JOURNAL OF COMBINATORIAL THEORY, Series A 45, 62–103 (1987)

Shifted Tableaux, Schur Q-Functions, and a Conjecture of R. Stanley

BRUCE E. SAGAN*

Department of Mathematics E1, University of Pennsylvania, Philadelphia, Pennsylvania 19104-6395

Communicated by the Managing Editors

Received February 28, 1986

We present an analog of the Robinson–Schensted correspondence that applies to shifted Young tableaux and is considerably simpler than the one proposed in [B. E. Sagan, *J. Combin. Theory Ser. A* **27** (1979), 10–18]. In addition, this algorithm enjoys many of the important properties of the original Robinson–Schensted map including an interpretation of row lengths in terms of k-increasing sequences, a jeu de taquin, and a generalization to tableaux with repeated entries analogous to Knuth's construction (*Pacific J. Math.* **34** (1970), 709–727). The fact that the Knuth relations hold for our algorithm yields a simple proof of a conjecture of Stanley. © 1987 Academic Press, Inc.

1. INTRODUCTION

The Robinson–Schensted correspondence [R, Se] is at the heart of many combinatorial results concerning representations of the symmetric group $\mathcal{S}_n$. This map gives a simple bijective proof of the degree formula (cf. Corollary 3.2 below) and has been generalized by White [W2] to prove the orthogonality relations for the complete character table. Another generalization due to Knuth [K] provides a combinatorial proof of a fundamental identity involving Schur functions. White [W1] and Thomas [T2] have shown that there is a connection between the Robinson–Schensted algorithm and the Littlewood–Richardson rule [LR] which computes induced characters. In addition, the algorithm is equivalent to Schützenberger's jeu de taquin [Sü1] and can be used to compute the lengths of monotone subsequences of a permutation [Se, Gr].

In [S] we proposed an analog of the Robinson–Schensted map involving shifted Young tableaux. The main purpose of this construction was to give a bijective proof of the degree formula for projective representations of $\mathcal{S}_n$. Unfortunately the analog was fairly complicated and did not seem to

* Work supported in part by a NATO post-doctoral grant administered by the NSF.

0097-3165/87 $3.00

have any of the other nice properties of the original correspondence. In this paper we will present a new shifted Robinson–Schensted map which is not only much simpler but also reflects many of the important aspects of its parent algorithm. In particular we will use the Knuth relations for the shifted correspondence to give an almost trivial proof of a conjecture of Stanley that had withstood attack by other means.

It should be mentioned that many of these results have also been obtained independently by Worley and appear in his MIT thesis [Wo]. Worley's methods, however, differ substantially from our own and his work should be consulted as an alternative approach to this subject.

2. Definitions and Notation

Let $\mathbb{P}$ denote the set of positive integers. Recall that a *partition* of $n \in \mathbb{P}$ is a positive integral vector $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_l)$ such that $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_l$ and $\sum \lambda_i = n$. We write $\lambda \vdash n$ to show that λ is a partition of n. The *length of* λ, $l(\lambda)$, is the number of components l. A partition of n is *strict*, written $\lambda \models n$, if $\lambda_1 > \lambda_2 > \cdots > \lambda_l$.

Given a strict partition $\lambda \models n$ the corresponding *shifted shape* λ is an arrangement of n cells in $l(\lambda)$ rows with row i containing λ_i cells and indented $i-1$ spaces. For example, if $\lambda = (4, 2, 1)$ then the shifted shape of λ is shown in Fig. 1. We denote by (i, j) the cell in the ith row and jth column of the shifted shape so that

$$\lambda = \{(i, j) \mid 1 < i < l(\lambda) \text{ and } i < j < i + \lambda_i - 1\}.$$

A set of cells that will play a crucial role in the sequel will be the *diagonal* $\{(i, i) \in \lambda\}$. All other cells in λ are *off-diagonal*.

A (*partial*) *shifted Young tableaux* P of shape λ is an injection $P\colon \lambda \to \mathbb{P}$ which is increasing along rows and columns. Informally we think of P as a placement of distinct elements of $\mathbb{P}$ in the cells of λ and let p_{ij} denote the element of P in cell (i, j). If we wish to refer only to the shape of P we will write expressions like $(i, j) \in \text{shape } P$. Sometimes it will be convenient to consider P as merely the set of entries p_{ij} so that we can use notation such

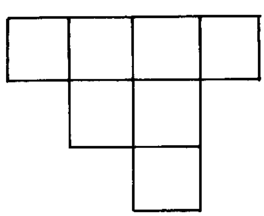

Fig. 1. The shifted shape (4, 2, 1).

as $x \in P$ to mean $x = p_{ij}$ for some (i, j). By way of illustration, one possible shifted tableau of shape $(4, 2, 1)$ is

$$P = \begin{matrix} 1 & 3 & 5 & 8 \\ & 4 & 7 & \\ & & 9 & \end{matrix} .$$

In this case $(2, 3) \in \text{shape } P$ but $(2, 4) \notin \text{shape } P$. Also $3 \in P$ because $p_{12} = 3$ but $2 \notin P$.

A partial shifted tableau is called *standard* if the numbers in the cells are precisely $1, 2, \ldots, n$, where $\lambda \models n$. The number of standard shifted tableaux of shape λ is denoted g_λ. As an example, $\begin{smallmatrix} 1 & 2 & 3 \\ & 4 & \end{smallmatrix}$ and $\begin{smallmatrix} 1 & 2 & 4 \\ & 3 & \end{smallmatrix}$ are the only standard shifted tableaux of shape $(3, 1)$ so $g_{(3,1)} = 2$.

In the following sections it will be useful to extend the positive integers to $\mathbb{P}_\infty = \mathbb{P} \cup \{\infty\}$ by adjoining a new symbol which is greater than every element of $\mathbb{P}$. It is now possible to embed a shifted tableau P in the full shifted plane $\{(i, j) \mid 1 \leqslant i \leqslant j\}$ by setting $p_{ij} = \infty$ for $(i, j) \notin \lambda$. We will assume from now on that all our tableaux have been so embedded without specific mention. However, we will often omit the extra ∞'s when working out examples.

All our algorithms will be written in a pidgin Pascal similar to that in Aho, Hopcroft, and Ullman [AHU]. There are two advantages to this approach. First of all, it will provide a rigorous description of the algorithms and leave us free to present a more informal synopsis in the text. Second, it will then be easier to code these procedures on an actual computer and investigate further properties of these algorithms.

3. The Shifted Robinson–Schensted Algorithm

Although the original Robinson–Schensted algorithm was first discovered by Robinson [R], our shifted analog will more closely follow the presentation of Schensted [Se] and we will henceforth use only the latter name. The function BUMP is fundamental to algorithms of this sort. It takes as input an integral vector $\mathbf{v}$ (padded with ∞'s) and an element $x \in \mathbb{P}$. The function scans the vector to find the first $\mathbf{v}[i]$ larger than x and replaces $\mathbf{v}[i]$ by x. Note that BUMP returns the value removed from $\mathbf{v}$ and also that if $\mathbf{v}$ is weakly increasing from left to right then after applying BUMP, $\mathbf{v}$ will still weakly increase.

```
function BUMP(x: ℙ; var v: array of ℙ): ℙ;
  var
    i: ℙ;
```

```
begin
    i := FIRST(v); {FIRST returns the index of the first component of v}
    while x ≥ v[i] do i := i + 1;
    BUMP := v[i];
    v[i] := x
end; {BUMP}
```

Given a shifted tableau P and a positive integer $x \notin P$ we can insert x into P to form a new shifted tableau as follows. Bump $y_1 = x$ into the first row of P which will remove an element y_2. Now we try to bump y_2 into the second row and continue until one of two things happens. Either some y_k will replace an ∞ at the end of a row and come to rest there, called a *Schensted move*, or y_k will replace some non-∞ diagonal element of P, a *non-Schensted move*. In the non-Schensted case we continue the bumping process along the columns: y_{k+1} is placed in column $k+1$ and so on until some y comes to rest at the end of a column. The function INSERT is coded below. It returns the new tableau and, as a side effect, sets a boolean variable *schen* indicating whether the insertion was a Schensted move or not. The reader can easily verify that the tableau property (increasing rows and columns) is preserved by INSERT.

```
function INSERT(x:ℙ; var P:SHIFT_TAB; var schen: boolean):
SHIFT_TAB;
    var
        k:ℙ;
        y_1, y_2,..., y_n:ℙ;
    begin
        schen := true;
        y_1 := x;
        k := 1;
        repeat begin
            y_{k+1} := BUMP(y_k, row k of P);
            k := k + 1
            end {repeat}
        until (y_k = ∞) or (y_{k-1} = P_{k-1,k-1});
        if y_k = ∞ then INSERT := P
        else begin
        repeat begin
            y_{k+1} := BUMP(y_k, column k of P);
            k := k + 1
            end {repeat}
        until y_k = ∞;
        schen := false;
```

(15) INSERT := P
end {**else**}
end; {INSERT}

If insertion of x into P yields P' we will write $I_x(P) = P'$ for short. As an example, the reader can verify that

$$I_4\begin{pmatrix} 1 & 2 & 3 & 6 \\ & 5 & 7 & \end{pmatrix} = \begin{array}{cccc} 1 & 2 & 3 & 4 \\ & 5 & 6 & \\ & & 7 & \end{array}$$

is a Schensted move and that

$$I_3\begin{pmatrix} 1 & 2 & 4 & 7 \\ & 5 & 6 & \end{pmatrix} = \begin{array}{ccccc} 1 & 2 & 3 & 6 & 7 \\ & 4 & 5 & & \end{array}$$

is non-Schensted.

We need one more concept before we can state the shifted Schensted correspondence. A shifted tableau is called *circled* if some of its elements have been distinguished by having rings placed around them, i.e., the entries of our tableau are taken from the set $\mathbb{P}^0 = \{①, 1, ②, 2, ③, 3, \ldots\}$ with the total order $① < 1 < ② < 2 < ③ < 3 < \cdots$ and the restriction that Ⓜ and m cannot both occur in the tableau. When we wish to refer to an integer without specifying whether it is circled or not we will write $\boxed{m}$. This convention will only apply to variables near the middle of the alphabet. With letters near the end, such as x and y, no assumption is to be made about whether the integers they represent are circled or not unless the contrary is specifically stated. Thus we could have $x =$ ⓝ or $x = n$ but n itself will always stand for an uncircled integer.

THEOREM 3.1 (shifted Schensted correspondence). *There is a bijection between permutations $\pi \in \mathcal{G}_n$ and pairs (P, Q) of standard shifted tableaux such that shape $P =$ shape $Q \models n$ and Q has a subset of its off-diagonal elements circled.*

Proof. If $\pi = x_1 x_2 \cdots x_n$ is one-line notation then we can find the corresponding pair (P, Q) using the procedure SHIFTED SCHENSTED. Starting with P and Q both empty, we insert the elements $x_1, x_2, \ldots, x_n$ into P. After inserting x_k, a new cell (i, j) will become filled at the end of a row and column of P. We then set

$$q_{ij} = \begin{cases} k & \text{if the insertion of } x_k \text{ was a Schensted move} \\ ⓚ & \text{if the insertion of } x_k \text{ was a non-Schensted move.} \end{cases}$$

Because $\pi \in \mathcal{G}_n$ we are assured that P and Q are standard with n cells. Also our choice of the cell (i, j) above ensures that at every stage shape $P =$ shape Q. Finally note that a non-Schensted move must involve the bumping of a diagonal element into the next column so q_{ii} will never be circled for any i.

```
procedure SHIFTED_SCHENSTED (x₁x₂···xₙ: PERMUTATION);
  var
    k: ℙ;
    schen: boolean;
    P: SHIFT_TAB;
    Q: CIRC_SHIFT_TAB;
  begin
    P := ∅;
    Q := ∅;
    for k := 1 to n do begin
      P := INSERT (x_k, P, schen);
      (i, j) := shape P − shape Q;
      {shape Q is the same as shape P before the insertion}
      if schen then q_ij := k else q_ij := ⓚ
  end {for}
end; {SHIFTED_SCHENSTED}
```

To show that this map is a bijection we construct its inverse step-by-step. The algorithms BUMPOUT, DELETE, and SHIFTED_SCHENSTED_INV accomplish this task.

```
function BUMPOUT (x: ℙ; var v: array of ℙ): ℙ;
  {BUMPOUT will only work properly if v[i] < x for some i}
  var
    i: ℙ;
  begin
    i := LAST (v);
    while x ⩽ v[i] do i := i − 1;
    BUMPOUT := v[i];
    v[i] := x
  end; {BUMPOUT}

function DELETE ((i, j): CELL; var P: SHIFT_TAB;
                                 var schen: boolean): ℙ;
  {DELETE returns the element removed from P}
  var
    k, l, y₁,..., yₙ: ℙ;
```

```
begin
  if schen then k := i else k := j;
  {k represents the number of BUMPOUTS needed}
  y_k := p_ij;
  if not (schen) then {in this case we must
                        start on the columns}
    repeat begin
      k := k - 1;
      y_k := BUMPOUT (y_{k+1}, column k of P)
      end {repeat}
    until y_{k+1} = p_kk
  for l := k - 1 downto 1 do {in both cases continue
                               along the rows}
    y_l := BUMPOUT (y_{l+1}, row l of P);
  DELETE := y_1
end; {DELETE}

procedure SHIFTED_SCHENSTED_INV (P: SHIFT_TAB;
                                  Q: CIRC_SHIFT_TAB);
var
  k: ℙ;
  x_1 x_2 ··· x_n: PERMUTATION;
begin
  for k := n downto 1 do begin
    (i, j) := cell containing [k] in Q;
    if q_ij = k then x_k := DELETE ((i, j), P, true)
             else x_k := DELETE ((i, j), P, false)
  end; {for}
end; {SHIFTED_SCHENSTED_INV}
```

If π is mapped to (P, Q) by the bijection above then we say that P is the *(shifted) P-tableaux corresponding to* π and write $P = P(\pi)$. A similar definition applies to Q. As an example we have listed the tableaux constructed when SHIFTED_SCHENSTED is applied to the permutation $\pi = 2\ 6\ 5\ 1\ 7\ 4\ 3$. The inverse construction is obtained by reading the list right to left

$$P\colon\ 2,\ \ 2\ 6,\ \ \begin{matrix}2 & 5\\ & 6\end{matrix},\ \ \begin{matrix}1 & 2 & 5\\ & 6 & \end{matrix},\ \ \begin{matrix}1 & 2 & 5 & 7\\ & 6 & & \end{matrix},\ \ \begin{matrix}1 & 2 & 4 & 7\\ & 5 & 6 & \end{matrix},\ \ \begin{matrix}1 & 2 & 3 & 6 & 7\\ & 4 & 5 & & \end{matrix} = P(\pi)$$

$$Q\colon\ 1,\ \ 1\ 2,\ \ \begin{matrix}1 & 2\\ & 3\end{matrix},\ \ \begin{matrix}1 & 2 & ④\\ & 3 & \end{matrix},\ \ \begin{matrix}1 & 2 & ④ & 5\\ & 3 & & \end{matrix},\ \ \begin{matrix}1 & 2 & ④ & 5\\ & 3 & ⑥ & \end{matrix},\ \ \begin{matrix}1 & 2 & ④ & 5 & ⑦\\ & 3 & ⑥ & & \end{matrix} = Q(\pi).$$

As an immediate corollary of this bijection we have a combinatorial proof of the degree formula for projective characters of $\mathscr{S}_n$.

COROLLARY 3.2 [Su]. $n! = \sum_{\lambda \models n} 2^{n-l(\lambda)} g_\lambda^2$.

Proof. $n!$ is the cardinality of $\mathscr{G}_n$ while g_λ^2 counts pairs of standard shifted tableaux of the same shape. The factor of $2^{n-l(\lambda)}$ accounts for the possible circlings of Q. ▮

4. THE LIFTING LEMMA

We will now relate the shifted Schensted correspondence to the ordinary (unshifted) one. In so doing we will prove a lemma that will permit us to lift many of the properties of the original algorithm to this new setting. First of all we recall some basic facts about left-justified tableaux.

Given a partition $\mu = (\mu_1, \mu_2, ..., \mu_l)$ of n, the *shape of* μ is an array of n cells in $l(\mu)$ left-justified rows with row i continuing μ_i cells. Such a shape will sometimes be called "left-justified," "ordinary," or "unshifted" to distinguish it from a shifted shape. The unshifted shape of $\mu = (4, 2, 1)$ is shown in Fig. 2. For reasons that will become clear shortly, the row indices start at 1 while the column indices start at 0 so that

$$\mu = \{(i, j) \mid 1 \leqslant i \leqslant l(\mu) \text{ and } 0 \leqslant j \leqslant \mu_i - 1\}.$$

A (*generalized*) *Young tableau* R of shape μ is a map $R: \mu \to \mathbb{P}$ which is weakly increasing along rows and strictly increasing along columns. A particular Young tableau of shape (4, 2, 1) would be

$$R = \begin{matrix} 1 & 1 & 2 & 2 \\ 2 & 3 & & \\ 4 & & & \end{matrix} .$$

To avoid confusion we will always use λ, P, and Q when referring to shifted shapes and tableaux, while usually reserving μ and R for their unshifted cousins.

Schensted's correspondence can be constructed using insertion by either rows or columns. However when it is applied to generalized tableaux, as

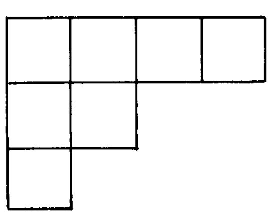

FIG. 2. The left-justified shape (4, 2, 1).

was done by Knuth [K], the column bumping must be modified to preserve column strictess, i.e., x must bump the first element larger than *or equal to* itself in the given column. This gives us a new **function** BUMPEQ which differs from BUMP only at line (2) must be changed to

```
(2')    while x > v[i] do i := i + 1;
```

Row insertion of x into a left-justified tableau R is like performing a Schensted move in the shifted case.

```
function ROW_INSERT (x: P; var R: GEN_TAB): GEN_TAB;
  var
    k: P;
    y_1, y_2,..., y_n: P
  begin
(1)   y_1 := x;
(2)   k := 1;
(3)   repeat begin
(4)     y_{k+1} := BUMP (y_k, row k of R);
(5)     k := k + 1;
        end {repeat}
(6)   until y_k = ∞;
(7)   ROW_INSERT := R
(8) end; {ROW_INSERT}
```

If applying ROW_INSERT to an element x and tableau R yields a tableau R' then we write $R_x(R) = R'$. Note that R's denoting row insertion will always bear subscripts while R's denoting tableaux will never have them.

Column insertion of x into R is done by the function COL_INSERT whose code is exactly like that of ROW_INSERT with line (4) replaced by

```
(4')    y_{k+1} := BUMPEQ (y_k, column k − 1 of R);
```

and with ROW changed to COL everywhere. Our shorthand for this operation will be $C_x(R) = R'$. The fundamental lemma relating these two operations is

LEMMA 4.1 [Se]. *Row and column insertion commute, i.e., for any tableau R and x, $y \in \mathbf{P}$:*

$$R_x \circ C_y(R) = C_y \circ R_x(R). \quad \blacksquare$$

If $\sigma = x_1 x_2 \cdots x_n$ is any sequence of positive integers then the ordinary

Knuth–Schensted correspondence associates with σ a pair of tableaux whose first member is

$$R = R(\sigma) = R_{x} \cdots R_{x_2} R_{x_1}(\varnothing).$$

This tableau is called the *R-tableau corresponding to* σ. R can also be built up by column insertions. In fact Lemma 4.1 yields

COROLLARY 4.2 [Se]. $R_{x_n} \cdots R_{x_2} R_{x_1}(\varnothing) = C_{x_1} \cdots C_{x_{n-1}} C_{x_n}(\varnothing)$.

The key concept connecting left-justified and shifted tableaux is that of a shift-symmetric tableau. Specifically, a *shift-symmetric tableau* is a generalized Young tableau R such that $R_{ij} = R_{ji-1}$ for all $i \geqslant j \geqslant 1$, e.g.,

$$R = \begin{matrix} 1 & 1 & 2 & 4 \\ 2 & 5 & 5 & \\ 4 & & & \end{matrix} .$$

The shift-symmetric tableaux are exactly those obtained by taking a partially shifted tableau P, transposing P, and then glueing together P and P^t along the diagonal. In the example above

$$P = \begin{matrix} 1 & 2 & 4 \\ & 5 & \end{matrix} .$$

Given tableaux P and R related in this way we write $P = \sqrt{R}$ (this is because the monomials corresponding to P and R satisfy this equation; see Sect. 8).

The next lemma is crucial to the rest of our discussion as it relates the shifted insertion operator to the two unshifted ones. Intuitively it says that the insertion of x into a shifted tableau is equivalent to both row and column inserting x into the corresponding shift symmetric tableau. The row bumping portion of I_x corresponds to the first half of R_x (we cannot see the second half because it is "below the diagonal") while the column bumping portion, if any, corresponds to the second half of C_x (the first half being obscured).

LEMMA 4.3 (Lifting Lemma). *Given a shifted tableau P and $x \notin P$ then $P = \sqrt{R}$ implies $I_x(P) = \sqrt{C_x R_x(R)}$. Equivalently, the following diagram commutes*

$$\begin{tikzcd} P \arrow[r, "I_x"] \arrow[d, "\text{square}"'] & P' \arrow[d, "\text{square}"] \\ R \arrow[r, "C_x R_x"'] & R' \end{tikzcd}$$

Proof. There are two cases.

Case 1. *I_x does not displace any diagonal element, finite or infinite.* If I_x applied to P causes elements to be bumped into cells $(1, j_1)$; $(2, j_2)$;...; (k, j_k) then $R_x(R)$ causes the same elements to be bumped into the corresponding cells. Finally, C_x causes these elements to be bumped into cells $(j_1, 0)$; $(j_2, 1)$;...; $(j_k, k-1)$ and the lemma follows.

Case 2. *I_x causes the displacement of a diagonal element.* (Note that this case includes both Schensted moves that end on the diagonal and all non-Schensted moves.) Let the cells disturbed by I_x be $(1, j_1)$; $(2, j_2)$;...; (d, d); $(i_{d+1}, d+1)$;...; (i_k, k). The insertion $R_x(R)$ starts with cells $(1, j_1)$;...; $(d-1, j_{d-1})$ but when y_d bumps into row d it will replace r_{dd-1}, since $y_d < r_{dd} = r_{dd-1}$ and $y_d > r_{d-1d} = r_{dd-2}$. Now y_{d+1} is the same as it was for I_x and is bumped into row $d+1$. Also the portion of row $d+1$ of R containing elements greater than y_{d+1} (the only ones y_{d+1} can bump) is the same as the corresponding portion of column $d+1$ of P. Hence R_x continues with the correct elements through the correct cells $(d+1, i_{d+1}-1)$;...; $(k, i_k - 1)$.

When performing C_x, note that although the zeroth column of $R' = R_x(R)$ many differ from the zeroth column of R (= first row of P), the only elements changed will be greater than the one which x must displace. Hence $y_1 = x$ enters cell $(j_1, 0)$ as desired. This continues to happen until y_d is to bump an element from column $d-1$ at which point $y_d = r'_{dd-1}$ so that y_d will replace r'_{dd-1} and we will have $y_{d+1} = r'_{dd-1}$. Hence $y_{d+1} = y_d$ will bump back onto the diagonal in the dth column of R' where it should be, to start the bumping which corresponds to the column insertion portion of I_x. ∎

Given $\pi = x_1 x_2 \cdots x_n$, let $\bar{\pi} = x_n \cdots x_2 x_1$ and consider the concatenation $\bar{\pi}\pi = x_n \cdots x_2 x_1 x_1 x_2 \cdots x_n$.

COROLLARY 4.4. *For $\pi \in \mathcal{G}_n$: $P(\pi) = \sqrt{R(\bar{\pi}\pi)}$.*

Proof. With π as above we have

$$\begin{aligned}
P(\pi) &= I_{x_n} \cdots I_{x_2} I_{x_1}(\varnothing) \\
&= \sqrt{C_{x_n} R_{x_n} \cdots C_{x_1} R_{x_1}(\varnothing)} && \text{by the lifting lemma and induction} \\
&= \sqrt{R_{x_n} \cdots R_{x_1} C_{x_n} \cdots C_{x_1}(\varnothing)} && \text{by Lemma 4.1} \\
&= \sqrt{R_{x_n} \cdots R_{x_1} R_{x_1} \cdots R_{x_n}(\varnothing)} && \text{by Corollary 4.2} \\
&= \sqrt{R(\bar{\pi}\pi)}. \quad ∎
\end{aligned}$$

Thus the behavior of the shifted P-tableau corresponding to π will mirror that of the unshifted R-tableau for the palindrome $\bar{\pi}\pi$.

5. Greene's Theorem

In his original paper [Se], Schensted gave an interpretation of the lengths of the first row and column of an R-tableau in terms of increasing and decreasing subsequences of the corresponding permutation. Later Greene [Gr] extended this interpretation to the whole shape as follows.

Given a sequence of positive integers $\sigma = x_1 x_2 \cdots x_n$ we say that σ is *strictly k-increasing* if σ is the union of k subsequences, each of which is strictly increasing (equivalently, contains no decreasing subsequence of length $k+1$). Note that "1-increasing" is the same as "increasing." Weakly k-increasing, strictly k-decreasing, and weakly k-decreasing sequences are defined analogously. Also let $a_k(\sigma)$ and $d_k(\sigma)$ denote the lengths of the longest weakly increasing and strictly decreasing subsequences of σ respectively. For example, $\sigma = 314159265$ has a 2-increasing subsequence $\sigma = 31415925 = 3459 \cup 1125$ which is of maximal length (since σ itself is not 2-increasing) so $a_2(\sigma) = 8$.

THEOREM 5.1 [Gr]. *Consider $\sigma \in \mathbb{P}^n$ and suppose the shape of $R(\sigma)$ is $\mu = (\mu_1, \mu_2, ..., \mu_l)$. Then for all $k \leqslant n$*

$$a_k(\sigma) = \mu_1 + \mu_2 + \cdots + \mu_k$$

$$d_k(\sigma) = \mu'_1 + \mu'_2 + \cdots + \mu'_k,$$

where $\mu' = (\mu'_1, \mu'_2, ..., \mu'_m)$ is the conjugate of μ, i.e., *μ'_j is the length of the* jth *column of μ.* ∎

The analog of this result for shifted tableaux is

COROLLARY 5.2. *Consider $\pi \in \mathcal{G}_n$ and suppose the shifted shape of $P(\pi)$ is $\lambda = (\lambda_1, \lambda_2, ..., \lambda_l)$. Then for all $k \leqslant l$*

$$a_k(\bar{\pi}\pi) = \lambda_1 + \lambda_2 + \cdots + \lambda_k + \binom{k+1}{2}$$

$$d_k(\bar{\pi}\pi) = \lambda_1 + \lambda_2 + \cdots + \lambda_k + \binom{k}{2}.$$

Furthermore for $k > l$ we have

$$a_k(\bar{\pi}\pi) = \lambda_1 + \cdots + \lambda_l + \lambda'_{l+1} + \cdots + \lambda'_k + \binom{l+1}{2}$$

$$d_k(\bar{\pi}\pi) = \lambda_1 + \cdots + \lambda_l + \lambda'_l + \cdots + \lambda'_{k-1} + \binom{l}{2},$$

where λ'_j is the length of the jth *column of the shifted shape.*

Proof. If $R = R(\bar{\pi}\pi)$ then by Corollary 4.4, R has shape $(\lambda_1 + 1, \lambda_2 + 2, ..., \lambda_l + l, \quad \lambda'_{l+1}, \quad \lambda'_{l+2}, ..., \lambda'_m)$ and conjugate shape $(\lambda_1, \lambda_2 + 1, ..., \lambda_l + l - 1, \lambda'_l, \lambda'_{l+1}, ..., \lambda'_m)$. Now apply Theorem 5.1. ▮

Note. (1) Here $\lambda' = (\lambda'_1, \lambda'_2, ..., \lambda'_m)$ is *not* the usual conjugate of λ.

(2) We have implicitly used the lifting lemma in the proof of Corollary 5.2 since Corollary 4.4 was proved using its aid.

We can use these ideas to explain all the entries of a tableau and not just its shape. Given any tableay R (shifted or not) with largest part m then R determines a sequence of tableaux $(R^{(1)}, R^{(2)}, ..., R^{(m)} = R)$, where $R^{(k)}$ is the tableau formed by all the elements up to size k in R, i.e.,

$$R^{(k)} = \{r_{ij} \in R \mid r_{ij} \leqslant k\}.$$

R also has an associated shape sequence $(\mu^{(1)}, \mu^{(2)}, ..., \mu^{(m)})$, where $\mu^{(k)} = \text{shape } R^{(k)}$. If, say,

$$R = \begin{array}{llll} 1 & 1 & 2 & 5 \\ 3 & 4 & 5 & 6 \\ 9 & & & \end{array}$$

then the associated tableau and shape sequences are

$$(R^{(1)}, R^{(2)}, R^{(3)}, ..., R^{(8)}, R^{(9)}) = \begin{array}{l}11\\ \\ \\ \end{array}, \begin{array}{l}112\\ \\ \\ \end{array}, \begin{array}{l}112\\3\\ \\ \end{array}, ..., \begin{array}{l}1125\\3456\\ \\ \end{array}, \begin{array}{l}1125\\3456\\9\end{array}$$

$$(\mu^{(1)}, \mu^{(2)}, \mu^{(3)}, ..., \mu^{(8)}, \mu^{(9)}) = ((2), (3), (3,1), ..., (4,4), (4,4,1)).$$

Clearly the shape sequence uniquely determines R and vice versa (in fact, some authors prefer to define tableaux in terms of their shape sequences [M]).

Similarly, given $\sigma = x_1 x_2 \cdots x_n \in \mathbb{P}^n$ with largest element m we have an associated sequence of subsequences $(\sigma^{(1)}, \sigma^{(2)}, ..., \sigma^{(m)})$, where $\sigma^{(k)} = x_{i_1} x_{i_2} \cdots x_{i_j}$ consists of all $x_i \leqslant k$. For $\sigma = 314159265$ we have $(\sigma^{(1)}, \sigma^{(2)}, \sigma^{(3)}, ..., \sigma^{(8)}, \sigma^{(9)}) = (11, 112, 3112, ..., 31415265, 314159265)$. The relationship between these concepts is given by

PROPOSITION 5.3. *The R-tableau for a sequence $\sigma \in \mathbb{P}^n$ with greatest element m satisfies $R(\sigma)^{(k)} = R(\sigma^{(k)})$ for all $k \leqslant m$.*

Proof. This is any easy reverse induction using the fact that if a largest element moves during an insertion it is always the last to be bumped. ▮

COROLLARY 5.4. *The P-tableau for $\pi \in \mathcal{G}_n$ satisfies*

$$P(\pi)^{(k)} = P(\pi^{(k)}) \qquad \text{for all } k \leqslant n.$$

Proof.

$$\begin{aligned} P(\pi)^{(k)} &= \sqrt{R(\bar{\pi}\pi)^{(k)}} && \text{by Corollary 4.4} \\ &= \sqrt{R(\bar{\pi}^{(k)}\pi^{(k)})} && \text{by Proposition 5.3} \\ &= P(\pi^{(k)}) && \text{by Corollary 4.4 again.} \quad \blacksquare \end{aligned}$$

Now for $\pi \in \mathcal{G}_n$ we let $\mathbf{a}(\pi) = (a_1(\bar{\pi}\pi), a_2(\bar{\pi}\pi), ..., a_n(\bar{\pi}\pi))$ and define the *Greene invariant of* π to be the vector $G(\pi) = (\mathbf{a}(\pi^{(1)}), \mathbf{a}(\pi^{(2)}), ..., \mathbf{a}(\pi^{(n)}))$. To illustrate, if $\pi = 2143$ then $G(\pi) = ((2), (3, 4), (4, 5, 6), (4, 7, 8))$. This is similar to, but not exactly the same as, the Greene invariant as defined by Worley [Wo]. We can now characterize those permutations having the same P-tableau in terms of their Greene invariants.

THEOREM 5.5. *Given π, $\sigma \in \mathcal{G}_n$ then $P(\pi) = P(\sigma)$ if and only if $G(\pi) = G(\sigma)$.*

Proof. $P(\pi) = P(\sigma)$ if and only if $P(\pi)$ and $P(\sigma)$ have the same shape vector $(\lambda^{(1)}, \lambda^{(2)}, ..., \lambda^{(n)})$. But $\lambda^{(k)} = \text{shape } P(\pi)^{(k)} = \text{shape } P(\pi^{(k)})$ from Corollary 5.4 and shape $P(\pi^{(k)}) = \text{shape } P(\sigma^{(k)})$ is equivalent to $\mathbf{a}(\pi^{(k)}) = \mathbf{a}(\sigma^{(k)})$ by Corollary 5.2. Since this holds for all k, the theorem follows from the definition of $G(\pi)$. $\blacksquare$

In the next two sections we will give two more characterizations of those permutations having the same P-tableau.

6. THE JEU DE TAQUIN

Schützenberger's "jeu de taquin" or "teasing game" [Sü1, Sü2] can be described as follows. Given left-justified shapes μ_1 and μ_2 with $\mu_2 \subseteq \mu_1$, the *skew shape* $\mu_1 - \mu_2$ is just the set-theoretic difference. If $\mu_1 = (4, 2, 1)$ and $\mu_2 = (1, 1)$ then the unshaded cells in Fig. 3 are the skew shape $\mu_1 - \mu_2$.

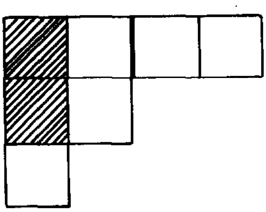

FIG. 3. The skew shape $(4, 2, 1) - (1, 1)$.

A *corner cell* of a shape μ in any (i, j) such that $(i+1, j)$, $(i, j+1) \notin \mu$, i.e., (i, j) is at the end of a row and column of μ. In the example above the corner cells of μ_1 are $(1, 4)$, $(2, 2)$, and $(3, 1)$, while μ_2 has corner cell $(2, 1)$.

As is to be expected, a *skew Young tableau* R of shape $\mu_2 - \mu_1$ is a function $R: \mu_1 - \mu_2 \to \mathbb{P}$ with the same row and column restrictions as for non-skew tableaux. Given R and any corner cell (i, j) of μ_2 the function MOVE creates a new skew tableau by filling (i, j) with either the element directly below or to the right, depending upon which is smaller. This action is called a *slide*. MOVE then proceeds to slide another element into the resulting hole until the empty cell propagates to a corner of μ_1 which is removed from R. It is easy to verify that MOVE preserves the (skew) tableau property

function MOVE (**var**(i, j): CELL; **var**R: SKEW_TAB): SKEW_TAB;
  var
    $r: \mathbb{P}$;
    temp: CELL;
    μ_1, μ_2: PARTITION;
  begin
    $\mu_1 - \mu_2 :=$ shape (R);
    $\mu_2 := \mu_2 - (i, j)$;
    if $r_{ij+1} < r_{i+1j}$ **then** $temp := (i, j+1)$ **else** $temp := (i+1, j)$;
    $r := r_{temp}$;
    while $r \langle\ \rangle \infty$ **do begin**
      $r_{ij} := r$;
      $(i, j) := temp$;
      if $r_{ij+1} < r_{i+1j}$ **then** $temp := (i, j+1)$ **else** $temp := (i+1, j)$
      $r := r_{temp}$
    end; {**while**}
    $\mu_1 := \mu_1 - (i, j)$;
  end; {MOVE}

A *diagonal strip* in the skew shape with cells $(n, 0)$, $(n-1, 1), \ldots,$ $(1, n-1)$. Given a sequence $\sigma = x_1 x_2 \cdots x_n \in \mathbb{P}^n$, place the elements of σ in a diagonal strip tableau R by setting $r_{n-i+1, i-1} = x_i$ for $i = 1, 2, \ldots, n$. We then play JEU_DE_TAQUIN by choosing corner squares at random and moving into them until R becomes a non-skew tableau. For example, if $\sigma = 3141$ then a typical game is shown in Fig. 4.

procedure JEU_DE_TAQUIN (σ: SEQUENCE);
  var
    R: SKEW_TAB;
    μ_1, μ_2: PARTITION;
    (i, j): CELL;

begin
R: diagonal strip corresponding to σ;
$\mu_1 - \mu_2 :=$ shape(R);
while $\mu_2 \langle\ \rangle \phi$ **do begin**
$(i, j) :=$ random corner cell of μ_2;
MOVE $((i, j); R)$
end{**while**}
end; {JEU_DE_TAQUIN}

If applying JEU_DE_TAQUIN to σ yields a tableau R then we write $\mathrm{J}(\sigma) = R$. It is not clear that J is well defined, however.

THEOREM 6.1 [Sü2, T1]. *For any $\sigma \in \mathbb{P}^n$ the tableau $J(\sigma)$ is independent of the order in which the corner cells are filled and in fact $J(\sigma) = R(\sigma)$.* ▮

The shifted analog of a diagonal strip consisting of cells (n, n), $(n-1, n+1), \ldots, (1, 2n-1)$. Starting with $\pi \in \mathscr{G}_n$ in a shifted diagonal strip, the shifted jeu de taquin is virtually identical to the unshifted one for sequences. The only difference is that, when sliding into a diagonal cell (i, i), there is no element in cell $(i+1, i)$ so we automatically use the $(i, i+1)$ entry. Thus we need only change the code of MOVE at lines (3) and (8) to read

if $(i = j)$ **or** $(r_{ij+1} < r_{i+1j})$ **then** $temp := (i, j+1)$ **else** $temp := (i+1, j)$

If π reduces to P under shifted jeu de taquin we write $\mathrm{SJ}(\pi) = P$.

THEOREM 6.2. *For any $\pi \in \mathscr{G}_n$ the shifted tableau* $\mathrm{SJ}(\pi)$ *is independent of the order in which the corner cells are filled and in fact* $\mathrm{SJ}(\pi) = P(\pi)$.

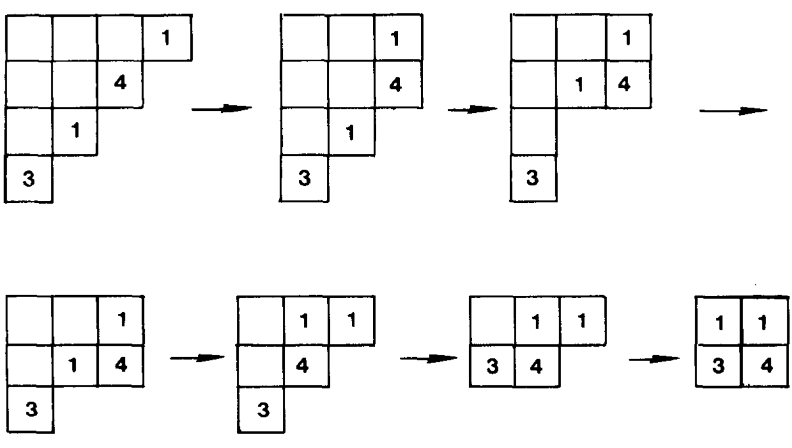

FIG. 4. The jeu de taquin.

Proof. Let P be the shifted diagonal strip for π and construct R by taking the diagonal strip for $\bar{\pi}\pi$ and applying the MOVE operation to cells $(n, n-1)$, $(n+1, n+2)$,..., $(2n-1, 0)$. Thus $P=\sqrt{R}$ with P and R reducing to $\mathrm{SJ}(\pi)$ and $\mathrm{J}(\bar{\pi}\pi)$ under their respective jeux. If we can show that $\mathrm{SJ}(\pi)=\sqrt{\mathrm{J}(\bar{\pi}\pi)}$ we will be done since $\sqrt{\mathrm{J}(\bar{\pi}\pi)}=\sqrt{R(\bar{\pi}\pi)}=P(\pi)$ by Theorem 6.1 and Corollary 4.4

By induction, it suffices to show that if $P=\sqrt{R}$ at some stage of the game and we move from P to P_1 then we can make moves in R to form R_1 so that $P_1=\sqrt{R_1}$. Let the move in P fill cell (i, j) and make the moves in R that fill (i, j) and $(j, i-1)$. If the move in P never passes through a diagonal cell then clearly $P_1=\sqrt{R_1}$ holds.

If diagonal cell (k, k) becomes empty in P then p_{kk+1} moves into it. Now in R, the move to fill (i, j) will cause (k, k) to become empty. But by shift-symmetry $r_{k+1k}=r_{k+1k+1}>r_{kk+1}$ so r_{kk+1} will slide into cell (k, k) as it should. This will continue to happen at all further diagonal cells on the path of this move. Hence in $R'=\mathrm{MOVE}(R, (i, j))$ we see that R' restricted to $\{(i, j) \mid j \geqslant i \geqslant 1\}$ is just P_1.

Now consider filling $(j, i-1)$ in R'. When the corresponding subdiagonal cell $(k, k-1)$ becomes empty then by shift-symmetry and the effect of the first move $r'_{k+1k-1}=r_{k+1k-1}=r_{kk+1}=r'_{kk}$. Thus r'_{k+1k-1} will slide into $(k, k-1)$ to preserve column-strictness and this move will stay below the diagonal sliding along the mirror image of the (i, j) slide. As before, $P_1=\sqrt{R_1}$. ∎

As an immediate corollary we have our second characterization of permutations having the same P-tableau.

COROLLARY 6.3. *Given $\pi, \sigma \in \mathscr{S}_n$ then $P(\pi)=P(\sigma)$ if and only if $\mathrm{SJ}(\pi)=\mathrm{SJ}(\sigma)$.*

7. KNUTH EQUIVALENCE

When Knuth generalized Schenstead's algorithm to tableaux with repeated entries [K] he also noted that there are certain adjacent transpositions in σ that leave the tableau $R(\sigma)$ invariant. Specifically, given a sequence $\sigma \in \mathbb{P}^n$ and three consecutive elements x, y, and z in σ (not necessarily in that order) then a *Knuth transposition* applied to σ can

(K1) replace xzy by zxy or vice versa if $x \leqslant y < z$,

(K2) replace yxz by yzx or vice versa if $x < y \leqslant z$.

We say that $\pi, \sigma \in \mathbb{P}^n$ are *Knuth equivalent*, written $\pi \equiv_{\mathrm{K}} \sigma$, if σ can be obtained from π by a sequence of Knuth transpositions. For example $3141592 \equiv_{\mathrm{K}} 3451129$ because

$$\begin{aligned} 3141592 &\equiv_K 3411592 && \text{by (K2)} \\ &\equiv_K 3411529 && \text{by (K2)} \\ &\equiv_K 3415129 && \text{by (K1)} \\ &\equiv_K 3451129 && \text{by (K1).} \end{aligned}$$

THEOREM 7.1 [K]. *Given $\pi, \sigma \in \mathbb{P}^n$ then $R(\pi) = R(\sigma)$ if and only if $\pi \equiv_K \sigma$.* ∎

For permutations $\pi = x_1 x_2 \cdots x_n \in \mathcal{G}_n$ the *shifted Knuth transpositions* are basically the same as the usual ones with the addition of a third option:

(SK1) replace xzy by zxy or vice versa if $x < y < z$,

(SK2) replace yxz by yzx or vice versa if $x < y < z$,

(SK3) replace $x_1 x_2$ by $x_2 x_1$,

i.e., the first two elements of a permutation can be switched regardless of size. Also, $\pi, \sigma \in \mathcal{G}_n$ are *shifted Knuth equivalent*, $\pi \equiv_{SK} \sigma$, if σ can be obtained from π by a sequence of transpositions of type (SK1), (SK2), or (SK3).

THEOREM 7.2. *Given $\pi, \sigma \in \mathcal{G}_n$ then $P(\pi) = P(\sigma)$ if and only if $\pi \equiv_{SK} \sigma$.*

Proof. We first prove sufficiency. If $\pi \equiv_{SK} \sigma$ implies $\bar{\pi}\pi \equiv_K \bar{\sigma}\sigma$ then we can use Theorem 7.1 and Corollary 4.4 to conclude that $P(\pi) = \sqrt{R(\bar{\pi}\pi)} = \sqrt{R(\bar{\sigma}\sigma)} = P(\sigma)$. By induction we may assume that σ is obtained from π by applying a single shifted Knuth transposition and need to show that it can be mimicked using ordinary Knuth transpositions.

Suppose the transposition is of type (SK1) with $\pi = x_1 \cdots xzy \cdots x_n$ and $\sigma = x_1 \cdots zxy \cdots x_n$. Thus

$$\begin{aligned} \bar{\pi}\pi &= x_n \cdots yzx \cdots x_1 x_1 \cdots xzy \cdots x_n \\ &\equiv_K x_n \cdots yxz \cdots x_1 x_1 \cdots zxy \cdots x_n && \text{by (K1)} \\ &\equiv_K x_n \cdots yzx \cdots x_1 x_1 \cdots zxy \cdots x_n && \text{by (K2) and the fact that } x \text{ and } y \text{ are distinct} \\ &= \bar{\sigma}\sigma. \end{aligned}$$

The second half of (SK1), replacing zxy by xzy, follows from symmetry and the argument for (SK2) is similar.

Now assume that the transposition is of type (SK3). If $x_1 < x_2$ then

$$\begin{aligned} \bar{\pi}\pi &= x_n \cdots x_2 x_1 x_1 x_2 \cdots x_n \\ &\equiv_K x_n \cdots x_1 x_2 x_1 x_2 \cdots x_n && \text{by (K1)} \\ &\equiv_K x_n \cdots x_1 x_2 x_2 x_1 \cdots x_n && \text{by (K2)} \\ &= \bar{\sigma}\sigma. \end{aligned}$$

If $x_1 > x_2$ then merely read the preceding four lines backwards. Hence the new shifted Knuth transposition for π is really a combination of two old ones using elements which are "invisible" since they live in $\bar{\pi}$.

It can also be proved directly that an (SK3) transposition leaves the P-tableau invariant. Merely note that $I_{x_2}I_{x_1}(\phi) = I_{x_1}I_{x_2}(\phi)$ since there is only one shifted tableau with two given elements. Hence

$$P(\pi) = I_{x_n} \cdots I_{x_3} I_{x_2} I_{x_1}(\phi) = I_{x_n} \cdots I_{x_3} I_{x_1} I_{x_2}(\phi) = P(\sigma).$$

Unfortunately we cannot lift the "only if" direction from Theorem 7.1. This is because it is not obvious that if $\bar{\pi}\pi \equiv_{\mathrm{K}} \bar{\sigma}\sigma$ then $\pi \equiv_{\mathrm{SK}} \sigma$ (the transpositions done in transforming $\bar{\pi}\pi$ to $\bar{\sigma}\sigma$ need not be done in symmetric pairs). However, a direct proof is available.

Given a shifted tableau P, the corresponding *word of* P, $\mathrm{w}(P)$, is the sequence obtained by reading the rows of P from left to right starting with the last row and working up. If

$$P = \begin{matrix} 1 & 3 & 5 & 8 \\ & 4 & 7 & \\ & & 9 & \end{matrix}\ ,$$

for example, then $\mathrm{w}(P) = 9471358$. Clearly the P-tableau of $\mathrm{w}(P)$ is P itself. Hence we need only show that if π is any other sequence with $P(\pi) = P$ then $\pi \equiv_{\mathrm{SK}} \mathrm{w}(P)$. By induction on the number of elements of P, the problem reduces to showing the following: if P' is a tableau such that $P' = I_x(P)$ for some x then $\mathrm{w}(P)\,x \equiv_{\mathrm{SK}} \mathrm{w}(P')$ where, as usual, juxtaposition denotes concatenation of sequences (or sequences and elements).

If $I_x(P)$ is a Schensted move then the insertion of x into P can be simulated using transposition of types (SK1) and (SK2) as in the unshifted case, proving the assertion. Since this simulation will be needed again later in the proof we will review the details. Suppose I_x displaces p_{1j} from the first row of P. Thus $p_{1j-1} < x < p_{1j} < p_{1j+1} < \cdots < p_{1\lambda_1}$ so that

$$\begin{aligned} \mathrm{w}(P)\,x = {} & p_{l1} \cdots p_{2\lambda_2} p_{11} \cdots p_{1j} \cdots p_{1\lambda_1} x \\ \equiv_{\mathrm{SK}} {} & p_{l1} \cdots p_{2\lambda_2} p_{11} \cdots p_{1j} x p_{1j+1} \cdots p_{1\lambda_1} \\ & \text{by repeated applications of (SK2)} \\ \equiv_{\mathrm{SK}} {} & p_{l1} \cdots p_{2\lambda_2} p_{1j} p_{11} \cdots p_{1j-1} x p_{1j+1} \cdots p_{1\lambda_1} \\ & \text{by repeated applications of (SK1).} \end{aligned}$$

But now the suffix (from p_{11} on) of this last word is the first row of $I_x(P)$ while p_{1j} is in the correct position to begin bumping into row two of P.

Hence this process can continue as long as only row insertions are required.

Now consider a non-Schensted insertion $I_x(P)$. If P has only one element y then $w\langle P\rangle x = yx \equiv_{SK} xy = \mathrm{w}(P')$ by (SK3). If P has more than one element then let $\mathrm{w}(P) = p_{l1} \cdots yz$ so that y and z are the last two elements in the first row of P. Also let $\lambda = (\lambda_1, \lambda_2, ..., \lambda_l)$ be the shape of P. We need to consider a couple of cases.

Case 1. *x bumps z from row* 1 *of P*. For any tableau P let $\hat{P}$ be the tableau obtained by removing the first row of P,

$$\begin{aligned} \mathrm{w}(P)\,x &\equiv_{SK} \mathrm{w}(\hat{P})\,zp_{11}\cdots yx && \text{by simulation of row insertion} \\ &\equiv_{SK} \mathrm{w}(I_z(\hat{P}))\,p_{11}\cdots yx && \text{by induction.} \end{aligned}$$

But $I_z(\hat{P})$ will either be $\hat{P}'$ or $\hat{P}'$ with a t adjoined in cell $(1, \lambda_1 + 1)$ in the case where $I_x(P)$ causes a t to be bumped into column $\lambda_1 + 1$. Next the insertions $I_{p_{11}}, ..., I_y$ will push all the elements in columns $1, ..., \lambda_1 - 1$ of $I_z(\hat{P})$ down one row. This yields a tableau P_1' which agrees with P' except along the λ_1th column which is still one row too high

$$\begin{aligned} \mathrm{w}(I_z(\hat{P}))\,p_{11}\cdots yz &\equiv_{SK} \mathrm{w}(P_1')\,x && \text{by induction} \\ &\equiv_{SK} \mathrm{w}(P') && \text{since } I_x(P_1') \text{ is Schensted} \end{aligned}$$

(it merely moves column λ_1 down one cell).

Case 2. *x bumps an element $<z$ from row* 1 *of P*. Thus $x < y < z$ and so using (SK1) and induction we have

$$\mathrm{w}(P)\,x = p_{l1}\cdots yzx \equiv_{SK} p_{l1}\cdots yxz \equiv_{SK} \mathrm{w}(P_1)\,xz \equiv_{SK} \mathrm{w}(I_x P_1)\,z,$$

where P_1 is P with z removed and the last column moved up a row. Now $I_z(P_1)$ is Schensted by definition of P_1. If we can show that I_z applied to $P_1' = I_x(P_1)$ is also Schensted then we will be done by the same reasoning as at the end of Case 1 (here, however, P_1' need not equal P_1 with its λ_1th column raised a cell).

To prove that $I_z(P_1)$ Schensted implies $I_z(P_1')$ Schensted, it suffices to show that the path of insertion of z into P_1' is the same as the path of insertion of z into P_1 except possibly for the first displacement which could be one cell to the right. The path of insertion, of course, consists of all cells where a bump takes place. In passing from P_1 to $P' = I_x(P_1)$ the contents of every cell decrease or stay the same. Thus if I_x causes no elements of the last column to be displaced then the two paths are identical. Suppose now that the path of I_x crosses column λ_1. It must do so on column insertion since $x < y$ and hence exactly one cell will be modified.

If that cell is $(1, \lambda_1)$ then the element from cell $(2, \lambda_1)$ of P ends up as entry $p'_{1\lambda_1+1}$ in P'_1. If I_z bumps $p'_{1\lambda_1+1}$ back to the second row, the path will continue down column λ_1 as before. Otherwise I_z will bump $p'_{1\lambda_1}$ which implies that $p'_{1\lambda_1} > z$ and so $p'_{1\lambda_1}$ must have come from cell $(i, \lambda_1 - 1)$ in P_1 where $i \geqslant 2$ ($i = 1$ would imply $p'_{1\lambda_1} = y < z$). Hence $p'_{2\lambda_1-1} < p'_{1\lambda_1}$ so that $p'_{1\lambda_1}$ must bump $p'_{2\lambda_1}$. But under these conditions the $(2, \lambda'_1)$ entries in P_1 and P' are identical so the paths are again the same from the second row on.

If the modified cell in column λ_1 is (i, λ_1) for $i \geqslant 2$ then the bumping path and elements bumped will agree through row $i-1$. Since P_1 is obtained by pushing up column λ_1 of P, we must have $p_{k\lambda_1-1} < p_{k-1\lambda_1}$ for all $k \geqslant 2$ and this property is preserved in P'_1. Thus $p'_{i\lambda_1-1} < p'_{i-1\lambda_1} < p'_{i\lambda_1}$ so that the path will continue through cell (i, λ_1) although it will displace a smaller element. But $p'_{i+1\lambda_1-1} < p'_{i\lambda_1} < p'_{i+1\lambda_1}$ so that $p'_{i+1\lambda_1}$ is bumped which is the same number from the same cell as in $I_z(P_1)$. Hence the insertion will push down the rest of the λ_1th column as desired. ∎

8. The Shifted Knuth Correspondence

We now generalize the construction of Section 3 to tableaux with repeated entries, thus obtaining a shifted version of Knuth's algorithm [K]. A *generalized shifted Young tableau*, T, is an assignment of elements of $\mathbb{P}^0$ to a shifted shape λ having the properties

(T1) T is weakly increasing along rows and columns and

(T2) For each integer $\boxed{m}$, there is at most one ⓜ in each row and at most one m in each column of T.

An example of a generalized shifted tableau is

$$T = \begin{array}{cccccccc} \textcircled{1} & 1 & 1 & \textcircled{2} & \textcircled{3} & 3 & 3 & \textcircled{4} \\ & 2 & 2 & 2 & \textcircled{3} & \textcircled{4} & 4 & 4 \\ & & \textcircled{3} & 3 & 3 & & & \\ & & & 4 & 4 & & & \end{array}$$

Note that conditions (T1) and (T2) imply that the cells of T containing $\boxed{m}$ form a union of *rim hooks* (connected skew shapes containing no 2×2 block of squares). Also from (T2) we see that the labeling of every $\boxed{m}$ in a rim hook is completely determined except for the one at the bottom left which can be circled or not.

As noted in Section 4, Knuth's original algorithm has as output a pair of *column strict* tableaux while shifted tableaux can have many copies of ⓜ in a column. Thus in describing the algorithms which follow it will be convenient to form *modified shifted tableaux* where the circled elements have been made distinct by attaching subscripts to them and insisting that

$$\textcircled{1}_1 < \textcircled{1}_1 < \textcircled{1}_3 < \cdots < 1 < \textcircled{2}_1 < \textcircled{2}_2 < \cdots < 2 < \cdots.$$

We also require that for each set of $\textcircled{m}$'s the subscripts increase from upper right to lower left. One modified tableau for the T above is

$$T = \begin{array}{cccccccc} \textcircled{1}_1 & 1 & 1 & \textcircled{2}_2 & \textcircled{3}_2 & 3 & 3 & \textcircled{4}_1 \\ & 2 & 2 & 2 & \textcircled{3}_3 & \textcircled{4}_2 & 4 & 4 \\ & & \textcircled{3}_4 & 3 & 3 & & & \\ & & & 4 & 4 & & & \end{array}$$

If we are to insert an $\textcircled{m}$ into T we must also modify it by giving it a subscript smaller than that of any $\textcircled{m}$ in T. Thus if we were to insert $\textcircled{2}$ into our running example then we would change it to $\textcircled{2}_1$. However, in the text we will often omit the subscripts to avoid notational clutter, assuming that the reader can reintroduce them as necessary.

With this bookkeeping device, the new INSERT function is exactly like the old one with two changes. First, we must use BUMPEQ instead of BUMP on the columns (this change occurs as line (11) of the code in Sect. 3). For example,

$$I_{\textcircled{3}}\left(\begin{array}{cccccccc} 1 & 1 & 1 & 2 & 3 & 3 & \textcircled{7} & 7 \\ & 2 & \textcircled{3} & 3 & \textcircled{5} & \textcircled{7} & 7 & \\ & & 4 & \textcircled{5} & 5 & & & \\ & & & 6 & 6 & & & \end{array}\right) = \left(\begin{array}{ccccccccc} 1 & 1 & 1 & 2 & \textcircled{3} & 3 & \textcircled{7} & 7 & 7 \\ & 2 & \textcircled{3} & 3 & 3 & 6 & \textcircled{7} & & \\ & & 4 & \textcircled{5} & 5 & & & & \\ & & & \textcircled{5} & 6 & & & & \end{array}\right)$$

because insertion of $\textcircled{3}_1$ into

$$\begin{array}{cccccccc} 1 & 1 & 1 & 2 & 3 & 3 & \textcircled{7}_1 & 7 \\ & 2 & \textcircled{3}_2 & 3 & \textcircled{5}_1 & \textcircled{7}_2 & 7 & \\ & & 4 & \textcircled{5}_2 & 5 & & & \\ & & & 6 & 6 & & & \end{array}$$

yields

$$\begin{array}{ccccccccc} 1 & 1 & 1 & 2 & \textcircled{3}_1 & 3 & \textcircled{7}_1 & 7 & 7 \\ & 2 & \textcircled{3}_2 & 3 & 3 & 6 & \textcircled{7}_2 & & \\ & & 4 & \textcircled{5}_1 & 5 & & & & \\ & & & \textcircled{5}_2 & 6 & & & & \end{array}$$

(we are justified in using the same symbol for this operator since I_x restricted to standard tableaux is the same as the one defined previously).

The second change only occurs if an $\boxed{m}$ bumps another $\boxed{m}$ from the diagonal and one (or both) is circled. In this case the circle is erased from the largest circled element of the pair before continuing with the column bumping. Examples of the two cases when this can occur are

$$I_1\begin{pmatrix}1 & 1 & \textcircled{2}\\ & 2 & \end{pmatrix} = \begin{matrix}1 & 1 & 1\\ & 2 & 2\end{matrix}$$

because the ② entering the second row on the diagonal loses its circle and

$$I_1\begin{pmatrix}1 & 1 & \textcircled{2}\\ & \textcircled{2} & \end{pmatrix} = \begin{matrix}1 & 1 & 1\\ & \textcircled{2} & 2\end{matrix},$$

since the ② leaving the second row diagonal cell has its circle erased (this is the larger of the two ②'s in the modified tableau). Hence the **else** clause at line (9) of INSERT should begin with

if $(y_{k-1} = \textcircled{m})$ **and** $(y_k = m)$ **then** UNCIRCLE (t_{k-1k-1});

if $(y_{k-1} = \textcircled{m})$ **and** $(y_k = \textcircled{m})$ **then** UNCIRCLE (y_k);

before starting the **repeat** (we assume the tableau is now named T, not P). It is easy to see that, even with the change, T remains a generalized tableau after insertion. It is also important to note that, for all $\boxed{m}$, the circling of the lowest leftmost $\boxed{m}$ in T is invariant under insertion. This fact will be important in the sequel.

Just as Knuth's algorithm puts pairs of generalized Young tableau in one-to-one correspondence with matrices having entries in $\mathbb{P} \cup \{0\}$, we can use *circled matrices* which are those with entries in $\mathbb{P}^0 \cup \{0\}$ and obtain

THEOREM 8.1 (shifted Knuth correspondence). *There is a bijection between circled matrices $A = (a_{ij})$ and pairs (T, U) of generalized shifted tableaux such that* shape $T =$ shape U, *U has no circles on its diagonal and j appears $\sum_i a_{ij}$ times in T while i appears $\sum_j a_{ij}$ times in U (we add elements of $\mathbb{P}^0$ ignoring their circles).*

Proof. We first convert A to two-line notation as follows. Take a_{ij} copies of the pair $\frac{i}{j}$ and put these columns in lexicographic order. If a_{ij} was circled then circle the first j in the string of $\frac{i}{j}$ columns. For example,

$$A = \begin{pmatrix}0 & \textcircled{3}\\ 2 & \textcircled{1}\\ 1 & 0\end{pmatrix} \text{ has a two-line notation } \begin{matrix}1 & 1 & 1 & 2 & 2 & 2 & 3\\ \textcircled{2} & 2 & 2 & 1 & 1 & \textcircled{2} & 1\end{matrix}.$$

If $\frac{m_1 m_2 \cdots m_n}{x_1 x_2 \cdots x_n}$ is the notation for A then INSERT the x_k into T and place the m_k in U so as to preserve shape, and flag which moves were non-

Schensted. Hence the only real difference between the procedures SHIFTED_KNUTH and SHIFTED_SCHENSTED is at line (7) which becomes

(7′) **if** *schen* **then** $u_{ij} := m_k$ **else** $u_{ij} := m_k$;

It is not clear that line (7′) preserves the tableau properties in U. However, this follows from

LEMMA 8.2. *Given a generalized shifted tableau T and $x_1, x_2 \in \mathbb{P}^0$, let (i_1, j_1) and (i_2, j_2) be the cells added on the boundary in passing from T to $T' = I_{x_1}(T)$ and from T' to $I_{x_2}(T')$, respectively. If $x_1 \leqslant x_2$ as subscripted elements then*

(i) *I_{x_1} Schensted implies that I_{x_2} is Schensted and $j_1 < j_2$,*

(ii) *I_{x_2} non-Schensted implies that I_{x_1} is non-Schensted and $i_2 > i_1$.*

Proof. (i) Under these conditions, the path for I_{x_2} always lies to the right of the path for I_{x_1} (proof as for the ordinary Knuth correspondence). Thus the second insertion cannot displace a diagonal element and $j_1 < j_2$.

(ii) If I_{x_2} is non-Schensted then I_{x_1} cannot be Schensted by part (i). Furthermore the row portion of the second insertion is still to the right of the first. This forces the column displacements for I_{x_1} to be above those for I_{x_2} and so $i_2 > i_1$. ∎

We now prove that U satisfies (T1) and (T2) by induction. Suppose $\begin{smallmatrix} m_1 \cdots m_{n-1} \\ x_1 \cdots x_{n-1} \end{smallmatrix}$ maps to (T', U') and consider the U corresponding to $T = I_{x_1}(T')$. If $m_{n-1} < m_n$ then $\boxed{m_n}$ is placed on the boundary of U' and is larger than any other entry so U is a tableau. If $m_{n-1} = m_n = m$ then by the lexicographic ordering we have $x_{n-1} \leqslant x_n$ and Lemma 8.2 applies. Thus when I_{x_n} is non-Schensted, all previous insertions obtained from pairs $\begin{smallmatrix} m \\ x_i \end{smallmatrix}$ must be non-Schensted by (ii). So only ⓜ's appear in U' and the new ⓜ added by the last insertion is in a lower row preserving the tableau properties. On the other hand, when I_{x_n} is Schensted there are two possibilities. Either $I_{x_{n-1}}$ is also Schensted, in which case an m is added to U' in a position to the right of all other m's by (i) so that (T1), (T2) hold. Otherwise $I_{x_{n-1}}$ is non-Schensted and, reasoning as before, the largest element in U' is ⓜ so an m can be added anywhere on the border.

To finish the proof of the theorem we must verify that one can recover a circled matrix from its pair of tableaux. The construction is similar to the inverse for the shifted Schensted map and only two additional points need discussion.

First of all, since elements may appear with repetition in U, we must know which of the possible largest elements to remove. If the maximum value is m (resp. $\textcircled{m}$) then by Lemma 8.2, again the rightmost (resp. lowest) one was the last to enter and should be removed first.

Second, we must be sure that we can reconstruct the circling in T if an $\boxed{m}$ from a column bumps out another $\boxed{m}$ on the diagonal. The diagonal $\boxed{m}$ can be either circled or not but the column $\boxed{m}$ must be uncircled (otherwise in the modified tableau we would have a smaller entry bumping out a larger one, which is impossible). Thus, there are two cases corresponding to those discussed for insertion. If both $\boxed{m}$'s are uncircled, then the entry displaced from the diagonal becomes circled before bumping into the row above. If only the diagonal element is circled then a circle is placed on the incoming m. Thus the insertion procedure is invertible as promised. ∎

In working out an example, rather than remodifying the tableau before each insertion, it is easier to modify the sequence $\pi = x_1 x_2 \cdots x_n$ once and remove the subscripts at the end. The *modified sequence* corresponding to π is obtained by labeling the $\textcircled{1}$'s in π right to left with 1, 2, 3,..., labeling the $\textcircled{2}$'s in like manner and continuing. So given

$$A = \begin{pmatrix} 0 & \textcircled{1} & \textcircled{2} \\ 1 & \textcircled{1} & 0 \\ \textcircled{2} & 0 & 0 \end{pmatrix}$$

we obtain the two-line array

$$\begin{matrix} 1 & 1 & 1 & 2 & 2 & 3 & 3 \\ \textcircled{2} & \textcircled{3} & 3 & 1 & \textcircled{2} & \textcircled{1} & 1 \end{matrix} \text{ which is modified to } \begin{matrix} 1 & 1 & 1 & 2 & 2 & 3 & 3 \\ \textcircled{2}_2 & \textcircled{3}_1 & 3 & 1 & \textcircled{2}_1 & \textcircled{1}_1 & 1 \end{matrix}.$$

Now the algorithm constructs the following tableaux

$$T\colon\ \textcircled{2}_2,\ \ \textcircled{2}_2\ \textcircled{3}_1,\ \ \textcircled{2}_2\ \textcircled{3}_1\ 3,\ \ 1\ \textcircled{2}_2\ \textcircled{3}_1\ 3,\ \ \begin{array}{cccc} 1 & \textcircled{2}_1 & \textcircled{3}_1 & 3 \\ & \textcircled{2}_2 & & \end{array},\ \ \begin{array}{ccccc} 1 & 1 & \textcircled{2}_1 & \textcircled{3}_1 & 3 \\ & \textcircled{2}_2 & & & \end{array},\ \ \begin{array}{ccccc} 1 & 1 & 1 & \textcircled{3}_1 & 3 \\ & \textcircled{2}_1 & 2 & & \end{array}$$

$$U\colon\ 1,\ \ 1\ 1,\ \ 1\ 1\ 1,\ \ 1\ 1\ 1\ \textcircled{2},\ \ \begin{array}{cccc} 1 & 1 & 1 & \textcircled{2} \\ & 2 & & \end{array},\ \ \begin{array}{ccccc} 1 & 1 & 1 & \textcircled{2} & \textcircled{3} \\ & 2 & & & \end{array},\ \ \begin{array}{ccccc} 1 & 1 & 1 & \textcircled{2} & \textcircled{3} \\ & 2 & \textcircled{3} & & \end{array}$$

so finally

$$(T, U) = \left(\begin{array}{ccccc} 1 & 1 & 1 & \textcircled{3} & 3 \\ & \textcircled{2} & 2 & & \end{array},\ \begin{array}{ccccc} 1 & 1 & 1 & \textcircled{2} & \textcircled{3} \\ & 2 & \textcircled{3} & & \end{array} \right).$$

For any finite multiset ("set" with repetitions) S of elements from $\mathbb{P}^0$, with maximum $\boxed{m}$, the *content of* S is $c(S) = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$, where a_i is the number of $\boxed{i}$'s in S. In particular, if S is any sort of array then we can talk about the content of the array. For the tableau T above we have

$c(T)=(1^3, 2^2, 3^2)$. Also define the *monomial corresponding to* T to be $m(T)=x_1^{x_1}x_2^{a_2}\cdots x_m^{a_m}$ and form the generating function

$$s_\mu(\mathbf{x})=\sum_R m(R),$$

where the sum is over all (left-justified) generalized Young tableaux of shape μ. $s_\mu(\mathbf{x})$ is called a *Schur function* and plays a crucial role in the ordinary representation theory of $\mathscr{G}_n$. Knuth's correspondence can be used to derive Cauchy's identity for the s_μ, namely

$$\prod_{i,j=1}^{\infty} \frac{1}{1-x_i y_j}=\sum_\mu s_\mu(\mathbf{x})\, s_\mu(\mathbf{y}),$$

where the sum is over all partitions μ and $s_\mu(\mathbf{y})$ is the Schur function in the indeterminates y_i, $i\in\mathbb{P}$.

The shifted analog of the Schur functions are a special case of the Hall–Littlewood polynomials called *Schur Q-functions* and defined by

$$Q_\lambda(\mathbf{x})=\sum_T m(T)$$

as T runs over all generalized shifted tableaux of shape λ. The Q_λ satisfy the following identity:

COROLLARY 8.3.

$$\prod_{i,j=1}^{\infty} \frac{1+x_i y_j}{1-x_i y_j}=\sum_\lambda \frac{1}{2^{l(\lambda)}}\, Q_\lambda(\mathbf{x})\, Q_\lambda(\mathbf{y}),$$

where the sum is over all strict partitions λ.

Proof. The left-hand side enumerates all circled matrices A. The denominator counts matrices with entries in $\mathbb{P}\cup\{0\}$, while the numerator accounts for the circles. For choosing the $x_i y_j$ from $1+x_i y_j$ amounts to circling a_{ij} and adding one to it (so that a_{ij} becomes positive). Otherwise the 1 is chosen, leaving the entry alone.

By Theorem 8.2. we will be done if we can show that the right side corresponds to (T, U) pairs. $Q_\lambda(\mathbf{x})$ counts all possible generalized shifted tableaux T by definition. To see why the remaining factor enumerates the U-tableaux, note that the elements on the diagonal of an arbitrary tableau are all at the lower left end of their rim hooks and so can be circled or not. Thus after collecting like monomials, all terms in the sum $Q_\lambda(\mathbf{y})$ will be divisible by $2^{l(\lambda)}$. It follows that $(1/2^{l(\lambda)})\, Q_\lambda(y)$ counts shifted tableaux with no diagonal circles. ∎

We remark that one can rederive Corollary 3.2 by equating the coefficients of $x_1 x_2\cdots x_n y_1 y_2\cdots y_n$ on both sides of Corollary 8.3.

9. Labeling Operators

We now concentrate on the *T-tableau corresponding to* σ, where $\sigma = x_1 x_2 \cdots x_n \in (\mathbb{P}^0)^n$, which is defined as $T(\sigma) = I_{x_n} \cdots I_{x_1}(\phi)$. It would be tempting to try to prove a lifting lemma for T-tableaux connecting the original and shifted versions of Knuth's algorithm. Unfortunately, when a generalized shifted tableau is shift-symmetrized the result is not column strict. One could fix this by modifying the rows as was done for the columns, but then all the entries of the tableau have essentially become distinct. Thus we might as well replace the generalized tableau by a standard one and then appeal to the results proved for the shifted Schensted map. This is the purpose of a labeling operator.

Given a generalized shifted tableau T, define the *labeling of* T to be the triple

$$\mathrm{L}(T) = (P, c, s),$$

where

(i) P is the standard shifted tableau defined as follows. If there are k ①'s in T then replace them with the numbers $1, 2, \ldots, k$ working from top to bottom. Now replace the 1's in T by $k+1$, $k+2, \ldots, k+l$ moving from left to right. Continue with the ②'s replaced by $k+l+1$, etc. until every element is labeled.

(ii) $c = c(T)$ is the content of T.

(iii) $s = s(T)$ is the set of all integers m such that the lower-leftmost occurrence of $\boxed{m}$ in T is circled. For example, when

$$T = \begin{array}{cccccccc} \textcircled{1} & 1 & 1 & \textcircled{2} & \textcircled{3} & 3 & 3 & \textcircled{4} \\ & 2 & 2 & 2 & \textcircled{3} & \textcircled{4} & 4 & 4 \\ & & \textcircled{3} & 3 & 3 & & & \\ & & & 4 & 4 & & & \end{array}$$

then

$$L(T) = \left(\begin{array}{cccccccc} 1 & 2 & 3 & 4 & 8 & 13 & 14 & 15, \\ & 5 & 6 & 7 & 9 & 16 & 19 & 20 \\ & & 10 & 11 & 12 & & & \\ & & & 17 & 18 & & & \end{array} (1^3, 2^4, 3^7, 4^6), \{1, 3\} \right).$$

The components c and s of the labeling appear so that we lose no information in passing from T to $L(T)$. Specifically,

Lemma 9.1. *The labeling operator for generalized shifted tableaux is injective.*

Proof. We must show that given a triple (P, c, s) we can reconstruct the tableau T from which it came. We will find the positions of the $\boxed{1}$'s, the case for general $\boxed{m}$ being similar. Since $c(T) = (1^{a_1}, 2^{a_2}, ..., m^{a_m})$ the cells labeled $1, 2, ..., a_1$ in P contained either ① or 1 in T. Let the cells containing these elements have coordinates $(i_1, j_1), (i_2, j_2), ..., (i_{a_1}, j_{a_1})$ respectively. Since P came from a generalized tableau we must have $i_1 < i_2 < \cdots < i_k \geqslant i_{k+1} \geqslant \cdots \geqslant i_{a_1}$ for some k. Thus the entries $1, 2, ..., k-1$ should be replaced by ① while $k+1, k+2, ..., a_1$ should become 1. The only question occurs with (i_k, j_k) but this is easily resolved for $t_{i_k j_k} =$ ① if and only if $1 \in s(T)$. ▮

Analogously, we have a *labeling operator for sequences* $\sigma = x_1 x_2 \cdots x_n \in (\mathbb{P}^0)^n$,

$$L(\sigma) = (\pi, c, s),$$

where

(i) $\pi \in \mathcal{G}_n$ is the permutation obtained from σ in almost the same way as P is obtained from T above. The only difference is that when replacing ⓜ's one works from right to left instead of top to bottom.

(ii) $c = c(\sigma)$ is the content of σ.

(iii) $s = s(\sigma)$ is the set of all integers whose leftmost occurrence is circled.

As an illustration, if $\sigma = 1\,1$ ② ① ② $2\,2$ ③ ① ②, then $L(\sigma) = (3\,4\,7\,2\,6\,8\,9\,10\,1\,5,\ (1^4, 2^5, 3^1),\ \{2, 3\})$. Reasoning as in Lemma 9.1, we have

LEMMA 9.2. *The labeling operator for sequences is injective.* ▮

Henceforth it will be assumed that all operators on standard tableaux and permutations have been extended to triples by letting them act on the first component, e.g., $P(\pi, c, s) = (P(\pi), c, s)$. The connection between the shifted Schensted and Knuth algorithms can now be stated.

THEOREM 9.3. *For any sequence* $\sigma \in (\mathbb{P}^0)^n$ *we have*

$$PL(\sigma) = LT(\sigma),$$

i.e., the labeling operator commutes with formation of tableaux by insertion.

Proof. We need to refine the labeling operators so that we can work at the bumping level. First consider sequence-tableau pairs (σ, T) and let

$$L(\sigma, T) = (\pi, P, c, s),$$

where

(i) π and P are obtained from σ and T respectively by, for each $\boxed{m}$, labeling the ⓜ's in σ followed by those in T (using the order defined above within the permutation or tableau) and then labeling the m's in T followed by those in σ.

(ii) $c = c(\sigma \cup T)$.

(iii) $s \subseteq \mathbb{P}$ with $i \in s$ exactly when T contains an $\boxed{i}$ and $i \in s(T)$ or T does not contain an $\boxed{i}$ and $i \in s(\sigma)$.

The reader should think of σ as the suffix of a larger sequence whose initial segment has been converted to T by insertion. In fact, if $\sigma = xw \cdots v$ then let $\hat{\sigma} = w \cdots v$ be σ with its first element removed and define $I_x(\sigma, T) = (\hat{\sigma}, I_x T)$.

It is clear from the definition that the labeling operator for pairs reduces to the ones for sequences or tableaux if $T = \phi$ or $\sigma = \phi$, respectively. Hence, to prove the theorem it suffices to show that $I_l L(\sigma, T) = L I_x(\sigma, T)$ for any σ, T with $x \in \mathbb{P}^0$ being the first element of σ and $l \in \mathbb{P}$ being the first element of $L(\sigma, T)$.

Now enlarge the sequence-tableau pairs to quadruples (σ, T, y, k), where $y \in \mathbb{P}^0$ and k is the index of the row or column of T into which y is being bumped. (There is never any ambiguity as it is impossible to find two different positions for y, one in the kth row and one in the kth column, both of which maintain the tableaux properties). Our final labeling operator is

$$L(\sigma, T, y, k) = (\pi, P, l, k, c, s)$$

defined by

(i) π and P are constructed from σ and T as above until we get to the value $y = \boxed{m}$ in T. To determine l, the label of y, we continue as follows. If we are bumping into row k then for $y =$ ⓜ (resp. $y = m$) first label all similar elements above (resp. below and in) the kth row, then y itself, and end with those below and in (resp. above) row k. If ⓜ (resp. m) is to be bumped into the kth column then label all similar elements to the right of and in (resp. left of) column k, label y, and continue with the entries to the left of (resp. right of and in) the kth column. Finally, label the rest of the elements of σ and T as before.

(ii) $c = c(\sigma \cup T \cup y)$.

(iii) When $y \neq \boxed{i}$ we determine whether $i \in s$ by using the same rules as for pairs (σ, T). In the contrary case, consider the cell into which y will be placed and see if there are any $\boxed{i}$'s already in T below, to the left of, or in this position. If so, then $i \in s$ exactly when $i \in s(T)$. Otherwise $i \in s$ if and only if y is circled.

If applying the appropriate bumping function to y and T yields output z and a new tableau T' then define

$$B_y(\sigma, T, y, k) = \begin{cases} (\sigma, T', z, k+1) & \text{if } z \neq \infty \\ (\hat{\sigma}, T', x, 1) & \text{otherwise, } \sigma = x\hat{\sigma}. \end{cases}$$

This reduces the problem to proving

$$B_l L(\sigma, T, y, k) = L B_y(\sigma, T, y, k) \tag{9.1}$$

where l is the label of y.

Assume that we are bumping into the kth row (the column case is similar). First of all, we must show that y and l are placed in the same positions in T and P, respectively. This will happen if $t_{kj} \leqslant y$ is equivalent to $p_{kj} \leqslant l$. If $y = ⓜ$ then $t_{kj} \leqslant y$ as modified integers (circled elements subscripted) implies that $t_{kj} \leqslant m - 1$. Thus t_{kj} is labeled before y as desired. If $y = m$ then it is also possible that $t_{kj} = m$ but then by condition (i) for quadruples, we label all the m's in row k before y so again $p_{kj} \leqslant l$. In the same way, it can be shown that $t_{kj} > y$ implies $p_{kj} > l$.

Now we must check that the label of y remains l after it is bumped into row k. If $y = ⓜ$ then, after bumping, y is labeled after labeling the ⓜ's in rows 1 through $k-1$. This is the same as the order in which these elements are labeled in (σ, T, y, k) by condition (i). Note that this holds even if y's circle is removed because of bumping an m on the diagonal. In that case, there are no ⓜ's below y, and y becomes the lowest leftmost m which is labeled directly after all the ⓜ's above. If $y = m$ then y will be placed after any other m's in row k and so will be the last m to be labeled in the rows below and including the kth. This again agrees with the label of y in $L(\sigma, T, y, k)$.

Similarly, one may check that the other elements of T and the element z coming out of row k receive the same labels before and after bumping. In particular, when $z = \infty$ this follows from the fact that condition (i) for $L(\hat{\sigma}, T', x, 1)$ reduces to the corresponding one for the pair $L(x\hat{\sigma}, T')$. Hence the first four components on both sides of Eq. (9.1) agree.

Clearly $c(\sigma \cup T \cup y)$ is invariant under bumping since only circling can change. To see that s remains the same, recall that the circling of the lowest leftmost entry of T is never changed by insertion. Also condition (iii) is constructed so that even when this element is in transition between rows, its status can still be ascertained by consulting s. This completes the verification of Eq. (9.1) and the proof of the theorem. ∎

10. Greene's Invariant

Given a modified sequence σ, an *increasing subsequence relative to $\bar{\sigma}\sigma$* is a subsequence $x_1 x_2 \cdots x_n$ of $\bar{\sigma}\sigma$ such that $x_1 \cdots x_i$ is a subsequence of $\bar{\sigma}$, $x_{i+1} \cdots x_n$ is a subsequence of σ, and

(A1) $x_1 \cdots x_i$ is strictly increasing,

(A2) $x_i x_{i+1} \cdots x_n$ is weakly increasing.

Dually, a *decreasing subsequences relative to $\bar{\sigma}\sigma$* satisfies

(D1) $x_1 \cdots x_i$ is weakly decreasing,

(D2) $x_i x_{i+1} \cdots x_n$ is strictly decreasing.

A *k-increasing* (resp. *k-decreasing*) *subsequence relative to $\bar{\sigma}\sigma$* is one which is the union of k increasing (resp. decreasing) subsequences relative to $\bar{\sigma}\sigma$. We also let $a_k(\bar{\sigma}\sigma)$ (resp. $l_k(\bar{\sigma}\sigma)$) denote the length of the longest k-increasing (resp. k-decreasing) subsequence relative to $\bar{\sigma}\sigma$.

For this section only we will let $L(\sigma)$, $L(T)$ denote the first component of the usual labeling triple.

THEOREM 10.1. *Let σ be a modified sequence, then*

$$a_k(\bar{\sigma}\sigma) = a_k(\overline{L(\sigma)}\, L(\sigma))$$

and

$$d_k(\bar{\sigma}\sigma) = d_k(\overline{L(\sigma)}\, L(\sigma)).$$

Proof. We will prove the assertion about a_k; the same method can be used to obtain the result for d_k. First we show that $a_k(\bar{\sigma}\sigma) \geqslant a_k(\overline{L(\sigma)}\, L(\sigma))$ by proving that for every weakly k-increasing subsequence of $\overline{L(\sigma)}\, L(\sigma)$ the corresponding subsequence in $\bar{\sigma}\sigma$ is k-increasing relative to $\bar{\sigma}\sigma$. This will follow if we can demonstrate that given $l \leqslant m$ in $\overline{L(\sigma)}\, L(\sigma)$ with l to the left of m then their parent elements x and y in $\bar{\sigma}\sigma$ satisfy the proper inequality needed for (A1)–(A2).

Note that we must always have $x \leqslant y$ because L is order preserving. Thus (A2) is satisfied and it suffices to show that when $x, y \in \bar{\sigma}$ then we cannot have $x = y$. But equal elements are labeled from right to left in passing to $\overline{L(\sigma)}$ and so x to the left of y would imply $l > m$, a contradiction. Hence (A1) holds as well.

Unfortunately, the same argument cannot be used to prove that $a_k(\bar{\sigma}\sigma) \leqslant a_k(\overline{L(\sigma)}\, L(\sigma))$. For example, if $\sigma = 11$ then $\bar{\sigma}\sigma = 1111$, and the subsequence consisting of the first, third, and fourth 1's is increasing relative to $\bar{\sigma}\sigma$. However, $\overline{L(\sigma)}\, L(\sigma) = 2112$ so that the corresponding subsequence 212 is not increasing. Notice, though, that we can transform our subsequence in $\bar{\sigma}\sigma$ to the one containing the 1's in positions two, three, and four and obtain the weakly increasing subsequence 112 in $\overline{L(\sigma)}\, L(\sigma)$. The rest of the proof will show that this type of transformation can always be done.

Let $x_1 x_2 \cdots x_n = \sigma_1 \uplus \sigma_2 \uplus \cdots \uplus \sigma_k$, where each σ_i is an increasing subsequence relative to $\bar{\sigma}\sigma$ and $\uplus$ denotes disjoint union. Also define the *left and right halves* of σ_i to be $\sigma_i^{\mathrm{L}} = \sigma_i \cap \bar{\sigma}$, $\sigma_i^{\mathrm{R}} = \sigma_i \cap \sigma$ for all i. Finally, let σ_1 be the subsequence whose last element in $\bar{\sigma}$ is the farthest to the left among all the σ_i.

By induction we can assume that $\sigma_2, ..., \sigma_k$ have been transformed so that the corresponding subsequences in $\overline{L(\sigma)}\, L(\sigma)$ are weakly increasing and that the elements of $\dot{\bigcup}_{i \geqslant 2} \sigma_i$ remain unchanged (although they may be shuffled into different subsequences) with one exception. We are permitted to replace a rightmost element of σ_i^{L}, $i \geqslant 2$, with an element farther to the right in $\bar{\sigma}$. These conditions insure that σ_1 will not get in the way of the transformation of $\sigma_2, ..., \sigma_k$.

Now considerations like those in the first half of the proof show that the only possible descent that can occur in the labeled version of σ_1 is between x and y_1 where x (resp. y_1) is the rightmost (resp. leftmost) element of σ_1^{L} (resp. σ_1^{R}). In fact this will only happen when $x = y_1 = n$ and y_1 is closer to the beginning of σ than x is to the end of $\bar{\sigma}$.

Let the labels of x and y_1 be l and m_1 respectively, $l > m_1$, and consider all n's in $\bar{\sigma}$ with labels less than or equal to m_1. If some such n is not a member of $\dot{\bigcup}_{i \geqslant 2} \sigma_i$ we can replace x by this n and we are done. Otherwise each such n is already in some σ_i, $i \geqslant 2$. But then by the pigeonhole principle there must be one σ_i such that σ_i^{R} begins with an element y_2 whose label, m_2, is greater than m_1. Form the new subsequences $\sigma_1 := \sigma_1^{\mathrm{L}} \cup \sigma_i^{\mathrm{R}}$ and $\sigma_i := \sigma_i^{\mathrm{L}} \cup \sigma_1^{\mathrm{R}}$, where $:=$ is, as usual, the Pascal assignment symbol. By construction, σ_i still corresponds to a weakly increasing subsequence in $\overline{L(\sigma)}\, L(\sigma)$. If the labeling of σ_1 is now weakly increasing, then we are done. Otherwise, keep repeating this process, starting with the new σ_1 and y_2. It is easy to verify that at some point one of the two conditions for termination of the loop becomes satisfied so that σ_1 becomes labeled appropriately. ▮

We now have the analog of Greene's theorem for generalized shifted tableaux.

THEOREM 10.2. *If σ is a modified sequence and* shape $T(\sigma) = (\lambda_1, \lambda_2, ..., \lambda_l)$ *then for* $k \leqslant l$

$$a_k(\bar{\sigma}\sigma) = \lambda_1 + \lambda_2 + \cdots + \lambda_k + \binom{k+1}{2},$$

$$l_k(\bar{\sigma}\sigma) = \lambda_1 + \lambda_2 + \cdots + \lambda_k + \binom{k}{2}.$$

Furthermore, for $k > l$ we have

$$a_k(\bar{\sigma}\sigma) = \lambda_1 + \cdots + \lambda_l + \lambda'_{l+1} + \cdots + \lambda'_k + \binom{l+1}{2}$$

$$l_k(\bar{\sigma}\sigma) = \lambda_1 + \cdots + \lambda_l + \lambda'_l + \cdots + \lambda'_{k-1} + \binom{l}{2},$$

where λ'_j is the length of the jth column of the shifted shape of λ.

Proof. By commutativity of the labeling operator (Theorem 9.3)

$$(\lambda_1, \lambda_2, \ldots, \lambda_l) = \text{shape } T(\sigma) = \text{shape } LT(\sigma) = \text{shape } PL(\sigma).$$

But we have just proved that $a_k(\bar{\sigma}\sigma) = a_k(\overline{L(\sigma)}\, L(\sigma))$ and $d_k(\bar{\sigma}\sigma) = d_k(\overline{L(\sigma)}\, L(\sigma))$ so the theorem follows from the corresponding result for P-tableaux, Corollary 5.2. ∎

Finally, given $\sigma \in (\mathbb{P}^0)^n$ let $\sigma^{(z)}$ denote the subsequence $x_1, x_2, \ldots, x_k$ of σ consisting of all $x_i \leqslant z$. When z is circled, say $z = ⓜ$, assume that z has the highest possible subscript so that $x \leqslant z$ for any other $x = ⓜ$ in σ. Setting $\mathbf{a}(\sigma) = (a_1(\bar{\sigma}\sigma), a_2(\bar{\sigma}\sigma), \ldots, a_n(\bar{\sigma}\sigma))$ then the *Greene invariant* of σ is defined to be

$$G(\sigma) = (\mathbf{a}(\sigma^{(①)}), \mathbf{a}(\boldsymbol{\sigma}^{(1)}), \ldots, \mathbf{a}(\boldsymbol{\sigma}^{(ⓜ)}), \mathbf{a}(\sigma^{(m)})),$$

where m is the largest value in σ. It will now be an easy matter for the reader to prove

THEOREM 10.3. *Given $\sigma_1, \sigma_2 \in (\mathbb{P}^0)^n$ then $T(\sigma_1) = T(\sigma_2)$ if and only if $G(\sigma_1) = G(\sigma_2)$.*

11. ENCORE LE JEU DE TAQUIN

To obtain the jeu de taquin for modified skew tableaux we need only make one change in the MOVE function of Section 6. Specifically, suppose that diagonal cell (i, i) is to be filled during a move in tableau T and that t_{ii+1}, t_{i+1i+1} both have the value $\boxed{m}$. By condition (T1) we must have $t_{ii+1} = ⓜ$ while t_{i+1i+1} may be circled or not. In either situation t_{ii+1} will move into square (i, i) and t_{i+1i+1} will follow into $(i, i+1)$. If $t_{i+1i+1} = ⓜ$ then we remove its circle before it slides up. Otherwise $t_{i+1i+1} = m$ and the circle of t_{ii+1} is erased after it moves left. Thus the following lines should be added just after line (6) in the code of MOVE

if $(i=j)$ **then begin**
$\boxed{m} := t_{ii}$;
if $(t_{i+1i+1} = ⓜ)$ **then** UNCIRCLE $(t_{i+1,i+1})$;
if $(t_{i+1i+1} = m)$ **then** UNCIRCLE (t_{ii})
end; {**then**}

Examples of these two cases are given in Fig. 5.

Let SJ devote the shifted jeu de taquin applied to modified sequences, then

THEOREM 11.1. *Given a modified sequence σ then the tableau* $\mathrm{SJ}(\sigma)$ *is independent of the order in which the cells are filled and* $\mathrm{SJ}\cdot L(\sigma) = L\cdot \mathrm{SJ}(\sigma)$.

Proof. Consider some given order of filling the cells in the diagonal tableau corresponding to σ. If we can show that following the same order in $L(\sigma)$ will always ensure $\mathrm{SJ}\cdot L(\sigma) = L\cdot \mathrm{SJ}(\sigma)$ then since $\mathrm{SJ}\cdot L(\sigma)$ is independent of order (Theorem 6.2) and L is injective (Lemma 9.1) we will have that $\mathrm{SJ}(\sigma)$ is also independent.

Reducing the problem from the game to the move level and from there to a single slide (cf. the proof of Theorem 9.3) it becomes apparent that we need to prove the following. Consider a skew modified tableau T with a hole at cell (i, j) and $L(T) = (P, c, s)$, labeling the elements of T as one would a non-skew hole-less tableau. It is important to note that for tableau with holes, deciding whether $m \in s$ is done by first finding the column of smallest index in which an $\boxed{m}$ appears and then going to the bottom $\boxed{m}$ in that column, i.e., "leftmost" takes precedence over "lowest." If sliding to fill (i, j) in T and P yield T' and P', respectively, then it suffices to show $L(T') = (P', c, s)$.

First, the elements used to slide into (i, j) in both T and P should come from the same cell. This is trivially true if (i, j) is on the diagonal since then there is only one choice. Otherwise $i \neq j$ and we must verify that $t_{ij+1} \geqslant t_{i+1j}$ if and only if $p_{ij+1} > p_{i+1j}$. This follows from the fact that L is monotone increasing except when $t_{ij+1} = t_{i+1j} = m$. But in that case the $(i+1, j)$ entry is labeled before the $(i, j+1)$ one, so $p_{i+1j} < p_{ij+1}$ as required.

Next, if $y = \boxed{m}$ is the element used to fill (i, j) in T we must show that y gets the same label l before and after the slide. Assume that y slides left (the

FIG. 5. Special diagonal slides.

case where y slides up is similar) and that y's circling is unchanged. If $y = \textcircled{m}$ then y's label must remain l because it stays in the same row. If $y = m$ then there can be no m in column j so that all m's to the left of y before the slide still bear the same relationship afterwards. Since m's are labeled left to right, y will again receive label l. Finally, suppose that $y = \textcircled{m}$ slides onto the diagonal and becomes uncircled. Then y must have been the lowest $\textcircled{m}$ in T and becomes the leftmost m in T'. In either situation y is labeled directly after all the other $\textcircled{m}$'s.

The same sort of considerations show that the labels of the other entries of T remain the same after the slide. Thus $L(T')$ has first component P'.

Lastly, we must check that c and s are left invariant. The fact that $c(T) = c(T')$ is immediate because sliding does not change content, verifying that $s(T) = s(T')$ is again a case by case argument. Returning to our assumption that y is sliding left, suppose that y is the lowest leftmost occurrence of $\boxed{m}$ in T. Since y is even farther left after the slide it continues to represent that position in T'. Also y's circling cannot change in this case, so $m \in s(T)$ if and only if $m \in s(T')$ as desired. If y is not the lowest leftmost $\boxed{m}$ then there are two possibilities. One is that there could be a $z = \boxed{m}$ to the left of y in T. But then z must be in the same row or lower than y before and hence after the slide. The second option is to have all other $\boxed{m}$'s in T directly below y also in the same column. This forces $y = \textcircled{m}$ and $t_{i+1j} < \textcircled{m}$ when $i \neq j$, which contradicts the fact that y is to slide into cell (i, j). Thus under these conditions we must have $i = j$ so that $y = \textcircled{m}$ is sliding onto the diagonal with $t_{ij+1} = \boxed{m}$ being the lowest leftmost $\boxed{m}$ in T. It is easy to check that whether $t_{ij+1} = \textcircled{m}$ or m, the presence or absence of that integer in s will be preserved. ∎

Commutativity of L and SJ provides the final link in the chain of theorems needed to prove

THEOREM 11.2. *Given a modified sequence σ then* $\mathrm{SJ}(\sigma) = T(\sigma)$.

Proof. Applying Theorems 11.1, 6.2, and 9.3 in turn yields

$$L \cdot \mathrm{SJ}(\sigma) = \mathrm{SJ} \cdot L(\sigma) = P \cdot L(\sigma) = L \cdot T(\sigma).$$

Hence $\mathrm{SJ}(\sigma) = T(\sigma)$ by injectivity of the labeling operator (Lemma 9.1). ∎

12. The Knuth Relations Revisited

Given a modified sequence σ then the possible *Knuth relations* that can be applied to σ in order to form another modified sequence are

(SK1) replace xzy by zxy or vice versa if $x \leqslant y < z$, (SK2) replace yxz by yzx or vice versa if $x < y \leqslant z$,

(SK3) for the first two elements of σ we may

replace $\boxed{m}\ \boxed{n}$ by $\boxed{n}\ \boxed{m}$ or vice versa if $m \neq n$,

replace $\boxed{m}$ m by $\boxed{m}\ \textcircled{m}$ or vice versa,

i.e., we can change the circling of the second element while keeping the first the same. Knuth equivalence of modified sequences is defined in the usual way using the three relations above. By way of example

$$\begin{aligned} 3\ \textcircled{1}_2\ 4\ \textcircled{1}_1 &\equiv_{SK} 3\ \textcircled{1}_2\ \textcircled{1}_1\ 4 && \text{by (SK2)} \\ &\equiv_{SK} \textcircled{1}_2\ 3\ \textcircled{1}_1\ 4 && \text{by the first half of (SK3)} \\ &\equiv_{SK} \textcircled{1}_2\ \textcircled{1}_1\ 3\ 4 && \text{by (SK2)} \\ &\equiv_{SK} \textcircled{1}_1\ 1\ 3\ 4 && \text{by the second half of (SK3).} \end{aligned}$$

This equivalence relation extends to triples (π, c, s) by saying that $(\pi_1, c_1, s_1) \equiv_{SK} (\pi_2, c_2, s_2)$ means $\pi_1 \equiv_{SK} \pi_2$, $c_1 = c_2$ and $s_1 = s_2$. With this convention, the labeling operator respects Knuth equivalence.

THEOREM 12.1. *Given modified sequences σ_1 and σ_2 then $\sigma_1 \equiv_{SK} \sigma_2$ if and only if $L(\sigma_1) \equiv_{SK} L(\sigma_2)$.*

Proof. Suppose first that $\sigma_1 \equiv_{SK} \sigma_2$ with $L(\sigma_1) = (\pi_1, c_1, s_1)$ and $L(\sigma_2) = (\pi_2, c_2, s_2)$. Since none of the three relations change the content or circling of the leftmost element $\boxed{m}$ in a sequence it is clear that $c_1 = c_2$ and $s_1 = s_2$.

To prove that $\pi_1 \equiv_{SK} \pi_2$ it suffices to show that if σ_1 can be transformed into σ_2 by applying a single SKi above, $i = 1, 2, 3$, then π_1 can be turned into π_2 by the corresponding SKi in Section 7. Suppose that the transportation is of type (SK1) and the elements x, y, z are labeled l, m, n, respectively. If $x < y < z$ then $l < m < n$ because L is monotone and SK1 applies to π_1 and π_2. If $x = y$ then they must both be uncircled since the sequences are modified. But uncircled elements are labeled left to right and x is to the left of y in both triples of (SK1). So $l < m$ and (SK1) can again be applied. The argument if (SK2) is used is similar and there is nothing to prove for (SK3) since the first two elements of a permutation can always be interchanged regardless of magnitude.

For the "if" direction we proceed as before, keeping the same notation. Take the case where π_1 and π_2 differ by a transposition of type (SK1). Since L is monotone and $l < m < n$ we must have $x \leqslant y \leqslant z$. But if $y = z$ then we must have $m > n$ because of the left-to-right labeling of uncircled elements and this is a contradiction. Thus $x \leqslant y < z$ and (SK1) can be applied to σ_1 and σ_2.

Omitting the (SK2) case, which presents nothing new, consider what happens if initial elements x and y with labels l and m are switched. If x and y are not versions (circled or uncircled) of the same integer then they can certainly be interchanged by the first half of (SK3). If, on the other hand, we have $x=y=n$ then $xy=nn$ is labeled lm, where $m=l+1$. Now $ml=l+1l$ and if the second n is to have a smaller label then y becomes ⓝ. Meanwhile $x=n$ cannot change since $s_1=s_2$ so nn is replaced by nⓝ which is consistent with the second half of (SK3). The same sort of reasoning can be applied to the other three cases where x and y have the same underlying integer. ∎

We conclude this section with a result that the reader has surely anticipated.

THEOREM 12.2. *Given modified sequences σ_1 and σ_2 then $\sigma_1 \equiv_{SK} \sigma_2$ if and only if $T(\sigma_1)=T(\sigma_2)$.*

Proof. We have the following string of equivalent statements:

$$\begin{aligned}
\sigma_1 \equiv_{SK} \sigma_2 \quad & \text{if and only if } L(\sigma_1) \equiv_{SK} L(\sigma_2) \text{ by Theorem 12.1} \\
& \text{if and only if } PL(\sigma_1)=PL(\sigma_2) \text{ by Theorem 7.2} \\
& \text{if and only if } LT(\sigma_1)=LT(\sigma_2) \text{ by Theorem 9.3} \\
& \text{if and only if } T(\sigma_1)=T(\sigma_2) \text{ by Lemma 9.1.}
\end{aligned}$$

13. STANLEY'S CONJECTURE

The *Schur P-functions* are defined by $P_\lambda(\mathbf{x})=(1/2^{l(\lambda)})\,Q_\lambda(\mathbf{x})$ and count generalized shifted Young tableaux with none circled on the diagonal. $P_\lambda(\mathbf{x})$ is a symmetric function in the variables x_i, $i\in\mathbb{P}$, a fact which is not obvious from the combinatorial interpretation but follows trivially from the equivalent algebraic definition, see MacDonald [M]. The $s_\mu(\mathbf{x})$ are also symmetric and form an integral basis for the space of all symmetric functions. Thus for a given partition λ of n we have $P_\lambda(\mathbf{x})=\sum_{\mu\vdash n} a_{\lambda\mu}s_\mu(\mathbf{x})$ for some scalars $a_{\lambda\mu}\in Z$. Stanley noted that the foregoing machinery could be used to prove a conjecture of his [St] that the $a_{\lambda\mu}$ are actually nonnegative.

THEOREM 13.1. *If $P_\lambda(x)=\sum_\mu a_{\lambda\mu}s_\mu(x)$ then $a_{\lambda\mu}\in\mathbb{P}\cup\{0\}$.*

Proof. Take any fixed generalized shifted tableau T of shape λ and consider all pairs (T, U_j) as U_j varies through all generalized tableaux of the same shape with no diagonal circles. Note that $P_\lambda(\mathbf{x})$ enumerates this set via the second component. Using the inverse shifted Knuth map, we obtain a set of matrices in modified two-line notation which form, by definition, a shifted Knuth equivalence class E. But the possible Knuth transpositions (K1)–(K2) for left-justified tableaux are a subset of those for generalized shifted tableaux (SK1)–(SK3). Thus E can be partitioned into ordinary Knuth equivalence classes $E = E_1 \uplus E_2 \uplus \cdots \uplus E_k$. Now applying the regular Knuth correspondence to E we see that each E_i becomes a set of pairs (R_i, S_{i_j}), where R_i is a fixed generalized left-justified tableau and S_{i_j} takes on all possible second tableau values. Hence if $\mu_i = \text{shape}\, R_i = \text{shape}\, S_{i_j}$ for all j then $s_{\mu_i}(\mathbf{x})$ counts the S_{i_j} and $P_\lambda(\mathbf{x}) = s_{\mu_1}(\mathbf{x}) + s_{\mu_2}(\mathbf{x}) + \cdots + s_{\mu_k}(\mathbf{x})$. ∎

For example, if $\lambda = (3, 1)$ and we choose $T = {}^{1\ 2\ 4}_{\ \ \ 3}$ then $E = \{1342, 3142, 1324, 3124, 3412, 4312, 4132, 1432\}$, where the top rows of the two-line arrays take on all possible values and so have been omitted for lack of space. Now

$$E = \{1342, 1324, 3124\} \uplus \{3142, 3412\} \uplus \{4312, 4132, 1432\}$$

whence the elements of each subset map to tableau pairs with first components

$$\begin{matrix} 1 & 2 & 4 \\ 3 & & \end{matrix}\ , \quad \begin{matrix} 1 & 2 \\ 3 & 4 \end{matrix}\ , \quad \text{and} \quad \begin{matrix} 1 & 2 \\ 3 & \\ 4 & \end{matrix}\ ,$$

respectively. Thus $P_{(3,1)}(\mathbf{x}) = s_{(3,1)}(\mathbf{x}) + s_{(2,2)}(\mathbf{x}) + s_{(2,1,1)}(\mathbf{x})$. It should be noted that a character-theoretic proof of Theorem 12.1 can be given. In fact, Stanley and Morris [St, Mo2] independently observed that the result follows easily from an earlier work of Morris [Mo1].

With a little more work, we can extract further information about the coefficients $a_{\lambda\mu}$. Take the shifted tableau T in the proof of Theorem 13.1 to be standard so that all the left-justified R_i will have entries $1, 2, \ldots, n$ as well. Given μ, the number of times s_μ occurs in the sum for P_λ is just the number of R_i having shape μ. We have proved

COROLLARY 13.2. *If $P_\lambda(x) = \sum_\mu a_{\lambda\mu} s_\mu(x)$ then $a_{\lambda\mu} \leqslant f_\mu$, where f_μ is the number of ordinary standard Young tableaux of shape μ.* ∎

14. Comments and Open Problems

We have seen that many of the properties of the Robinson–Schensted correspondence have analogs for the shifted case. However, the original algorithm is so rich in structure that there is still much to be done.

(1) The reader may have noticed that most of our theorems have dealt with the first tableau of the pairs (P, Q) and (T, U). There are many results for the unshifted correspondence that involve both arrays but these do not seem to carry over easily to our setting. For example, if $\pi \in \mathscr{G}_n$ then a result of Robinson [R] states that applying Robinson–Schensted to π^{-1} interchanges the tableaux. On the other hand, there does not seem to be any simple connection between $(P(\pi), Q(\pi))$ and $(P(\pi^{-1}), Q(\pi^{-1}))$. It should be noted that MacLarnan [Ml] has developed a whole family of shifted Schensted-like correspondences, all of which have the property that, with appropriate circling changes, $P(\pi^{-1}) = Q(\pi)$ and $Q(\pi) = P(\pi^{-1})$. Unfortunately, these maps do not enjoy other important properties such as the Knuth relations.

(2) The difficulty in working with Q-tableaux is the absence of a lifting lemma relating the left-justified and shifted cases. A candidate for such a result is the following conjecture of Shore [Wo]. Given $Q(\pi)$, consider the corresponding shift-symmetrized tableau where the circling below the diagonal has been reversed. Now delete all circled elements and perform slides to eliminate the holes. Shore conjectures that the final array will be the second tableau of the ordinary Robinson–Schensted map. For example, suppose $\pi = 3142$ so that

$$\pi \to (P(\pi), Q(\pi)) = \left(\begin{matrix} 1 & 2 & 4 \\ & 3 & \end{matrix}, \quad \begin{matrix} 1 & \textcircled{2} & 3 \\ & 4 & \end{matrix} \right)$$

and

$$\pi \to (R(\pi), S(\pi)) = \left(\begin{matrix} 1 & 2 \\ 3 & 4 \end{matrix}, \quad \begin{matrix} 1 & 3 \\ 2 & 4 \end{matrix} \right)$$

by the shifted and left-justified algorithms, respectively. Shore's method for obtaining $S(\pi)$ from $Q(\pi)$ is shown in Fig. 6.

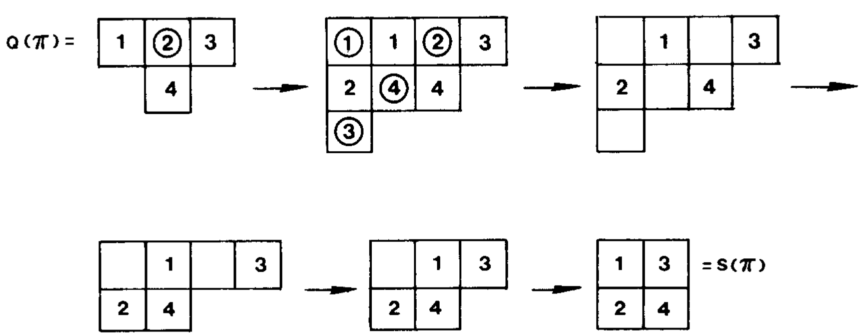

FIG. 6. Shore's transformation.

(3) Sometimes even results involving just the first tableau are elusive. In Schensted's original paper [Se] he proves that $R(\bar{\pi}) = R(\pi)^{t}$, where t denotes transposition. It is not at all clear what is meant by the transpose of a shifted tableau or how to apply Schensted's theorem to $P(\bar{\pi})$.

(4) Viewing π as a permutation matrix, we can let the dihedral group of the square act on π and ask: what happens to the corresponding tableaux? Problems (1) and (3) are the special cases of this approach obtained by considering two reflections. Gansner [G] and Schützenberger [Sü2] have answered this question for the ordinary algorithm. What can be said for P and Q?

(5) Lascoux and Schützenberger [LS] have been able to interpret various algorithms in this area of research by using their plactic monoid $\mathcal{M}$. $\mathcal{M}$ can be viewed as the quotient of the free monoid on $\mathbb{P}$ (or any ordered alphabet) by the Knuth relations (K1)–(K2). It is tempting to try the same thing in the shifted case, but we run into problems immediately because (SK3) is not context-free. It would be interesting to find the correct definition for a shifted plactic monoid and use it to obtain the results of this paper.

(6) White [W2, W3] and White and Stanton [SW] have generalized the Robinson–Schensted algorithm to rim-hook tableaux and hybrid tableaux (those which are partially column-strict and partially rim-hook). These correspondences can be applied to give a combinatorial proof of the orthogonality relations for the full character table of $\mathcal{S}_n$ by using the Murnagham–Nakayama formula [J1] to interpret the characters. We have recently developed a shifted rim-hook tableaux correspondence and hope it will provide, in conjunction with the projective Murnagham–Nakayama rule developed by Humphreys [H], a bijective proof of the projective character relations. These results will appear elsewhere.

(7) A number of authors [T1, W1, W3] have shown that there is an intimate connection between the Robinson–Schensted map and the celebrated Littlewood–Richardson rule for computing products of Schur functions. As yet, no analogous formula is known for finding the expansion of $Q_\lambda(\mathbf{x})\, Q_\mu(\mathbf{x})$. Perhaps the correspondence we have presented here can point the way. In particular, Worley [Wo] has developed a method for computing the coefficients of $Q_\lambda(\mathbf{x})\, S_\mu(\mathbf{x})$, where $S_\mu(\mathbf{x}) = \det(Q_{(\mu_i - i + j)}(\mathbf{x}))$. Worley has also used a shifted analog of Schützenberger's evacuation operator [Sü2] (a relative of the jeu de taquin) to obtain various other properties of symmetric function expansions conjectured by Stanley. Thus the problem of the preceding section is not the only one to succumb to this machinery.

(8) The Hall–Littlewood polynomials $Q_\lambda(\mathbf{x}; t)$ are symmetric functions

in the variables x_i, $i \in \mathbb{P}$, and a parameter t [M]. On setting $t=0$ or $t=-1$ one obtains the Schur functions $s_\lambda(\mathbf{x})$ or $Q_\lambda(\mathbf{x})$, respectively. Also, the $Q_\lambda(\mathbf{x}; t)$ satisfy the identity

$$\prod_{i,j=1}^{\infty} \frac{1-tx_i y_j}{1-x_i y_j} = \sum_\lambda \frac{1}{b_\lambda(t)} Q_\lambda(\mathbf{x}; t)\, Q_\lambda(\mathbf{y}; t)$$

(where $b_\lambda(t)$ is a polynomial in t) which generalizes both Cauchy's identity and Corollary 8.3. Perhaps it is possible to find a Knuth-type algorithm for proving this equation which will have the left-justified and shifted maps as special cases.

(9) Another approach to subsuming both correspondences would be to work in the covering group $\tilde{\mathscr{G}}_n$ of $\mathscr{G}_n$ [Su]. $\tilde{\mathscr{G}}_n$ is a central extension of $\mathscr{G}_n$ by the Schur multiplier, in this case of order 2, so that $|\tilde{\mathscr{G}}_n| = 2n!$. The ordinary representations of $\tilde{\mathscr{G}}_n$ correspond to both the ordinary and projective representations of $\mathscr{G}_n$. James [J2] has suggested that by finding a Robinson–Schensted map for $\tilde{\mathscr{G}}_n$ one might obtain the other two "for free."

(10) Robinson–Schensted correspondences have been developed for other groups, notably the classical Weyl groups [BV] and the symplectic group [B]. These involve various different types of arrays such as domino and oscillating tableaux. Can projective analogs be found?

It is hoped that this selection of problems will whet the reader's interest in the subject.

Note added in proof. Haimon (private communication) has recently proved Shore's conjecture in (2) above by using a correspondence that is in some sense dual to the one we have presented. The dual correspondence applied to π^{-1} yields the pair $(Q(\pi), P(\pi))$ thus providing another way of approaching the questions raised in (1).

References

[AHU] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, "Data Structures and Algorithms," Addison-Wesley, Reading, MA, 1983

[B] A. Berele, A Schensted-type correspondence for the symplectic group, preprint.

[BV] D. Barbasch and D. Vogan, Primitive ideals and orbital integrals in complex classical groups, *Math. Ann.* **259** (1982), 153–199.

[G] E. Gansner, "Matrix Correspondences and the Enumeration of Plane Partitions," Ph.D. thesis, MIT, 1978.

[Gr] C. Greene, An extension of Schensted's theorem, *Adv. in Math.* **14** (1974), 254–265.

[H] J. F. Humphreys, Blocks of projective representations of symmetric groups, preprint.

[J1] G. D. James, "The Representation Theory of the Symmetric Groups," Lecture Notes in Math. Vol. 682, Springer, Berlin, 1978.

[J2] G. D. James, private communication.

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

[LS] A. LASCOUX AND M. P. SCHÜTZENBERGER, Le monoide plaxique, *in* "Combinatorial Mathematics, Optimal Designs and Their Applications" (J. Srivastava, Ed.), pp. 251–255, Ann. Discrete Math. Vol. 6, North-Holland, Amsterdam/New York, 1980.

[LR] D. E. LITTLEWOOD AND A. R. RICHARDSON, Group characters and algebra, *Philos. Trans. Roy. Soc. London Ser. A* **233** (1934), 49–141.

[M] I. G. MACDONALD, "Symmetric Functions and Hall Polynomials," Oxford Univ. Press, Oxford, 1979.

[Ml] D. MACLARNAN, Ph.D. thesis, Univ. of Calif., San Diego, 1985.

[Mo1] A. O. MORRIS, The projective characters of the symmetric group—An alternative proof, *J. London Math. Soc.* **19** (1979), 57–58.

[Mo2] A. O. MORRIS, private communication.

[R] G. DE B. ROBINSON, On representations of the symmetric group, *Amer. J. Math.* **60** (1968), 745–760.

[S] B. E. SAGAN, An analog of Schensted's algorithm for shifted Young tableaux, *J. Combin. Theory Ser. A* **27** (1979), 10–18.

[Se] C. SCHENSTED, Longest increasing and decreasing subsequences, *Canad. J. Math.* **13** (1961), 179–191.

[Su] I. SCHUR, Über die Darstellung der symmetrischen und der alternienden Gruppe durch gebrochene lineare Substitutionen, *J. Reine Angew. Math.* **139** (1911), 155–250.

[Sü1] M.-P. SCHÜTZENBERGER, Quelques remarques sur une construction de Schensted, Math. Scand. **12** (1963), 117–128.

[Sü2] M.-P. SCHÜTZENBERGER, La correspondence de Robinson, *in* "Combinatoire et représentation du groupe symmétrique, Strasbourg 1976" (D. Foata, Ed.), Lecture Notes in Math. Vol. 579, pp. 59–113, Springer, Berlin, 1977.

[St] R. P. STANLEY, private communications.

[SW] D. STANTON AND D. WHITE, A Schensted correspondence for rim hook tableaux, *J. Combin. Theory Ser. A*, in press.

[T1] G. P. THOMAS, On a construction of Schützenberger, *Discrete Math.* **17** (1977), 107–118.

[T2] G. P. THOMAS, On Schensted's construction and the multiplication of Schur functions, *Adv. in Math.* **30** (1978), 8–32.

[W1] D. WHITE, Some connections between the Littlewood–Richardson rule and the construction of Schensted, *J. Combin. Theory Ser. A* **30** (1981), 237–247.

[W2] D. WHITE, A bijection proving orthogonality of the characters of S_n, *Adv. in Math.* **50** (1983), 160–186.

[W3] D. WHITE, Hybrid tableaux and the Littlewood–Richardson rule, preprint.

[Wo] D. WORLEY, "A Theory of Shifted Young Tableaux," Ph.D. thesis, MIT, 1984

582a/45/1-8