Notes on Knuth equivalence

Bill Casselman
University of British Columbia
cass@math.ubc.ca

The Robinson-Schensted correspondence is a bijection between permutations in $\mathfrak{S}_n$ and pairs of Young tableaux of the same shape and size n. What properties of a permutation can be read off from the corresponding tableaux?

It is largely expository in nature, but how I treat some aspects of Knuth equivalence in the last section is apparently somewhat new.

Contents

1. The Schensted process

The Schensted process starts with an input sequence σ of distinct positive integers and produces two tableaux P_σ and Q_σ. It does this by applying a basic insertion step to each item in the input sequence.

The process is most efficiently described in terms of recursion. If P is a tableau, let P_1 be its first row and $P_{>1}$ the tableau made up of subsequent rows. I write $P = P_1/P_{>1}$. If σ is any array, I write σ_1 for the first element and $\sigma_{>1}$ for the remainder, and write $\sigma = \sigma_1 \cdot \sigma_{>1}$. I write $[\ldots]$ for a single row of a tableau.

There is one basic routine **insert**(P, σ) in which P is a Young tableau and σ a sequence of distinct positive integers. It produces a new tableau.

insert(P, σ):
- If $\sigma = \emptyset$, return P.
- If $\sigma \neq \emptyset$ and $P = \emptyset$, set $P = [\sigma_1]$ and return **insert**$(P, \sigma_{>1})$.
- Suppose that σ has one element and that P is not empty. There are two cases: (a) The element σ_1 is larger than all the elements in P_1. Then set $P'_1 = P_1 \cdot \sigma_1$ and return $P'_1/P_{>1}$. (b) There exists an element of P_1 larger than σ_1, and in fact a first one (reading left to right). Suppose it is s. Let P'_1 be the result of replacing s by σ_1, and set $P'_{>1} =$ **insert**$(P_{>1}, s)$. Return $P'_1/P'_{>1}$.
- If σ has more than one element, then set $P =$ **insert**$(P, (\sigma_1))$ and return **insert**$(P, \sigma_{>1})$.

The original process then produces **insert**$(\emptyset, \sigma)$.

A slightly more complicated version produces a second tableau Q_σ of the same shape. This is normally constructed at the same time as P—when the m-th item is taken from input, the integer m is placed at the site on the second tableau corresponding to the newly occupied site on the first. There is also a remarkable well known duality:

1.1. Lemma. *If $\tau = \sigma^{-1}$, then the matrix Q_σ is the same as as the tableau P_τ.*

For example, with the input sequence $(4, 2, 1, 3)$ we get (with modified entries marked):

Step	P	Q	Remaining input
			4 2 1 3
1	$\begin{array}{c}4\end{array}$	$\begin{array}{c}1\end{array}$	2 1 3
2	$\begin{array}{c}2\\4\end{array}$	$\begin{array}{c}1\\2\end{array}$	1 3
3	$\begin{array}{c}1\\2\\4\end{array}$	$\begin{array}{c}1\\2\\3\end{array}$	3
4	$\begin{array}{cc}1&3\\2&\\4&\end{array}$	$\begin{array}{cc}1&4\\2&\\3&\end{array}$	

2. Descents and tableaux

Items in P_σ are moved around as the input is read (whereas those in Q_σ remain where they are placed initially). *How does the final placement in P_σ of an item in the input sequence σ relate to its initial one?*

2.1. Proposition. *The final position of an item in the tableau P_σ is either the same as its initial position, or lies below and weakly to its left.*

Proof. Suppose i bumps j at some position in P_σ. Consider the position just below that one. Either it is empty, or holds $k > j$. Either way, j will be inserted weakly to its left. This continues, row after row. (In effect, the proof is by induction on the number of rows.)

This is illustrated by the figure below.

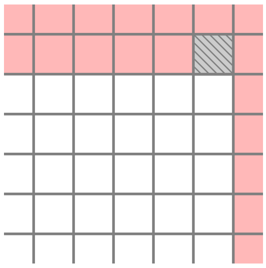

In the rest of this section I want to explain the possible relative positions of the numbers i, $i + 1$ in the Schensted tableau P_σ produced by a permutation σ.

I call attention first of all to some very general restrictions. In the figure below, it is shown that by the very definition of a strict tableau that there are certain geometrical conditions on entries.

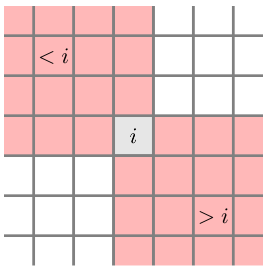

In considering how i and $i + 1$ are placed, there is the further restriction that there are no integers between i and $i + 1$. In the figure on the left below the number i has been placed, and the illegal positions for $i + 1$ are the shaded squares. In that on the right, it is $i + 1$ whose position is fixed, and the illegal positions for i are shaded.

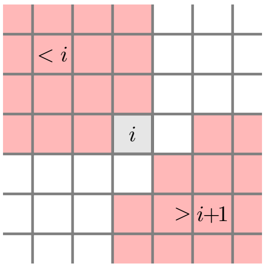

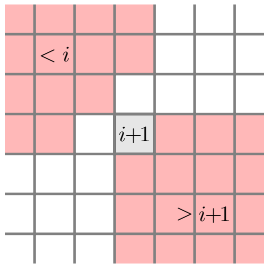

Speaking roughly, we can say that $i+1$ must lie either north-east or south-west of i. There is a very simple criterion for which of these occurs. Every tableau T defines a linear order on its entries—a variation of lexicographic. I say that $i \prec_T j$ if either (a) i lies in a row below that of j or (b) it lies in the same row but left of it. (This is a matter only of location, not value.)

2.2. Theorem. *Suppose σ to be in $\mathfrak{S}_n$. If $T = P_\sigma$ then $i \prec_T i+1$ if and only if i comes before $i+1$ in the array (σ_i).*

Proof. • Suppose that i comes before $i+1$ in the input sequence. It must be shown (i) that when $i+1$ is first inserted in a tableau that already contains i it will be placed either in a row above i, or in the same row but to its right; and (ii) that subsequent insertions will not change the fact that $i \prec_T i+1$.

Claim (i) is immediate, since the initial insertion of $i+1$ will place it in the first row of the tableau. As for (ii), subsequent insertions will never bounce $i+1$ from its position as long as i is just to its left, since there are no integers between i and $i+1$. Otherwise, $i+1$ might be bounced, but then it will be inserted in a lower tableau that already contains i, in which case we are back facing (i).

• Suppose that $i+1$ comes before i in the input sequence. It suffices to show that when i is inserted into a tableau containing $i+1$ it will be placed in a row strictly above $i+1$, and that subsequent insertions will not change this property.

Suppose i is inserted into a tableau containing $i+1$. Either it will be inserted into a row containing $i+1$, or it will be placed in a row above $i+1$. In the first case it will bounce $i+1$, which will wind up in a row below i.

Now consider any subsequent insertion into a tableau in which i is placed in a row above $i+1$. Such an insertion may bounce i, in which case it will be inserted into the row just below. We can invoke induction on the number of rows of the tableau. ∎

Let s_i be the swap of i, $i+1$. Recall that the length of a permutation is the number of the inversions in its array—the number of pairs (i,j) with $\sigma(i) < \sigma(j)$. Recall also:

> *Multiplying σ on the left by s_i swaps the symbols $i, i+1$ in the array of σ, whereas multiplying on the right swaps the entries in the locations $i, i+1$.*

Thus $\ell(s_i\sigma) = \ell(\sigma)+1$ if i appears before $i+1$ in the array of σ, and $\ell(s_i\sigma) = \ell(\sigma)-1$ if it appears after it.

2.3. Corollary. *If $T = P_\sigma$, then $\ell(s_i\sigma) > \ell(\sigma)$ if and only if $i \prec_T i+1$.*

The s_i such that $\ell(s_i\sigma) < \ell(\sigma)$ make up the **left descent set** of σ. (This is a well defined notion in any Coxeter group.) In other words, we can read off the left descent set of σ from the tableau P_σ. We can read off the right descent set of σ, which is the same as the left descent set of σ^{-1}, from Q_σ.

Remark. These results suggest a natural question. Since there are no integers between i and $i+1$, swapping them in a tableau will produce a tableau unless i and $i+1$ are next to each other in the same row or column. Suppose we do this in Q_σ but leave P_σ fixed. *What is the permutation with this new tableau pair?* Experiments lead me to think there is no simple answer to this question in general, but later on we shall see an important case in which there is one. Sometimes, but not always, it is σs_i, which amounts to an interchange of i and $i+1$ in the array of σ. The problem is that there are only special circumstances in which the tableaux of σs_i have the same shape as those of σ. *What are those circumstances?*

3. Knuth equivalence

Under what circumstances do two sequences (x_i) and (y_i) give rise to the same P?

Some simple examples are instructive. Two possible tableau pairs with equal P are:

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

The first comes from insertion of 132, the second from 312. Similarly, 213 and 231 have the same left tableau.

These are the basic cases, in a precise sense. Suppose that $a < b < c$ are three integers. The reasoning above shows that if $x = acb$ and $y = cab$ then $P_x = P_y$. Similarly for bca and bac. Following Knuth, I define an equivalence to be that generated by the relations

$$\begin{aligned} \boldsymbol{ac}b &\equiv_\kappa \boldsymbol{ca}b \\ b\boldsymbol{ac} &\equiv_\kappa b\boldsymbol{ca}\,. \end{aligned}$$

In less formal terms, whenever an item to left or right of a neighbouring pair fits in between them, we may swap the elements of this pair to obtain equivalent sequences. This can be extended to an equivalence relation on arbitrary arrays: I'll call x and y **Knuth-equivalent** and write $x \equiv_\kappa y$ if x and y can be connected by a chain of such elementary exchanges. For example, we have a sequence of exchanges

$$\begin{gathered}
[6, 3, \boldsymbol{1}, \boldsymbol{7}, \dot{2}, 5, 8, 4] \\
[6, 3, 7, 1, 2, \dot{5}, \boldsymbol{8}, \boldsymbol{4}] \\
[6, 3, 7, 1, \boldsymbol{2}, \boldsymbol{5}, \dot{4}, 8] \\
[6, 3, 7, \boldsymbol{1}, \boldsymbol{5}, \dot{2}, 4, 8] \\
[6, \boldsymbol{3}, \boldsymbol{7}, \dot{5}, 1, 2, 4, 8] \\
[6, 7, 3, 5, 1, 2, 4, 8]
\end{gathered}$$

An easy computation will show that both $(6, 3, 1, 7, 2, 5, 8, 4)$ and $(6, 7, 3, 5, 1, 2, 4, 8)$ are associated by the Schensted process to the tableau

$$\begin{array}{|c|c|c|c|}\hline 1 & 2 & 4 & 8 \\ \hline 3 & 5 \\ \cline{1-2} 6 & 7 \\ \cline{1-2}\end{array}\ . \tag{3.1}$$

The sequence of transformations above therefore illustrates one implication of the following basic result, originally from §6 of [Knuth:1970].

3.2. Theorem. *We have $x \equiv_\kappa y$ if and only if $P_x = P_y$.*

The proof will be long, and in several steps.

Step 1. To every P corresponds its **canonical sequence** p_P. It is the sequence obtained by scanning the rows of P bottom to top, left to right. For example, to the tableau (3.1) corresponds the sequence $[6, 7, 3, 5, 1, 2, 4, 8]$.

3.3. Lemma. *If p is the canonical sequence of the tableau P, then $P_p = P$.*

Proof. By induction on the number of rows in the tableau. The Lemma is trivial for a tableau of one row.

Let s_i be the sequence obtained from the i-th row of T by reading it left to right. If T has n rows, the canonical sequence of T is then the concatenation $s_n s_{n-1} \ldots s_1$. By induction, we may assume the Lemma to be true for the tail of T, whose canonical sequence is $s_n \ldots s_2$. Applying recursion, it suffices to show that inserting s_1 into the tail of T changes its first row to the first row of T, with extrusion s_2. This is straightforward.

Step 2. Next I'll prove the easier half of the Theorem:

(a) *If $P_x = P_y$ then $x \equiv_\kappa y$.*

Proof. This will involve a new interpretation of Schensted's insertion algorithm, which itself amounts to a basic connection between Knuth equivalence and Schensted's algorithm.

To start the proof, let $P = P_x$. It suffices to prove the claim in the special case that y is the canonical sequence p_P or, in other words, to prove that x can be connected by a chain of elementary exchanges to p_P.

The basic point is this:

Schensted's algorithm amounts to a succession of elementary exchanges.

If x is any array, let $\mathrm{NF}(x)$ (for **n**ormal **f**orm of x) be the canonical sequence of P_x. There are three important facts relating the normal form to Schensted row insertion.

3.4. Lemma. *Suppose that insertion of the array x into an empty row produces the row r and extrudes the sequence z.*

(a) The normal form $\mathrm{NF}(r)$ is just r;
(b) the normal form $\mathrm{NF}(x)$ is equal to $\mathrm{NF}(z) \cdot r$;
(c) the array x is Knuth-equivalent to $z \cdot r$.

Here $x \cdot y$ is the concatenation of the arrays x and y.

Proof. The first two of these is an immediate consequence of definitions.

The last amounts to the new interpretation of row insertion. Suppose we are inserting a single item c into a row

$$x_1 \dots x_m \, .$$

There are two possibilities for how insertion proceeds. One is that $x_m < c$, in which case we tack c onto the end, and bounce nothing. In the other, we find i such that $x_{i-1} < c < x_i$ and the sequence $x = (x_j)$ changes to

$$x_1 \dots x_{i-1} c x_{i+1} \dots x_m \, ,$$

and we bounce x_i. In this second case, I claim that the array

$$x_i \cdot x_1 \dots c \dots x_m$$

is Knuth-equivalent to $x \cdot c$.

To verify this, first of all set $j = m$, and as long as $c < x_{j-1} < x_j$ we can change $x_{j-1} x_j c$ to $x_{j-1} c x_j$, and decrement j. (Here and elsewhere I adopt the harmless convention that $x_{-1} = 0$.)

At the end of this phase we are looking at

$$x_{i-1} x_i c$$

with $x_{i-1} < c < x_i$. We leave c fixed in place (it has in effect just bounced x_i) and now proceed similarly to shift x_i all the way to the left. This proves the claim, and by induction on the length of z proves (b) above.

But if we combine (a) and (b) with an induction hypothesis on the length of x for the Proposition, we get

$$x \equiv_\kappa z \cdot r \equiv_\kappa \mathrm{NF}(z) \cdot r = \mathrm{NF}(x) \, .$$

For example, if we insert 4 into $(1, 2, 5, 8, 9)$ the 4 will bounce the 5. First we insert 4:

$$1\,2\,5\,8\,9 \cdot \boldsymbol{4}$$
$$1\,2\,5\,8\,\boldsymbol{9}\,\boldsymbol{4}$$
$$1\,2\,5\,\boldsymbol{8}\,\boldsymbol{4}\,9$$
$$1\,2\,5\,\boldsymbol{4}\,8\,9$$

and then bounce 5:

$$1\,\mathbf{2}\,\mathbf{5}\,4\,8\,9$$
$$\mathbf{1}\,\mathbf{5}\,2\,4\,8\,9$$
$$\mathbf{5}\cdot 1\,2\,4\,8\,9\,.$$

Step 3. Now the converse.

(b) *If x is a sequence (p_i). and $y \equiv_\kappa x$, then $P_y = P_x$.*

Proof. The proof I'll give here is straightforward if unilluminating. In the next section I'll give an alternate proof that I like better, but there is some virtue in the direct route outlined here, which amounts to the one hinted at in the answer to exercise 5.1.4.5 of [Knuth:1975]. A rather different proof of the Lemma can be found in Chapter 3 of [Fulton:1997].

It must be shown that if x and y are related by an elementary interchange then $P_x = P_y$. Applying an induction argument, it suffices to show that inserting x and y into a single row produce (a) the same new row as well as (b) extrusions that are Knuth-equivalent. It even suffices to assume that x and y are one of the Knuth triples acb etc.. The proof goes according to cases. There are several of these, and laying them all out is somewhat tedious, if automatic.

I shall track the insertions of the triples acb, cab, bac, bca into the row $x_1 \ldots x_n$. I shall allow $n = 0$ by following the harmless convention that $x_0 = 0$. I shall also use a trick suggested somewhere by Knuth—I introduce several very, very large integers and assume an arbitrary number of them at the end of every row. If I then perform the Schensted process, all of these will get placed again at the ends of rows. They can then finally be removed, leaving exactly the same tableau we would have had without using them. The point of this is to shrink the number of cases, since now every insertion will bounce something.

Define three integers i, j, k by the conditions

$$x_{i-1} < a < x_i, \quad x_{j-1} < b < x_j, \quad x_{k-1} < c < x_k\,.$$

The different cases we have to consider are distinguished by how i, j, k relate to each other. Since $a < b < c$ we know at least that $i \le j \le k$.

I shall now look at inputs acb, cab.

CASE $\boldsymbol{k = i}$

Here $x_{i-1} < a < b < c < x_i < x_{i+1}$.

Extrusion	Row	Input
	$x_1 \ldots x_n$	acb
x_i	$\ldots a x_{i+1} \ldots$	cb
$x_i x_{i+1}$	$\ldots ac \ldots$	b
$x_i x_{i+1} c$	$\ldots ab \ldots$	
	$x_1 \ldots x_n$	cab
x_i	$\ldots c x_{i+1}$	ab
$x_i c$	$\ldots a x_{i+1} \ldots$	b
$x_i c x_{i+1}$	$\ldots ab \ldots$	

Since $c < x_i < x_{i+1}$ the extruded triples are equivalent.

CASE $\boldsymbol{k = i + 1}$

Here $x_{i-1} < a < x_i < c < x_{i+1}$.

	$x_1 \ldots x_n$	acb
x_i	$\ldots a x_{i+1} \ldots$	cb
$x_i x_{i+1}$	$\ldots ac \ldots$	b
$x_i x_{i+1} c$	$\ldots ab \ldots$	

	$x_1 \ldots x_n$	cab
x_{i+1}	$\ldots x_i c$	ab
$x_{i+1}x_i$	$\ldots ac \ldots$	b
$x_{i+1}x_i c$	$\ldots ab \ldots$	

Again since $x_i < c < x_{i+1}$ the extruded triples are equivalent.

CASE $\boldsymbol{k \geq i+2}$
Here $x_{i-1} < a < x_i < x_{i+1} \leq x_{k-1} < c < x_k$.

	$x_1 \ldots x_n$	acb
x_i	$\ldots a x_{i+1} \ldots$	cb
$x_i x_k$	$\ldots a x_{i+1} \leq x_{k-1} c \ldots$	b

	$x_1 \ldots x_n$	cab
x_k	$\ldots x_i \ldots c x_{k+1}$	ab
$x_k x_i$	$\ldots a x_{i+1} \leq x_{k-1} c \ldots$	b

But now b might bounce anything, call it y, from x_{i+1} through c. Since in all these cases $x_i < y < x_k$, the extrusions are again equivalent.

We are through with inputs acb, cab. I'll leave the pair bac, bca as an exercise.

I am now going to rephrase Theorem 3.2. Let $\mathfrak{X}_n$ be the set of all injective maps from $[1, n]$ to the positive integers. The group $\mathfrak{S}_n$ acts on it on the right—the action is determined by the condition that s_i swaps the items in positions i and $i+1$. The space $\mathfrak{X}_n$ becomes a principal homogeneous space for this action.

Let $\mathfrak{S}_{i,i+1}$ be the subgroup of $\mathfrak{S}_n$ generated by s_i and s_{i+1}. It is isomorphic to $\mathfrak{S}_3$. Let $[\mathfrak{X}_n/\mathfrak{S}_{i,i+1}]$ be the set of arrays x with $x_i < x_{i+1} < x_{i+2}$.

3.5. Lemma. *Suppose $i \le n-2$. Every element x in $\mathfrak{X}_n$ can be expressed uniquely as $x = y \cdot s$, with y in $[\mathfrak{X}_n/\mathfrak{S}_{i,i+1}]$ and s in $\mathfrak{S}_{i,i+1}$.*

Proof. Since multiplication on the right by elements of $\mathfrak{S}_{i,i+1}$ permutes (x_i, x_{i+1}, x_{i+2}).

There is hence a dictionary between subsequences of neighbouring a, b, c with $a < b < c$ and elements of $\mathfrak{S}_{i,i+1}$.

(3.6)

pattern in $x_{i,i+1,i+2}$	factor s in $\mathfrak{S}_{i,i+1}$
abc	1
acb	s_{i+1}
cab	$s_i s_{i+1}$
bac	s_i
bca	$s_{i+1} s_i$
cba	$s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$.

The point for us now is that Knuth exchanges in positions i through $i+2$ can be interpreted in terms of these factorizations. I'll say that an array x is **eligible** for a Knuth exchange in positions $[i, i+2]$ if and only if the factor s in $\mathfrak{S}_{i,i+1}$ lies in $\{s_i, s_{i+1}, s_i s_{i+1}, s_{i+1} s_i\}$. Or, conversely, if it is not either 1 or $s_i s_{i+1} s_i$.

This can also be interpreted in terms of the following graph of Bruhat order of $\mathfrak{S}_{i,i+1}$, in which Knuth-equivalent permutations are linked by doubled edges. I call them **twins**. The identity and the involution at the top do not have twins.

(3.7)

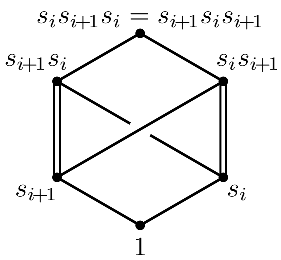

.

Let $\mathcal{D}_R(i, i+1)$ be the subset of σ in $\mathfrak{S}_n$ whose right factor in $\mathfrak{S}_{i,i+1}$ is eligible. For each eligible array x, let $\mathrm{TW}(x)$ be its twin. Ther map $x \mapsto \mathrm{TW}(x)$ is an involution of $\mathcal{D}_R(i, i+1)$.

The eligible arrays in $\mathfrak{X}$ can also be characterized by the requirement that

$$\{s_i, s_{i+1}\} \cap R_x$$

be a singleton.

3.8. Proposition. *Let $x = yw$ with y in $[\mathfrak{X}_n/\mathfrak{S}_{i,i+1}]$ and s in $\mathfrak{S}_{i,i+1}$. A Knuth exchange at sites i to $i+2$ replaces s by its twin, if it exists.*

To summarize, there is a right Knuth exchange $x \mapsto x^{|i,i+1}$ defined on a certain subset $\mathcal{D}_R(i, i+1)$ of W for every $1 \le i \le n-2$. Suppose i to be in this range, and suppose x to be a permutation with $x_i < x_{i+1} < x_{i+2}$, Then the Knuth exchange swaps

$$\begin{aligned} xs_i &\longleftrightarrow xs_is_{i+1} \\ xs_{i+1} &\longleftrightarrow xs_{i+1}s_i\,. \end{aligned}$$

One can also define a left Knuth exchange $x \mapsto {}^{i,i+1|}x$ with domain $\mathcal{D}_L(i, i+1)$, those x for which

$$\{s_i, s_{i+1}\} \cap L_x$$

is a singleton.

4. More about Knuth equivalence

In this section I'll analyze Knuth equivalence more closely, offering among other things a second proof of Theorem 3.2.

The material in this section seems to be well known, but details appeared only recently. There is an interesting discussion of the history of this topic in `mathoverflow` at

`http://mathoverflow.net/questions/139432/`
`has-reifegerstes-theorem-on-rsk-and-knuth-relations-received-a-slick-proof-by-n`

There is also a separate thread in the literature more directly related to Kazhdan-Lusztig cells. The earliest reference I am aware of is [Barbasch-Vogan:1982] (p. 172), although their account is a bit vague about precise effects. Later accounts include [Ariki:2000] and [Du:2005], but the relevant discussions there are rather sketchy.

Theorem 3.2 asserts that x and y are permutations and if $y = x^{|i,i+1}$ then $P_x = P_y$ (as I'll reprove in a little while). Theorem 2.2 applied to x^{-1}, in conjunction with the well known symmetry, tells us that we can determine from Q_x alone whether x is in $\mathcal{D}_R(i, i+1)$. I'll summarize here what we know.

As I have mentioned earlier, each tableau T determines an order on its entries, in terms of their locations. We have $i \prec_T j$ if either (1) i lies in a row below j or (2) i lies in the same row as j but to its left. Thus Theorem 2.2 says that if $T = P_x$ then i comes before $i+1$ in the array (x_j) if and only if $i \prec_T i+1$. If x is replaced by x^{-1} then P is replaced by Q. Translating this accordingly:

4.1. Proposition. *Suppose $1 \le i \le n-2$, x in $\mathfrak{S}_n$, $T = Q_x$. Then x lies in $\mathcal{D}_R(i, i+1)$ if and only if one of the following is valid:*

1(a) $i+1 \prec_T i \prec_T i+2$;
1(b) $i+2 \prec_T i \prec_T i+1$;
2(a) $i \prec_T i+2 \prec_T i+1$;
2(b) $i+1 \prec_T i+2 \prec_T i$.

These are just the different cases required by Theorem 2.2. For example, $xs_i < x$, $xs_{i+1} > x$ requires

$$i+1 \prec_T i,\ i+2\,.$$

Since $\prec_T$ is a linear order, we must then have either $i \prec_T i+2$ or $i+2 \prec_T i$. This gives cases 1(a) and 2(b). These cases are illustrated in the following diagrams.

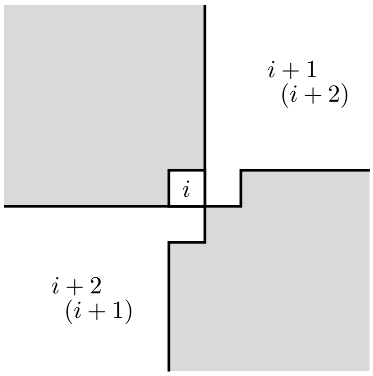

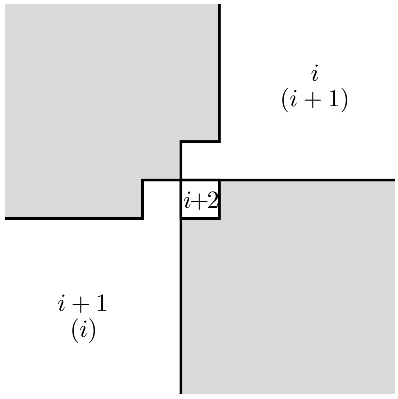

The next natural question is, *how is Q_y related to Q_x?* It should not surprise us that Q_x and Q_y differ only in the locations of some of the indices i, $i+1$, $i+2$.

4.2. Theorem. *Suppose $1 \le i \le n-2$, x in $\mathfrak{S}_n$, and suppose x is in $\mathcal{D}_R(i, i+1)$. In these circumstances, let $y = x^{|i,i+1}$. Then $P_x = P_y$, and the tableau Q_y is derived from Q_x by swapping the two extreme items in the relevant list of Proposition 4.1.*

Thus 1(a) and 1(b) are swapped, as are 2(a) and 2(b).

Before the proof, I'll look at some examples. Suppose the input sequence is $x = [4, 1, 6, 2, 5, 3]$. In the following list, b is $\dot{b}$, a and c in bold face.

$[4, 1, 6, 2, 5, 3]$	$\begin{array}{lll} 1 & 3 & 5 \\ 2 & 4 \\ 6 \end{array}$
$[\dot{4}, \boldsymbol{6}, \boldsymbol{1}, 2, 5, 3]$	$\begin{array}{lll} 1 & 2 & 5 \\ 3 & 4 \\ 6 \end{array}$
$[4, \boldsymbol{6}, \boldsymbol{1}, \dot{2}, 5, 3]$	$\begin{array}{lll} 1 & 2 & 5 \\ 3 & 4 \\ 6 \end{array}$
$[4, 1, \boldsymbol{2}, \boldsymbol{6}, \dot{5}, 3]$	$\begin{array}{lll} 1 & 3 & 4 \\ 2 & 5 \\ 6 \end{array}$
$[4, 1, 6, \boldsymbol{5}, \boldsymbol{2}, \dot{3}]$	$\begin{array}{lll} 1 & 3 & 6 \\ 2 & 4 \\ 5 \end{array}$

There is something slightly subtle about this result—the swap in Q_x is not necessarily the same as the swap in x. If x contains bac at positions $i, i+1, i+2$ with $a < b < c$, and y replaces this by bca, then in all cases Q_y is obtained from Q_x by swapping $i+1$ and $i+2$. But if x contains acb at the same positions and y replaces this by cab, then it can happen that either i and $i+1$ or $i+1$ and $i+2$ are swapped. These phenomena can be seen in the figures above.

What is going on reflects a fundamental difference between configurations of type (1), in which i is in the middle, and those of type (2), in which $i+2$ is in the middle. The first are stable, in the sense that successive insertions to thos eof $i, i+1, i+2$ will not change the type of the configuration. But the second can change (necessarily permanently) to a type (1) configuration. This happens when either i is bounced and reinserted into a row in which $i+1, i+2$ occur, in which case $i+1$ is bounced, or $i+1$ is bounced and reinserted into a row in which $i, i+2$ occur, in which case $i+2$ is bounced.

Proof of Theorem 4.2. Notation will be simpler if I interpret the result as an assertion about the inverse of x. Lemma 1.1 tells us that if $y = x^{-1}$ then $P_y = Q_x$ and $Q_y = P_x$. This allows us an easy translation. For example, if the Knuth triple acb occurs in positions i, $i+1$, $i+2$ of x then

$$x^{-1}\colon \begin{cases} a \mapsto i \\ c \mapsto i+1 \\ b \mapsto i+2\,. \end{cases}$$

In the Schensted process for x^{-1} the indices a, b, c are met in that order, so in x we first encounter i, then $i+2$, and then $i+1$.

Replacing x by x^{-1}, we now have two things to prove.

(1) Suppose x to contain in order $i+1$, i, $i+2$ while y differs from x only in swapping $i+2$ and $i+1$. We wish to show that $Q_x = Q_y$ and that P_y is obtained from P_x by swapping $i+1$ and $i+2$.

(2) Suppose x to contain in order i, $i+2$, $i+1$ while y differs from x only in swapping i and $i+1$. The first thing we wish to show is that $Q_x = Q_y$. As for P_x and P_y, I apply Lemma 1.1, applied to x. Depending on whether (1) or case (2) of of that Proposition occurs, we wish to show that either $i+1$ and $i+2$ or i and $i+1$ are swapped.

The basic idea in both cases is the same. Let $x_{\le m}$ be the sequence of x_i for $i \le m$, and similarly for $y_{\le m}$. Also, let $P_{x,n}$ be the tableau corresponding to $x_{\le n}$, and similarly $P_{y,n}$, $Q_{x,n}$, $Q_{y,n}$. By convention, any of these is an empty tableau for $n = 0$.

First, case (1). Suppose $i+1$, i, and $i+2$ to occur in x at positions k, ℓ, m, so y holds $i+2$, i, and $i+1$ at those same positions. We read in x and y item by item. I claim that as we do this $Q_{x,n}$ is always equal to $Q_{y,n}$, and that $P_{x,n}$ differs from $P_{y,n}$ only in that where $P_{x,n}$ holds $i+1$ (resp. $i+2$) and $P_{y,n}$ holds $i+2$ (resp. $i+1$).

These claims are certainly true for $n < k$. What happens for $n = k$? For x we insert $i+1$ into the top row of $P_{x,k-1}$ and for y we insert $i+2$ into the top row of $P_{y,k-1}$. But these top rows are the same, say r, and $r_j < i+1 < r_{j+1}$ if and only if $r_j < i+2 < r_{j+1}$ since r does not intersect $[i, i+2]$. Therefore $i+1$ is inserted in the top row of $P_{x,k-1}$ at the same location as $i+2$ is inserted in that of $P_{y,k-1}$, and the same item is bounced into the common lower rows of both. Thus $P_{x,k}$ differs from $P_{y,k}$ only in that $i+2$ is located in $P_{y,k}$ where $i+1$ is located in $P_{x,k}$ and $Q_{x,k} = Q_{y,k}$. Since $x_n = y_n$ for n in $[k+1, \ell-1]$, this remains true up through $n = \ell - 1$.

This illustrates the basic principle: inserting j into a tableau that does not contain $j+1$ has the same effect as inserting $j+1$ into one that does not contain j.

What happens at $n = \ell$? Well, i will bounce $i+1$ from $P_{x,\ell-1}$ if and only if it bounces $i+2$ from $P_{y,\ell-1}$, and it will bounce it to the same location. So our claim remains valid for $n = \ell$. In effect, $i+1$ and $i+2$ behave exactly the same as input, and Theorem 2.2 may be applied to both. This guarantees that our claim remains valid for $n < m$.

What happens for $n = m$? Well, i is located NE of $i+1$ in $P_{x,m-1}$ and it is located NE of $i+2$ in $P_{y,m-1}$, so inserting $i+1$ in $P_{y,m-1}$ has exactly the same effect as inserting $i+2$ into $P_{x,m-1}$, and the claim remains valid for $n = m$. The remaining input does not affect this.

Case (2) is essentially the same, except that in the final insertion of $n = m$ something new can happen, if the first row of $P_{x,m-1}$ contains i, $i+2$. In this case, $i+1$ bumps $i+2$, which is then inserted in a lower tableau, leaving i, $i+1$ in the first row. What happens for y? Since the claim is valid for $n = m-1$, the first row of y contains $i+1$, $i+2$, and i bumps $i+1$, leaving i, $i+2$ in the first row. We are now back in case (1). ▯

There is an important consequence—one can tell just from the tableau Q_x what right Knuth transforms are possible, and how to determine Q_y if y is such a transform. Equivalently, one can tell from P_x what left Knuth transforms are possible, and how to effect them.

If c is a chain of pairs (s_i, s_{i+1}), we can define the domain $\mathcal{D}_L(c)$ as well as the operator $x \mapsto {}^{c|}x$ from $\mathcal{D}_L(c)$ to W by induction—if c is the juxtaposition of $(i, i+1)$ and d then $\mathcal{D}_L(c)$ is the subset of $\mathcal{D}_L(i, i+1)$ such that ${}^{i,i+1|}x$ lies in $\mathcal{D}_L(d)$, and then ${}^{c|}x = {}^{i,i+1|d|}x$.

Define $\equiv_r$ to mean right Knuth-equivalence.

4.3. Corollary. *If x and y are two permutations such that $x \equiv_r y$, then x is in $\mathcal{D}_L(c)$ if and only if y is, and then ${}^{c|}x \equiv_r {}^{c|}y$.*

This is relevant in understanding that for $\mathfrak{S}_n$ the Knuth equivalence classes coincide with the cells defined by [Kazhdan-Lusztig:1979].

5. References

1. Susumu Ariki, 'Robinson-Schensted algorithm and left cells', *Advanced Studies in Pure Mathematics* **28** (2000), 1–20.

2. Dan Barbasch and David Vogan, 'Primitive ideals and orbital integrals in complex classical groups', *Mathematische Annalen* **259** (1982), 153–199.

3. Jie Du, 'Robinson-Schensted algorithm and Vogan equivalence', *Journal of Combinatorial Theory* **112** (2005), 165–172.

4. William Fulton, **Young diagrams**, Cambridge University Press, 1997.

5. A. M. Garsia and T. J. McLarnan, 'Relations between Young's natural and the Kazhdan-Lusztig representations of S_n', *Advances in Mathematics* **69** (1988), 32–92.

6. David Kazhdan and George Lusztig, 'Representations of Coxeter groups and Hecke algebras', *Inventiones Mathematicae* **53** (1979), 165–184.

7. Donald Ervin Knuth, 'Permutations, matrices, and generalized Young tableaux', *Pacific Journal of Mathematics* **34** (1970), 709–727. Also to be found in the collection **Selected papers on discrete mathematics**, CSLI, Stanford, 2003.

8. ——, **Fundamental algorithms**, volume I of **The Art of Computer Programming**, Addison-Wesley, 1973.

9. ——, **Sorting and Searching**, volume III of **The Art of Computer Programming**, Addison-Wesley, 1975.

10. Jacob Post, 'Combinatorics of arc diagrams, Ferrers fillings, Young tableaux and lattice paths', Ph. D. Thesis, Simon Fraser University, 2009.

11. Astrid Reifegerste, 'Permutation sign under the Robinson-Schensted correspondence', *Annals of Combinatorics* **8** (2004), 103–112.

12. C. Schensted, 'Longest increasing and decreasing subsequences', *Canadian Journal of Mathematics* **12** (1963), 117–128.

13. Robert Steinberg, 'An occurrence of the Robinson-Schensted correspondence', *Journal of Algebra* **113** (1988), 523–528.