed{m_3} \\ \boxed{m_4} & \end{array}.$$

Just as in case (I), we can assume $A121 = C1121$ so that $A212 = C1212$ is a C_l-word.

(III) $\hat{p}(T) = A121$ and T has the final segment

$$\begin{array}{ccc} & & \boxed{m_3} \\ \boxed{m_2} & \boxed{m_4} & \end{array}.$$

Then we can assume this is part of a final segment

$$\begin{array}{ccc} & \boxed{m_1} & \boxed{m_3} \\ \boxed{m_2} & \boxed{m_4} & \end{array}$$

so that $A121 = C2121$. But then $A212 = C2212$ and this contradicts the hypothesis that $A212$ was a winnowed expression, for 22 could only come from winnowing 22 or 202 and these are not reduced. □

As corollaries, we have the expected two theorems.

Theorem 5.16. *The operation $T \mapsto \hat{p}(T)$ is a bijection from standard tableaux of shape C_l to winnowed expressions of order l.*

Theorem 5.17. *All winnowed expressions of order l are connected by certain Coxeter-type relations, which are the C_l-relations from Table 6, or more restrictively, from Table 5.*

The last theorem is purely a statement about reduced expressions for $w_{B_{l+1}}$. It does not seem apparent whether this theorem can be proved easily without recourse to the tableau-promotion machinery used here.

Incidentally, winnowed expressions have a direct combinatorial interpretation: a winnowed expression of order l is a minimum-length sequence of adjacent transpositions of $l+1$ symbols realizing the identity permutation and such that each symbol is carried to the left-most position at some intermediate stage.

6. Reduced expressions of general Coxeter group elements

The A_l and B_l correspondences from Section 5 can be used to derive canonical non-negative integers m_λ^α for any element α of the relevant Coxeter group and each normal shape λ such that the number of reduced decompositions of α equals $\sum_\lambda m_\lambda^\alpha f_\lambda$, where f_λ denotes the number of standard tableaux of shape λ. For A_l, this was worked out by Greene and Edelman. In this section we work out the analogue for B_l.

For A_l, the numbers m_λ^α are the coefficients expanding Stanley's symmetric functions of [10] into Schur functions. This was noted by Greene and Edelman, and also by Lascoux and Schützenberger [6], whose 'nilplactic monoid' provides an alternative approach to the Greene–Edelman theory for A_l and an interpretation of the 'Stanley functions' in terms of Schubert polynomials. As Stanley observed, the analogous functions for B_l fail to be symmetric. Nevertheless he conjectured that the function associated to w_{B_l} is symmetric and equal to the Schur function $s_{R_{l,l}}$. Here we give a brief treatment of Stanley functions in order to prove this conjecture and at the same time explain their failure to be symmetric in general for B_l.

Throughout this section, w stands for w_{A_l} or w_{B_l}, depending upon the context.

Proposition 6.1. *The bijections $T \mapsto \hat{p}(T)$ of Theorems 5.10 and 5.12 have the property that the final segment of T containing $\{k, \ldots, n\}$ determines the corresponding segment $(r_k, \ldots, r_n)$ of $\hat{p}(T)$, while the initial segment $(r_1, \ldots, r_k)$ of $\hat{p}(T)$ determines the corresponding initial segment of T, for any $1 \leq k \leq n$.*

Proof. The relationship of the final segments is clear from the definition of $\hat{p}(T)$. As for the initial segments, suppose $\hat{p}(T) = AB$ and $\hat{p}(T') = AB'$. B and B' are reduced decompositions of the same element $A^{-1}w$, so we may as well assume they differ by a single Coxeter relation. If that relation is of the form $aba = bab$ in the case of A_l or $0101 = 1010$ in the case of B_l, then T and T' differ only by an elementary dual equivalence in the corresponding segment, so have the same initial segment. Otherwise the relation is either $ac = ca$, or, in the case of B_l, possibly $aba = bab$. But then according to Proposition 5.4 we may obtain T' from T by applying p^t for some $t \leq |B| - 2$ (or $|B| - 3$), exchanging n and $n-1$ in the corner cells labelled a and c, (or a and b), and applying p^{-t}. In this process, the initial segment of T corresponding to A is subjected to forward and reverse slides, and no other changes. Since this segment begins and ends with normal shape, it ends the same as it begins. □

This proof apparently cannot be adapted to Theorem 5.16. While we could replace the Coxeter relations with Table 6, in the absence of Coxeter group structure we cannot conclude that B and B' differ by these relations just because AB and AB' do. Nevertheless the following seems quite probable.

Conjecture 6.2. Proposition 6.1 applies to the bijection of Theorem 5.16 as well.

In view of Proposition 6.1, we can meaningfully extend the map $\hat{p}^{-1}$ carrying reduced expressions of w into tableaux to a map carrying arbitrary reduced expressions into tableaux.

Definition. Let E be a reduced expression for some element in A_l or B_l. Extend E to a reduced expression EX for w, and define $\theta(E)$ to be the initial segment of size $|E|$ in the tableau $\hat{p}^{-1}(EX)$. Thus $\theta(E)$ is a normal tableau whose shape is contained in A_l or B_l and $\theta(E)$ is well-defined, independently of the choice of X, by Proposition 6.1.

Theorem 6.3. *Let $\alpha \in A_l$ or B_l. Then as E varies over all reduced expressions for α, the multiset of tableaux $\theta(E)$ contains each tableau T of each shape λ with multiplicity m_λ^α depending on λ but not on T.*

Proof. Fix a reduced expression X of $\alpha^{-1}w$. Then E is a reduced expression for α if and only if EX is a reduced expression for w. Hence the tableaux $\theta(E)$ consist, as a multiset, of the initial segments of all tableaux S of shape A_l or B_l whose final segments have short promotion sequence X. In particular, the required m_λ^α is the number of such final segments with shape A_l/λ or B_l/λ, in the respective cases. □

Corollary 6.4. *The number of reduced expressions for α is equal to $\sum_\lambda m_\lambda^\alpha f_\lambda$.*

The Stanley symmetric functions are defined as follows. Let $\{x_1, x_2, \ldots\}$ be a countably infinite set of variables. For any subset $D \subseteq \{1, \ldots, n-1\}$ define a monomial $x_{i_1} \cdots x_{i_n}$ of degree n to be *D-admissible* if $i_k \leq i_{k+1}$ for all k, and $i_k < i_{k+1}$ when $k \in D$. The *Gessel quasi-symmetric function* $Q_D(x)$ is then the sum of all D-admissible monomials. $Q_D(x)$ is in general a nonsymmetric polynomial of degree n. If we define for a tableau T of size n the *descent set* $D(T)$ by $k \in D(T)$ when $k+1$ occurs in a strictly lower row of T than k, it is easy to see that the sum

$$\sum_{\text{sh } T = \lambda} Q_{D(T)}(x)$$

is the generating function for *column-strict* tableaux of shape λ, which is a symmetric function. In fact it is (by definition) the Schur function $s_\lambda(x)$.

Definition. Let $\alpha \in A_l$. The *Stanley symmetric function* $F_\alpha(x)$, homogeneous of degree $l(\alpha)$, is given by

$$\sum_E Q_{D(E)}(x),$$

where E ranges over reduced expressions for α, and $D(E)$ is the *descent set* of E, i.e., the set of k such that $E_k > E_{k+1}$.

Proposition 6.5. *For T of shape A_l, with $E = \hat{p}(T)$, $D(T) = D(E)$.*

Proof. The property $k \in D(T)$ is unaffected by the operation p^{n-k-1}, provided we do not renormalize the tableau entries in performing p^{n-k-1}, since this operation only subjects the entries $\{1, \ldots, k+1\}$ to slides, and slides preserve descents in the unshifted plane. After this operation, k and $k+1$ occupy corner cells labelled E_k and E_{k+1}, so k is a descent in T if and only if it is a descent in E. $\square$

Corollary 6.6. *$F_\alpha(x)$ is symmetric and its expansion into Schur functions is $F_\alpha(x) = \sum_\lambda m_\lambda^\alpha s_\lambda$.*

Slides in the shifted plane do not preserve descents, so Proposition 6.5 fails for B_l. However, there is a bijection between shifted tableaux of shape B_l and unshifted tableaux of square shape $R_{l,l}$, see [2, 12]. From $R_{l,l}$ to B_l, this bijection is given by regarding the unshifted T as shifted and computing $\nabla \leftarrow T$ by shifted *jeu-de-taquin*. From this description it follows that the bijection commutes with promotion (this is most easily seen by considering p^{-1}). It is also true that $n-1 \in D(T)$ if and only if $n-1 \in D(\nabla \leftarrow T)$. If $E = \hat{p}(\nabla \leftarrow T)$, then $k \in D(E)$ if and only if $n-1 \in D(E')$, where E' is E permuted cyclically $n-1-k$ places to the right. Hence $n-1 \in D(E')$ if and only if $n-1 \in D(p^{n-1-k}(\nabla \leftarrow T))$ if and only if $n-1 \in D(p^{n-1-k}(T))$ if and only if $k \in D(T)$. Thus $D(E) = D(T)$ and we have proved the following.

Corollary 6.7. *For $\beta = w_{B_l}$, $F_\beta(x)$ is equal to the Schur function $s_{R_{l,l}}$.*

References

[1] P. Edelman and C. Greene, Balanced tableaux, Adv. Math. 63 (1) (1987) 42–99.

[2] M. Haiman, On mixed insertion, symmetry, and shifted Young tableaux, J. Combin. Theory Ser. A 50 (2) (1989) 196–225.

[3] M. Haiman and D. Kim, A characterization of generalized staircases, Discrete Math. 99 (this Vol.) (1992) 115–122.

[4] D.E. Knuth, Permutations, matrices, and generalized Young tableaux, Pacific J. Math. 34 (1970) 709–727.

[5] W. Kraskiewicz, Reduced decompositions in hyperoctahedral groups, Manuscript, Mathematical Institute, Torun, Poland.

[6] A. Lascoux and M.-P. Schützenberger, Structure de Hopf de l'anneau de cohomologie et de l'anneau de Grothendieck d'une variete de drapeaux, C. R. Acad. Sci. Paris Sér. I Math. 295 (1982) 629–633.

[7] B.E. Sagan, Shifted tableaux, Schur Q-functions, and a conjecture of R. Stanley, J. Combin. Theory Ser. A 45 (1) (1987) 62–103.

[8] M.-P. Schützenberger, Promotion des morphisms d'ensembles ordonnes, Discrete Math. 2 (1972) 73–94.

[9] M.-P. Schützenberger, La correspondance de Robinson, in: D. Foata, ed., Combinatoire et Représentation du Groupe Symétrique—Strasbourg, 1976, Lecture Notes in Math. 579 (Springer, Berlin, 1977) 59–113.

[10] R.P. Stanley, On the number of reduced decompositions of elements of Coxeter groups, European J. Combin. 5 (1984) 359–372.

[11] R.P. Stanley, private communication.

[12] D.R. Worley, A theory of shifted Young tableaux, Ph.D. Thesis, M.I.T., 1984.