EXPLAINING THE SHAPE OF RSK

SIMON RUBINSTEIN-SALZEDO

1. Introduction

There is an algorithm, due to Robinson, Schensted, and Knuth (henceforth RSK), that gives a bijection between permutations $\sigma \in S_n$ and pairs (S, T) of standard Young tableaux of the same shape λ, where λ is a partition of n.

Let us illustrate how the RSK algorithm is performed. Let $\sigma \in S_8$ be the permutation given in cycle decomposition as $(186)(457)$. We first write the permutation in a different form $a_1 \cdots a_8$, where $a_i = \sigma(i)$. So, σ written in this manner becomes 82357146.

We obtain (S, T) incrementally from σ. First, we put the 8 in a box, the beginning of the first tableau S. Hence, after one step, S looks like

$$\begin{array}{|c|}\hline 8 \\ \hline \end{array}.$$

For future numbers, we try to fit them at the end of the first row, but sometimes this move is not allowable. When this occurs, there is a unique spot on the first row where the next number is allowable. If there is already a number there (i.e., if the new number is less than at least one of the numbers currently residing in the top row), then the number in that spot is bumped to the second row, and the same procedure is repeated until some bumped number finds a home at the right end of a row.

In this example, then, the 2 bumps the 8 down, so after step two S looks like

$$\begin{array}{|c|}\hline 2 \\ \hline 8 \\ \hline \end{array}.$$

The next number to place is the 3, which is allowed to follow the 2. Similarly, the 5 and 7 cause us no difficulties, so after step five, S looks like

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

Date: October 18, 2011.

Now, the 1 bumps the 2 down to the second row, and the 2 bumps the 8 down to the third row, so after step six, S looks like

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

On step seven, the 4 bumps the 5 down to the second row, where it can find a comfortable home, so then S is

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

Finally, on the last step, the 6 bumps the 7 down to row two, giving us the final version of S, which is

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

So, that's the story for S. To construct T, we number the evolution of the shape of S: at each stage, exactly one new box is added to S. So, the box of S constructed on the r^{th} step is labeled r. Retracing the above process then tells us that

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

There are many possible questions that one can ask about the RSK algorithm. For example, for some $\lambda \vdash n$, for which σ is the shape of S (or T) equal to λ? Or, for a given standard Young tableau X, when is $S(\sigma)$ equal to X?

More generally, the tableaux S and T are some encrypted form of σ. How do we decode them in order to extract useful information?

All these questions have elegant answers that we will discuss in this essay.

Notation. Let $\lambda \vdash n$ be a partition. The i^{th} part of λ will be denoted by λ_i. We'll write λ^* for its conjugate partition of n. That is, λ_i^* is the number of parts of λ with length at least i.

2. Monotone subsequences

We'll begin with a warm-up before tackling Greene's full theorem about shape of the tableaux from [1]. The following theorem is due to Schensted:

Theorem 1. *(Schensted, [3]) The length of the longest increasing subsequence of σ is equal to λ_1 (the number of columns of S). Similarly, the length of the longest decreasing subsequence of σ is equal to λ_1^* (the number of rows of S).*

Example. Let σ be the permutation 82357146. A longest increasing subsequence is 2357, and a longest decreasing subsequence is 821. (Neither of these is unique.) Hence, the longest increasing subsequence has length 4, and the longest decreasing subsequence has length 3, as predicted by Schensted's Theorem.

Before starting the proof, we introduce the notion of basic subsequences. For a positive integer j, the j^{th} basic subsequence is a chronology of the numbers that, at some point in the construction of S under the RSK algorithm, occupy the j^{th} column in the first row. Note that each basic subsequence is decreasing.

Example. Using 82357146 again, we find that the first basic subsequence is 821, the second is 3, the third is 54, and the fourth is 76.

Proof. Since each basic subsequence is decreasing, an increasing sequence can only contain at most one member of each basic subsequence. Thus, the length of the longest increasing subsequence is less than or equal to λ_1. Now, for each member r of the j^{th} basic subsequence, we can find some member of the $(j-1)^{\text{th}}$ basic subsequence which is less than r; one way of doing this is to pick the number that occupied the $(j-1)^{\text{th}}$ spot on the first row when r was first inserted. Hence, if we pick an element in the $(\lambda_1)^{\text{th}}$ basic subsequence and work back from there, we obtain an increasing subsequence of length λ_1.

To prove the part about descending sequences, we claim that, if we reverse a permutation σ to obtain τ (so, in the example $\sigma = 82357146$ we've been using all along, we obtain $\tau = 64175328$), then $S(\tau) = S(\sigma)^*$ is the conjugate tableau. We omit the proof because it consists only of a rather tedious computation. Now, a decreasing sequence of σ is the same as an increasing sequence of τ, so the result follows. ■

3. Knuth's equivalence relation

Another question we can ask, that will also be useful for future analysis, is when two permutations $\sigma, \tau \in S_n$ give rise to the same first tableau S. Knuth provided an answer to this question.

Definition 2. Suppose $\sigma \in S_n$ is a permutation, and there exist three numbers $1 \leq x < y < z \leq n$ so that σ contains a consecutive subsequence of x, y, and z (not necessarily in that order), with x and z adjacent to each other. Let $\tau \in S_n$ be the permutation obtained from σ by interchanging x and z. Then we say that σ and τ are an elementary Knuth equivalent pair. More generally, let $\sim$ be the equivalence relation on S_n generated by the Knuth equivalent pairs described above. Then if $\sigma \sim \tau$, we say that σ and τ are Knuth-equivalent.

Remark 3. Knuth equivalence can also be defined in terms of the Bruhat-Chevalley ordering, if we give the symmetric group the structure of a Coxeter group. It is likely that there is interesting combinatorics to be found in relating the RSK algorithm to this ordering and perhaps algebraic groups in general. Unfortunately, I do not know what it is.

Theorem 4. *[2] Two permutations $\sigma, \tau \in S_n$ have the same S if and only if they are Knuth-equivalent.*

Example. For our favorite permutation $\sigma = 82357146$, we know that

$$S = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 5 & 7 \\ \cline{1-3} 8 \\ \cline{1-1} \end{array}, \qquad T = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 5 \\ \hline 2 & 7 & 8 \\ \cline{1-3} 6 \\ \cline{1-1} \end{array}.$$

For $\tau = 82351746$ (we've just switched the fifth and sixth digits, which is an allowable operation under Knuth equivalence), we have

$$S = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 5 & 7 \\ \cline{1-3} 8 \\ \cline{1-1} \end{array}, \qquad T = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 7 & 8 \\ \cline{1-3} 5 \\ \cline{1-1} \end{array}.$$

While it is true that we have switched the fifth and sixth digits to get from σ to τ, and the numbers 5 and 6 have been interchanged in $T(\sigma)$ and $T(\tau)$, we ought not read too much into this behavior: I do not know of any simple explanation in the variance of T as we replace σ with a Knuth-equivalent τ.

The value of Knuth equivalence is that it allows us, for each standard Young tableau S, to pick out a *canonical* permutation $\sigma \in S_n$ so that $S(\sigma) = S$. From a tableau S, we start at the bottom and read each row in order. With $\sigma = 82357146$, we know that

$$S(\sigma) = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 5 & 7 \\ \cline{1-3} 8 \\ \cline{1-1} \end{array},$$

so the canonical Knuth equivalent permutation is 82571346. (We'll also call this permutation the reading word of S.) The following sequence of elementary Knuth equivalent pairs demonstrates that these are indeed Knuth equivalent:

82357146
82351746
82351476
82315476
82135476
82153476
82513476
82513746
82517346
82571346

I do not know if this is the shortest sequence of elementary equivalent pairs.

Proof. (Sketch.) It is a routine, if somewhat unpleasant, computation to show that if σ and τ are Knuth equivalent, then $S(\sigma) = S(\tau)$: it suffices to check the result for elementary Knuth equivalent pairs. The other direction is more exciting. Suppose $S(\sigma) = S(\tau)$. It suffices to show the result when σ is a canonical representative of a Knuth equivalence class. For lack of better options, we use the following slightly unwieldy notational convention: if $\sigma \in S_n$ and $1 \le k \le n$, let $(\sigma, k.5)$ denote the element of S_{n+1} so that $(\sigma, k.5)(i) = \sigma(i)$ if $\sigma(i) \le k$, $(\sigma, k.5)(i) = \sigma(i)+1$ if $\sigma(i) > k$, and $\sigma(n + 1) = k + 1$.

By induction, it suffices to show that if $\sigma \in S_n$ is a canonical representative of a Knuth equivalence class, and $\sigma' \in S_{n+1}$ is the reading word of $(\sigma, k.5)$, then $(\sigma, k.5)$ and σ' are Knuth equivalent. Thus, we must only check that if $a_1 < \cdots < a_{j-1} < x < a_j < \cdots < a_k$, then the permutation $a_1 \cdots a_k x$ is Knuth equivalent to $a_1 \cdots a_j x a_{j+1} \cdots a_k$ and to $a_j a_1 \cdots a_{j-1} x a_{j+1} \cdots a_k$, which is straightforward to verify. ■

4. The shape of RSK

Using the notion of Knuth equivalence discussed in the previous section, we can state and prove Greene's generalization to Schensted's Theorem, as described in [1].

Definition 5. Let $\sigma \in S_n$ be a permutation. For a positive integer k, we say that a subsequence of σ is k-descending if it does not contain an increasing sequence of length $k + 1$. Similarly, we say that a subsequence is k-ascending if it does not contain a decreasing sequence of length $k + 1$. For a positive integer k, we let $d_k(\sigma)$ be the maximum length of a k-descending subsequence. Similarly, we let $a_k(\sigma)$ be the maximum length of a k-ascending subsequence.

Another way to detect k-descending subsequences is to note that they are the unions of k (possibly empty) decreasing subsequences.

Example. Let σ be the permutation 82357146. As we saw above, a longest 1-descending sequence is 821, and a longest 1-ascending sequence is 2357. A longest 2-descending sequence is 82514, and a longest 2-ascending sequence is 2357146. A longest 3-descending sequence is 8257146.

Theorem 6. *For a positive integer k, we have*

$$d_k(\sigma) = \lambda_1^* + \cdots + \lambda_k^*,$$
$$a_k(\sigma) = \lambda_1 + \cdots + \lambda_k.$$

Proof. Our method of attack is as follows: We first show that the lengths of maximal k-ascending and k-descending sequences are constant on Knuth equivalence classes. Then, we prove the result for the canonical representatives in each Knuth equivalence class.

To prove the first assertion, we must show the following: Suppose that $x < y < z$, and σ contains a consecutive subsequence of the form y, z, x or z, x, y. We must show that switching x and z does not change the lengths of maximal k-ascending or k-descending sequences. Since all possible cases are roughly symmetric, we'll just show that one of these moves does not increase $a_k(\sigma)$. Let's work with the case that y, z, x is a consecutive subsequence. Let τ be the permutation obtained by switching x and z. Suppose $a_k(\tau) > a_k(\sigma)$. This means that τ contains a k-ascending subsequence γ of length at least $a_k(\sigma)+1$. Hence, γ must contain both x and z, or else it would also be a k-ascending subsequence of σ. Write γ as the union of k ascending sequences: $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$. Then one of the γ_i must contain both x and z, or else nothing would have changed from switching x and z. But γ_i is still an increasing subsequence if we replace x by y. Thus σ also has a k-ascending subsequence of length $a_k(\tau)$ if γ did not contain y. Hence, γ must have contained y.

So, let's suppose that γ contains y. Switch x and z back again. By hypothesis, this subsequence contains a decreasing subsubsequence δ of length $k+1$, and δ must contain z and x (in that order). But, replacing z by y in δ gives us a decreasing subsubsequence of γ of length $k+1$, which contradicts the assumption that γ was k-ascending. The other cases are similar.

So, that shows that $a_k(\sigma)$ and $d_k(\sigma)$ are constant on Knuth-equivalence classes. Let's now show the result for the canonical representatives of these equivalence classes. That is, we must show that if σ is a canonical representative, then $a_k(\sigma) = \lambda_1 + \cdots + \lambda_k$. (The case of k-decreasing sequences follows from this and the observation we made earlier, that reversing the sequence transposes S.)

First, $a_k(\sigma) \geq \lambda_1 + \cdots + \lambda_k$, since we can start reading S from the k^{th} row and going up; this gives a k-ascending sequence since each row of S is increasing. Now, partition σ into λ_1 decreasing subsequences by reading the columns from bottom to top. A k-ascending sequence can only intersect a decreasing sequence in at most k elements, so if γ is any k-ascending sequence, it can only meet the i^{th} column in at most $\min(k, \lambda_i^*)$ elements. This number is maximized in all cases by choosing the sequence given above. This proves Greene's extension of Schensted's Theorem. ■

It is important to note that, while Greene's Theorem gives us an interpretation of the shape of S, it does not give any meaning to the individual parts of λ or λ^*. That is, λ_2 on its own (for example) has no natural combinatorial interpretation beyond its interpretation in conjunction with λ_1.

References

[1] Curtis Greene, *An extension of Schensted's theorem*, Advances in Math. **14** (1974), 254–265. MR 0354395 (50 #6874)

[2] Donald E. Knuth, *Permutations, matrices, and generalized Young tableaux*, Pacific J. Math. **34** (1970), 709–727. MR 0272654 (42 #7535)

[3] C. Schensted, *Longest increasing and decreasing subsequences*, Canad. J. Math. **13** (1961), 179–191. MR 0121305 (22 #12047)

Department of Mathematics, Stanford University, Stanford, CA 94305
E-mail address: simonr@math.stanford.edu