The Robinson-Schensted-Knuth correspondence (RSKc), informal lecture notes

Erik Aas

May 6, 2012

Problems

Erdos-Szekeres Theorem: Any permutation σ of $[r^2+1] = \{1, 2, \ldots, r^2, r^2+1\}$ contains a monotone subsequence of length $r+1$.
This theorem can be proved by induction in a not-so-illuminating way.

Proof. (sketchy) Given σ, suppose the theorem is false and all monotone subsequences have length $\leq r$ (by induction some monotone subsequence has length exactly r), and remove all the elements that are beginnings of longest monotone subsequences. There can be at most r beginnings of longest increasing sequences and likewise for beginnings of longest decreasing sequences, since a collection of beginnings of longest increasing sequences themselves form a decreasing subsequence, which by assumption has length $\leq r$. It is easy to see that at most $r+r-1$ elements (if the number of beginnings of increasing subsequences and of decreasing subsequences are both r, then the one of these must be the beginning of both an increasing and a decreasing subsequence) are removed by doing this. Now there are at least $r^2 + 1 - (2r - 1) = (r - 1)^2 + 1$ elements left and by induction there is a monotonic subsequence of length r among those, which is a contradiction. □

Knowledge of the RSKc trivializes the theorem above, as we shall see. Not surprisingly, the proof we get by using the correspondence is also more conceptual. In particular, we will be able to understand where the squaring in '$r^2 + 1$' comes into play rather than just making the induction work. (Of course, we could get rid of the '+1':s in the statement of the theorem by saying "the shortest (varying σ) possible longest monotone subsequence of the r^2-permutation σ has length r" instead.)

Problems: Find the longest common subsequence of two permutations π, σ of $[n]$, or, equivalently, find the longest increasing subsequence of $\sigma^{-1}\pi$.

The longest common subsequence problem has seen applications to genomics and department dinners. The special case considered here, that of permutations or of strings of equal length that use all characters from a given alphabet exactly once, does not easily extend to the general problem of finding the longest common subsequence of two strings. However, the RSKc will give us an efficient algorithm for solving this problem in this special, running in time $O(n \log n)$. This is substantially better than the best known algorithm for the general case. Strings having repeated characters and not having the same length (or same multiset of characters used) in contrast to permutations, present different problems when trying to extend from permutations to the general setting.

The Cauchy Identity. The following holds,

$$\sum_{\lambda} s_{\lambda}(x)s_{\lambda}(y) = \prod_{i,j\geq 1} \frac{1}{1 - x_i y_j}.$$

Here, s_{λ} is the Schur function. It will be defined below when the weight of a tableau is defined.

This was proved in a lecture earlier in the course in a way different from the proof to be given here. Also, there the Schur function was defined differently. I will not prove the definitions are equivalent. The Cauchy identity is included here because the RSKc being *weight-preserving* (as will be explained), seems to be the most natural combinatorial proof of this identity. In turn, the RSKc being weight-preserving is immediate from the definition to be given below.

Definitions

I will mention three correspondences, which form a chain of extensions (as do their names): the Schensted correspondence (Sc), the Robinson-Schensted correspondence (RSc) and the Robinson-Schensted-Knuth correspondence (RSKc).

The RSc was mentioned in Kurt's lectures, and it is a correspondence between permutations and pairs of Young tableaux. The RSKc generalizes 'permutation' to 'generalised permutation' and 'pair of standard young tableaux' to 'pair of semistandard young tableaux'. The Sc is an interesting special case of the RSc, the algorithm producing only the P tableau as opposed to RSc which produces both the P and Q tableaux.

[I now define the RSc. Below is the example I will use to do that.]

Example: Let $\pi = 75823146$.

$P(\pi)$	$Q(\pi)$
7	1
5 7	1 2
5 8 7	1 3 2
2 8 5 7	1 3 2 4
2 3 5 8 7	1 3 2 5 4
1 3 2 8 5 7	1 3 2 5 4 6
1 3 4 2 8 5 7	1 3 7 2 5 4 6
1 3 4 6 2 8 5 7	1 3 7 8 2 5 4 6

This process can be inverted, showing that this is indeed a bijection. It incidentally proves $\sum_\lambda \left(f^\lambda\right)^2 = n!$, where f^λ is the number of Young tableaux with strictly increasing rows and columns using all elements from $[n]$ exactly once.

We now make the last generalisation, obtaining the RSKc. Observe that a permutation can be identified with its permutation matrix, and that such matrices are square matrices with nonnegative integer entries. The RSKc is a bijection from the set of nonegative integer matrices to the set of semistandard Young tableaux (ssyt) with elements from $\mathbb{Z}_+$. A tableau is semi-standard if its rows and columns are weakly decreasing. The first step is to associate to each such matrix a *generalised permutation* in the following manner:

[example used:]

$$\begin{pmatrix} 0 & 2 & 1 \\ 1 & 3 & 0 \\ 2 & 1 & 0 \end{pmatrix}$$

will correspond to

$$\begin{pmatrix} 1 & 1 & 1 & 2 & 2 & 2 & 2 & 3 & 3 & 3 \\ 2 & 2 & 3 & 1 & 2 & 2 & 2 & 1 & 1 & 2 \end{pmatrix}$$

We always assume the top row in a generalized permutation is sorted, and that ties are broken by increasing values in the lower row. The RSKc is simply the natural extension of the RSc for these kinds of arrays (a permuation can also be thought of as a two-line array whose top row form a range of distinct integers starting with 1). We never bump out another element with an element of equal value; we instead find the smallest strictly larger to bump out. We insert the values in the top row into Q and the values in the bottom row into P, inserting a pair in the same column at the same time, starting from the left. However, we will only use the generalized permutations as the things 'on the left side' of the correspondence, not integer matrices.

Definition: The weight of a generalized permutation is the monomial $\prod_j x_{u_j} y_{l_j}$, where u_i are the elements of the upper row and l_i are the elements of the lower row. (So the variables are $x_1, x_2, \ldots$, $y_1, y_2, \ldots$, and the weight is a monomial in x and y.) The weight of a semi-standard Young tableau is $\prod_{\textbf{cell c of the tableau}} x_{\textbf{value in c}}$. It is a monomial in the variables $x_1, x_2, \ldots$.

We now define the Schur function $s_\lambda(\mathbf{x}) = s_\lambda(x_1, x_2, \ldots)$ by

$$s_\lambda(\mathbf{x}) = \sum_{U \text{ is a standard Young tableau with shape } \lambda} \textbf{weight}(U)(\mathbf{x}).$$

It is immediate from the definition that the weight of the generalised permutation π in $\mathbf{x}, \mathbf{y}$ is equal to the weight of $P(\pi)$ in $\mathbf{x}$ times the weight of $Q(\pi)$ in $\mathbf{y}$, since the same values occur with the same multiplicities in the tableaux and the gpermutation.

Properties

Here are some interesting properties of the RSc.

Theorem: Suppose π RS corresponds to (P, Q). Then

- **π^{-1} RS corresponds to (Q, P).**
- **π^r, the reversal of π (the reversal of 41235 is 53214), RS corresponds to $(P^t, ?)$, the transpose of P and tableau we will saying nothing about.**
- **The length of a longest increasing subsequence of π is equal to the number of elements in the top row of $P(\pi)$.**

A theorem of Curtis Greene proves a more general form of the last statement:

Theorem (Greene) Let π be a permutation and λ the partition given by $P(\pi)$. Then, for each k, $\lambda_1 + \ldots + \lambda_k$ is equal to the greatest

length of any subsequence of π that can be written as the union of k increasing subsequences of π.

This is a suitable generalisation for proving the third point above by induction. We will not do a full proof of any of these statements, but indicate the connection between the definition given here and the 'geometric' definition (due to Viennot) given by Kurt in the lectures. From the geometric definition the three properties will be clear.

Observe that the theorem of Greene does not mention $Q(\pi)$ at all. In fact, $P(\pi)$ in a sense encodes the essential information about increasing subsequences of π, while $Q(\pi)$ tells us exactly which of the permutations with similar 'increasing-subsequences properties' the pair (P, Q) corresponds to (or equivalently, what properties π^{-1} has). If we identify all permutations having the same P tableau, we obtain what is called the *plactic monoid* (monoid referring to the operation of concatentation; usually the plactic monoid is defined for general words over some alphabet, not just for permutations of $[n]$). An interesting result on the plactic monoid is that equivalence in this monoid is the same as Knuth equivalence, defined as considering any two permutations connected by moves of the form $xzy \leftrightarrow zxy$ or $yxz \leftrightarrow yzx$, where $x < y < z$, to be the same. These moves can be seen to preserve the length of the longest increasing subsequence (but not the sequences themselves, of course). Here xyz is understood as three letters standing next to each other in the permutation (so '2143576' $\leftrightarrow$ '2143756' is a valid move). This definition is the basis of 'plactic sorting', another way of finding a longest common subsequence of two strings. It also holds that each class in the plactic monoid contains exactly one involution (a permutation all whose cycles have length at most 2). This is one way of proving the identity $\sum_{\lambda \vdash n} f^\lambda = \#($ involutions on $[n])$. (Observe that an involution on n is automatically a permutation and that the right hand side can be simplified slightly.)

Recall the construction of the *shadow diagram* of an n-permutation from the lectures, and that the (x-)coordinates of the vertical lines of this diagram define the first row of $Q(\pi)$, while the (y-)coordinates of the horizontal define the first row of $P(\pi)$.

The northeast corners of the lines of the shadow diagram defines a new set of dots in $[n]^2$, whose shadow diagram similarly induces the second rows of $P(\pi)$ and $Q(\pi)$, and so on.

The way to see the connection with the definition given here (usually described as the *bumping procedure*) is as follows:

[The following might be easier to read during the actual seminar than standalone.]

Think of the bumping construction of $(P(\pi), Q(\pi))$ as a process proceeding in steps (insertion of $\pi_1, \pi_2, \ldots, \pi_n$). These steps correspond to the growth of the shadow diagram from the left to the right. Suppose π_k is about to be inserted into $(P(\pi), Q(\pi))$ bumping-wise. Consider the part of the shadow diagram whose x-values are strictly less than k. We will now relate the column $x = k$ in the shadow diagram to what happens in the first row of $(P(\pi), Q(\pi))$. If π_k is inserted at the end of the last row, this corresponds to a vertical ray emanating from (k, π_k) in the shadow diagram. If π_k bumps some value y out of the first row, this corresponds to the vertical line going up from (k, π_k) intersecting a line with constant y-value y 'coming from the left'. This intersection will give a northeast corner for the second shadow diagram, which will in the next stage (when drawing the shadow diagram of the northeast corners) correspond to the insertion of y into the second row of P.

Observe also that $P(\pi)$ is associated only with x-values of the permutation diagram and $Q(\pi)$ is associated only with y-values of the permutation diagram. Thus interchanging x and y -values, that is, taking the inverse of π, reasonably corresponds to interchanging P and Q.

Similar considerations prove the second point above.

Finally it's easy to see that the number of shadow lines in the shadow diagram is equal to the length of a longest increasing subsequence, as observed in Kurt's lectures.

Solutions to problems

The Erdos-Szekeres theorem follows almost immediately from the properties above; we know that the length of a longest increasing subsequence is the length of the first row and (since an increasing subsequence of the reversal of a permutation is a decreasing subsequence of the original permutation, or directly from the geometric definition) that the length of a longest decreasing subsequence is the number of elements in the first column of the P tableau. In order for a permutation to have both these numbers at most r, its corresponding tableaux must fit into a square of side r, which in particular means it can have at most r^2 elements.

From this, a natural generalisation (with analogous proof) is easily seen to hold: any permutation whose decreasing subsequences are bounded in length by a and whose increasing subequences are bounded in length by b can have at most ab elements. Any permutation on $ab + 1$ violates at least one of those conditions. In fact the RSKc allows us to give a structured description of all permutations extremal with respect to these properties; they are the ones with rectangular tableaux.

Using the bumping procedure to calculate the top row of the P tableau of a permutation clearly does not need knowing the other rows of the tableau. Thus we can simply keep track of the top row as we construct the P tableau of $\sigma^{-1}\pi$ in order to find the length of the longest common subsequence of π and σ. Since the top row is sorted, we can perform binary search in it leading to a $O(n \log n)$ time $O(n)$ space algorithm.

The proof of the Cauchy identity is a computation

$$\prod_{i,j\geq 1} \frac{1}{1 - x_i y_j} = \prod_{i,j\geq 1} \sum_{k\geq 0} x_i^k x_j^k =$$

$$\sum_{\pi \text{ is a generalised permutation with entries from } \mathbb{Z}_+} \mathbf{weight}(\pi)(\mathbf{x}, \mathbf{y}) =$$

$$\sum_{T \text{ and } U \text{ are ssyt with same shape}} \mathbf{weight}(T)(\mathbf{x})\mathbf{weight}(U)(\mathbf{y}) =$$

$$\sum_{\lambda} \left(\sum_{U \text{ has shape } \lambda} \mathbf{weight}(U)(\mathbf{x}) \right) \left(\sum_{T \text{ has shape } \lambda} \mathbf{weight}(T)(\mathbf{y}) \right) =$$

$$\sum_{\lambda} s_{\lambda}(\mathbf{x}) s_{\lambda}(\mathbf{y}).$$

Here we use the fact that the RSKc is a weight-preserving bijection in the middle equality.

References

M. Lothaire: *Algebraic Combinatorics on Words*
Bruce Sagan: *The Symmetric Group*