Longest Increasing Subsequences and Random Matrices

Anand Sarwate

Abstract. We investigate the longest increasing subsequence of a permutation, and relate its length to random matrices. We provide a simple card game as a vehicle for computing the length of the longest increasing subsequence in a permutation. We discuss the relationship between permutations and combinatorial structures called Young tableaux via the Schensted correspondence. By relating representations of the symmetric group and the unitary group, we express the distribution of the longest increasing subsequence in terms of a matrix integral over the unitary group.

1. Introduction. In this paper, we investigate a simple problem, finding the longest increasing subsequence of a permutation, and we relate it to random matrices. Although our investigation is straightforward, it passes through important areas of combinatorics and algebra, and uses important constructions in both of those areas. Our presentation is largely modeled on that of Aldous and Diaconis [**1**], but is shorter and designed for a wider audience. Two semesters of abstract algebra and some exposure to measure theory should permit following the ideas presented here.

More precisely, the problem is as follows. Let π be an arbitrary permutation of the integers $1, 2, \ldots, n$. Denote by $\pi(i)$ the ith element of the permutation. Then an *increasing subsequence* $(i_1, i_2, \ldots, i_k)$ of π is a subsequence satisfying

$$\begin{gathered} i_1 < i_2 < \cdots < i_k, \\ \pi(i_1) < \pi(i_2) < \cdots < \pi(i_k). \end{gathered}$$

We denote by $l(\pi)$ the length of the *longest* increasing subsequence. We denote by L_n the integer valued random variable that takes on the value $l(\pi)$ when the permutation π is drawn from a uniform distribution. The problem is to express the distribution of L_n in terms of random unitary matrices.

Others have extended the results in this paper. Rains [**7**, pp. 5–6] expresses matrix integrals on the orthogonal and symplectic groups to the distribution of the longest increasing subsequences of permutations given some restrictions. Odlyzko and Rains [**6**, p. 2] give explicit computations and Monte Carlo simulations to illustrate the behavior of L_n. Aldous and Diaconis [**1**, p. 416] provide several different perspectives on the asymptotic behavior of L_n as n becomes large.

In Section 2, we discuss patience sorting, a simple card-game model for computing $l(\pi)$, and provide some Monte Carlo simulations to show the empirical distribution of L_n. In Section 3, we introduce Young tableaux, a combinatorial construction with applications in representation theory and geometry. We construct the Schensted correspondence, which makes explicit the link between the set of Young Tableaux and the symmetric group. The Schensted correspondence is also used to compute the distribution of L_n.

In Section 4, we provide a quick summary of some facts from representation theory and express the distribution of L_n in terms of characters of the irreducible representations of the symmetric group S_n. In Section 5, we discusses power-sum symmetric functions and their role in linking the symmetric and unitary groups. These functions allow us to express the distribution of L_n in terms of the matrix integral of the trace of a power of a unitary matrix. In the Appendix, we provide a MATLAB program that computes the length of the longest increasing subsequence in a permutation.

2. Patience Sorting. Patience sorting is a card game first invented in the 1960s. Its full history is given in Aldous and Diaconis [**1**, p. 417]. Suppose we shuffle a deck of cards, numbered $1, 2, \ldots, n$, which we shuffle into a random order. We turn up cards one at a time and place them into piles according to the following rules:

1. A card may always be used to start its own pile to the right of all other piles.
2. A card may instead be placed on top of a card with a higher number.

The goal is to end with as few piles as possible.

Although this game seems uninteresting at first, there is a simple strategy for playing that yields an effective way of computing $l(\pi)$, called the *greedy strategy.* It involves placing the current card on the leftmost pile possible. This strategy can be explained with an example. Say we have a deck of ten cards in the following order:

$$8\ 3\ 7\ 9\ 2\ 5\ 4\ 1\ 10\ 6.$$

The greedy strategy plays the game in the following manner:

$$\begin{array}{c} \\ \\ 8 \end{array}\qquad \begin{array}{c} \\ 3 \\ 8 \end{array}\qquad \begin{array}{cc} & \\ 3 & \\ 8 & 7 \end{array}\qquad \begin{array}{ccc} & & \\ 3 & & \\ 8 & 7 & 9 \end{array}\qquad \begin{array}{ccc} 2 & & \\ 3 & & \\ 8 & 7 & 9 \end{array}\qquad \begin{array}{ccc} 2 & & \\ 3 & 5 & \\ 8 & 7 & 9 \end{array}$$

$$\begin{array}{ccc} & & \\ 2 & 4 & \\ 3 & 5 & \\ 8 & 7 & 9 \end{array}\qquad \begin{array}{ccc} 1 & & \\ 2 & 4 & \\ 3 & 5 & \\ 8 & 7 & 9 \end{array}\qquad \begin{array}{cccc} 1 & & & \\ 2 & 4 & & \\ 3 & 5 & & \\ 8 & 7 & 9 & 10 \end{array}\qquad \begin{array}{cccc} 1 & & & \\ 2 & 4 & & \\ 3 & 5 & 6 & \\ 8 & 7 & 9 & 10 \end{array}$$

It turns out that patience sorting is an easy way for us not only to calculate $l(\pi)$, but also to find an increasing subsequence in π of length $l(\pi)$.

Theorem 2-1. *Let π be an ordering of a deck of cards numbered $\{1, 2, \ldots, n\}$. Then the greedy strategy results in exactly $l(\pi)$ piles.*

Proof: We first prove that the number of piles is at least $l(\pi)$, and then show that the number of piles is at most $l(\pi)$.

Let $i_1 < i_2 < \cdots < i_k$ be an increasing subsequence of π. Then $\pi(i_j)$ must always lie in a pile to the right of $\pi(i_{j+1})$ since we are only allowed to place cards on top of cards of higher value. Therefore any valid strategy, and particularly the greedy strategy, results in at least $l(\pi)$ piles.

We can use patience sorting to find an instance of an increasing subsequence, thereby showing that we have at most $l(\pi)$ piles. Every time we place a card c in a pile i that is not the first pile, draw an arrow from c to the top card d of the preceding pile $i-1$. We know $d < c$; otherwise, in our greedy strategy, c would be placed on top of d. Note that each arrow goes from a later card to an earlier card. This construction is illustrated in Figure 2-1 using the example above.

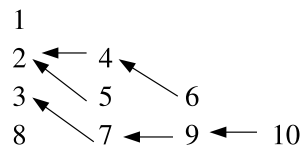

Figure 2-1. Constructing an increasing subsequence.

If we have k piles, call the the top card in the rightmost pile a_k. Then follow the arrow from a_k to a_{k-1} and so on. In this way we construct an increasing subsequence in π. Therefore we have at most $l(\pi)$ piles. □

The game is called patience *sorting* because, at the end, we can sort the cards into order. Card 1 will be at the top of some pile. Removing it, we are left with a patience-sorted deck of cards numbered $\{2, 3, \ldots, n\}$; so Card 2 must now be at the top of some pile. Proceeding as before, we can sort the deck.

In the Appendix we provide a MATLAB script to compute the distribution of L_n using this method. Histograms for $n = 10$ and $n = 20$ using 10^7 samples are shown in Figure 2-2. As we can see, the distribution is heavily clustered at the lower end of the graph.

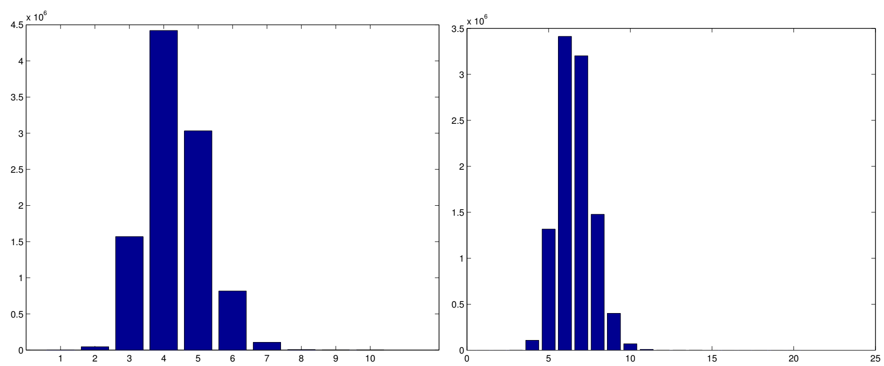

Figure 2-2. Histogram of number of piles for patience sorting on decks of 10 cards (left) and 20 cards (right) using 10^7 samples.

3. Young Tableaux. We now turn to another way of looking at permutations using diagrams called *Young tableaux*. These are special structures, which are used in invariant theory and group representations of the symmetric group S_n. They are also important in combinatorics and algebraic geometry, and in the theory of symmetric functions. What is most important for our problem is an important construction, called the Schensted correspondence, which relates multiset permutations to ordered pairs of Young tableaux.

Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ be a set of integers such that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_k > 0$ and $\sum \lambda_i = n$. Then we say λ is a *partition of* n and write $\lambda \vdash n$. We denote the number of entries in λ by $|\lambda|$, which in this case is equal to k. In the special case where $\lambda = (1, 1, \ldots, 1)$, we write $\lambda = 1^n$.

A *Ferrers diagram* with *shape* λ is a set of cells as shown on the left in Figure 3-1, where row i has λ_i cells. A *standard Young tableau* is a Ferrers diagram with the numbers $1, 2, \ldots, n$ in the cells such each number is used once, and the entries increase along each row and down each column.

1	2	7	10
3	5	9	
4	8		
6	11		

Figure 3-1. A Ferrers diagram for $(4, 3, 2, 2)$, and example of a Young tableau with shape λ.

The *hook length* h_c of a cell c is the number of cells to the right of c in its row plus the number of cells below c in its column plus one for the cell c itself. Thus, in the tableau in Figure 3-1, the hook length of the cell containing 5 is 4, and the hook length of the cell containing 8 is 2.

Denote by d_λ the number of standard Young tableaux of shape λ. Proposition 3-1 gives a surprisingly simple way of computing d_λ.

Proposition 3-1 (Hook Formula). *The number d_λ is given by the formula*

$$d_\lambda = \frac{n!}{\prod_c h_c}.$$

Unfortunately, no simple combinatorial proof of the hook formula exists. A number of outlines of existing proofs are given in [**8**, p. 266], and a probabilistic proof is given in [**5**, p. 54].

For example, if we look at the tableau in Figure 3-1, then we can calculate the number of tableau of shape $(4, 3, 2, 2)$:

$$d_{(4,3,2,2)} = \frac{11!}{7 \cdot 6 \cdot 3 \cdot 1 \cdot 5 \cdot 4 \cdot 1 \cdot 3 \cdot 2 \cdot 2 \cdot 1} = 1320.$$

This is far too many to explicitly verify, so we can look at an easier example. Consider $\lambda = (3, 2) \vdash 5$. The hook-length formula tells us we can construct $120/(4 \cdot 3 \cdot 1 \cdot 2 \cdot 1) = 5$ standard Young tableaux. These are shown in Figure 3-2. It is an exercise to prove there exist no more standard Young tableaux of that shape.

1	2	3
4	5	

1	3	5
2	4	

1	2	5
3	4	

1	2	4
3	5	

1	3	4
2	5	

Figure 3-2. All standard Young tableaux of shape (3,2).

Young tableaux are related to the permutation group by a construction called the Schensted correspondence, or perhaps more appropriately, the Robinson–Schensted–Knuth correspondence. For a more extensive discussion of the nomenclature, see Fulton [**5**, p. 38]. The Schensted correspondence is a bijection between the set of permutations and ordered pairs of Young tableaux. Consult Fulton [**5**, pp. 33–57] or Stanton and White [**10**, pp. 85–92] for more extensive treatments of the correspondence.

Theorem 3-2 (Schensted correspondence). *There exists a natural bijection between permutations $\pi \in S_n$ and ordered pairs of standard Young tableaux (P, Q) of the same shape $\lambda \vdash n$.*

Proof: We construct the pair (P, Q) from a permutation π to show that any permutation can be mapped to an ordered pair of Young tableaux. Next we prove that every (P, Q) arises from a unique permutation. Then we provide a means of recovering π from an arbitrary pair (P, Q) to prove every ordered pair corresponds to a unique permutation. In the spirit of the previous section, we denote the mth *card* in our permutation by $\pi(m)$.

Let π be a permutation of the numbers $1, 2, \ldots, n$. It helps to represent π in a two-line form, where the first line are the integers from 1 to n and the second line is the permutation:

$$\pi = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ 8 & 3 & 7 & 9 & 2 & 5 & 4 & 1 & 10 & 6 \end{pmatrix}. \tag{3-1}$$

We construct the Q-tableau from the first line and the P-tableau from the second by inserting one card at a time moving from left to right. Denote by P_m the tableau created after inserting m cards, and Q_m analogously. We start with 1 in the single cell of the left-most column of Q_1, and $\pi(1)$ in the single cell of the left-most column of P_1. To insert the m-th card, we use the following rules:

1. If $\pi(m)$ is larger than all cards in the current column of P_{m-1}, append $\pi(m)$ to the end of the current column.
2. If $\pi(m)$ is not larger than all cards in the current column of the P_{m-1}, replace the smallest card b such that $b > \pi(m)$ with $\pi(m)$. Now use Rules 1 and 2 to insert b into the next column to the right.
3. After $\pi(m)$ has been inserted into the P_{m-1}, we obtain P_m, which has the same shape as P_{m-1} except for a single added cell. Create Q_m by adding that same cell to Q_{m-1} with the number m in that cell.

As we can see, this algorithm produces two standard Young tableaux of the same shape from the permutation π. For example, in Figure 3-3 we construct the two tableaux for Permutation (3-1).

We must describe how to construct a permutation π from an arbitrary pair of standard Young tableaux. We essentially perform the reverse operation to the column insertion described in the first half of the proof. Again, we let P_m be the P-tableau with n cells, and Q_m the Q-tableau with m cells. The procedure is outlined below:

1. Remove the largest entry m in Q_m to form Q_{m-1}. Find the corresponding entry b in P_m and remove it from the tableau.
2. If b is in the first column of P, then set $\pi(m) = b$.
3. If b is not in the first column, we insert b into the column to the left. Find the largest entry c of this column such that $b > c$. Put b in the position of c and then place c according to rules 2 and 3.

It is clear that this procedure simply reverses the previous construction. Figure 3-4 provides an example of constructing π from a pair of Young tableaux. It is easy to see that π yields the given pair of tableaux under the map from permutations to tableaux described above.

We have exhibited a map that takes each permutation to an ordered pair of Young tableaux. The inverse takes each pair of tableaux to a permutation in S_n. $\square$

The Schensted correspondence has a number of interesting and beautiful properties, two of which are quickly described here to give a feeling for how useful the correspondence is.

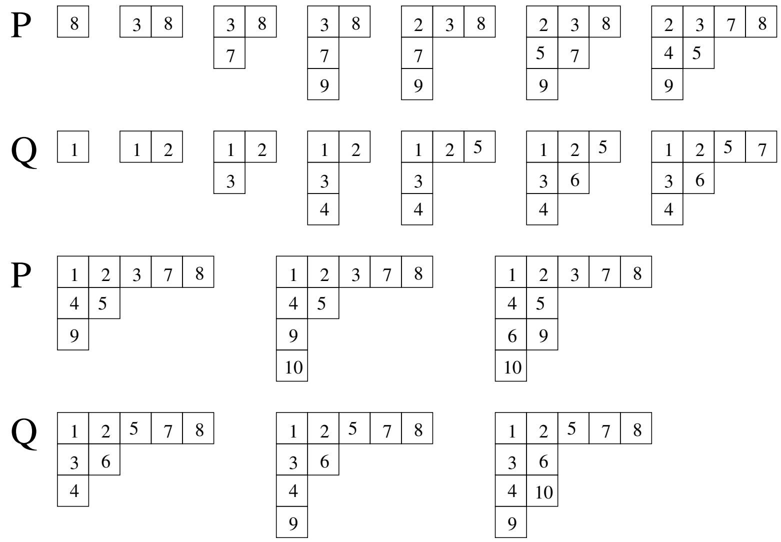

Figure 3-3. Successive steps of the Schensted correspondence of Permutation (3-1).

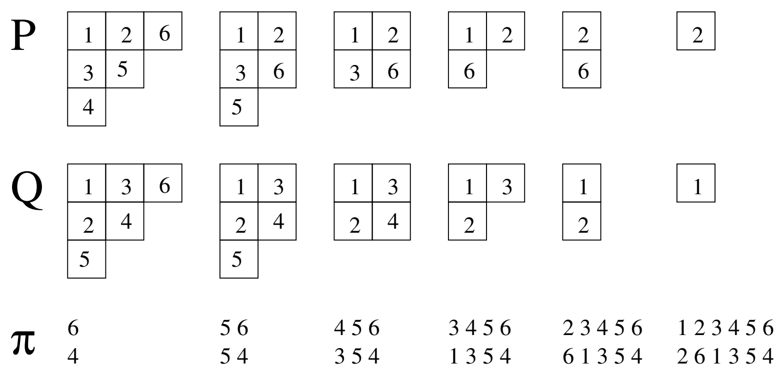

Figure 3-4. The Schensted correspondence from tableaux to permutation.

Proposition 3-3. *If π corresponds to (P, Q) under the Schensted correspondence, then π^{-1} corresponds to (Q, P).*

This proposition is not too difficult to prove, and a good proof is provided in [**10**, pp. 93–104].

An *involution* of $\{1, 2, \ldots, n\}$ is an element π of S_n such that $\pi = \pi^{-1}$. Thus Proposition 3-3 gives us the following corollary.

Corollary 3-4. *The number of involutions of* $\{1, 2, \ldots, n\}$ *is*

$$\sum_{\lambda \vdash n} d_\lambda.$$

Schensted's original motivation in constructing his correspondence was to investigate ways of computing $l(\pi)$ for a given permutation. Although this method is more tedious to perform than patience sorting, the following result is more useful to us because it allows us, via representation theory, to connect the distribution of L_n to random matrices. The proof of Proposition 3-5 is omitted here due to its length, and can be found in [**10**, pp. 93–104].

Proposition 3-5. *Given a permutation $\pi \in S_n$, the number of rows in the corresponding P-tableau under the Schensted correspondence is equal to $l(\pi)$. The number of columns is equal to the length of the longest decreasing subsequence.*

We can now obtain an explicit formula for the distribution of L_n by counting the number of Young tableaux with a certain number of rows.

Corollary 3-6. *The distribution of L_n can be expressed as follows:*

$$P(L_n = l) = \frac{1}{n!} \sum_{\lambda \vdash n,\ |\lambda| = l} (d_\lambda)^2.$$

4. Representations of S_n. In this section, we review some basics of representation theory, and describe the link between the representations of the symmetric group S_n and standard Young tableaux. For the basics of group representations, consult the algebra textbook by Artin [**2**, pp. 307–344]. The books by Diaconis [**3**, pp. 5–16] and Sagan [**9**, pp. 1–51] cover the basics as well as particular results on the representations of the symmetric group.

Let G be a group and V be a finite-dimensional vector space over the complex numbers $\mathbb{C}$. Then a *representation of* G is a homomorphism $\rho \colon G \to GL(V)$, where $GL(V)$ is the general linear group of all invertible linear transformations from V to itself. The space V is called the G-module associated with ρ because G acts on the space V. The *degree* of ρ is defined to be $\dim(V)$.

The representation $\rho \colon G \to GL(V)$ that sends all elements of G to the identity is called the *trivial representation*, and has degree 1. If $G = S_n$, then we can let ρ map G to the set of $n \times n$ permutation matrices. This is called the *defining representation* of S_n, and has degree n.

A G-module V is called *irreducible* if V has no proper subspace W that is invariant under the action of G. Mashke's theorem states that every G-module can be written as the direct sum of irreducible G-modules. The number of irreducible representations is equal to the number of conjugacy classes of G.

Consider the group S_n. The *cycle type* of a permutation is a list $(\lambda_1, \lambda_2, \dots, \lambda_k)$ of the cycle lengths in the permutation. The conjugacy classes of S_n consist of permutations having the same cycle type. Note that a cycle type is the same as a partition of n. We can in fact identify with each partition $\lambda \vdash n$ an irreducible module S^λ of S_n, called its *Specht module*.

The *character* of a representation ρ is the function $\chi \colon G \to \mathbb{C}$ defined by the equation

$$\chi(g) = tr(\rho(g)).$$

The character is constant on the conjugacy classes of G. Let $e \in G$ be the identity element. Then every representation must map e to the identity matrix in $GL(V)$. Thus $\chi(e) = tr(\rho(e)) = \dim(V)$. In the case where $G = S_n$ we denote the characters of S^λ by χ^λ. The most important fact for us is stated below in Lemma 4-1 without proof; consult [**9**, pp. 66–74].

Lemma 4-1. *Let λ be a partition of n, let S^λ be the corresponding Specht module, and let e be the identity element of S_n. Then*

$$d_\lambda = \dim(S^\lambda).$$

This is a powerful result, which relates the combinatorial enumeration of standard Young tableaux to the representations of the symmetric group. The proof involves a construction, which turns the standard Young tableaux of shape λ into a basis for the Specht module S^λ. Combining Lemma 4-1 and Corollary 3-6 we obtain the following expression for the cumulative distribution of L_n.

Proposition 4-2. *The cumulative distribution of L_n can be expressed as follows:*

$$P(L_n \leq l) = \frac{1}{n!} \sum_{\lambda \vdash n, |\lambda| \leq l} (\chi^\lambda(e))^2.$$

5. Random Matrices and L_n. We are now ready to connect the results from the previous sections to what we know about random matrices. We have expressed the density of L_n in terms of the characters of the symmetric group. Frobenius established an explicit relationship, via power-sum symmetric functions, between the characters of the symmetric group and the characters of the unitary group. By exploiting this relation, we can express the quantity in Proposition 4-2 in terms of an inner product of two power-sum symmetric functions, which is also a way of calculating a matrix integral over the unitary group.

The power-sum symmetric functions play an important role in the following analysis, since they are the link between random matrices and random permutations. Let U be a $l \times l$ matrix with eigenvalues $x_1, x_2, \ldots, x_l$. Then the power sum symmetric function P_j is given by

$$P_j(U) = \sum_{i=1}^{l} x_i^j.$$

If $\lambda \vdash n$, then we define

$$P_\lambda = \prod_{j=1}^{|\lambda|} P_{\lambda_j}.$$

We define the inner product of two power-sum symmetric functions as the expectation of their conjugate product using normalized Haar measure on the unitary group:

$$\langle P_\lambda, P_\mu \rangle = E_{U \in U(l)}[P_\lambda(U)\overline{P_\mu(U)}] = \frac{1}{n!} \int_{U \in U(l)} P_\lambda(U)\overline{P_\mu(U)}\, dU. \tag{5-1}$$

The *Schur functions* $\{s_\mu \mid \mu \vdash n\}$ are the irreducible characters of the unitary group of $l \times l$ matrices $U(l)$. Their inner product is defined analogously:

$$\langle s_\mu, s_\nu \rangle = E_{U \in U(l)}[s_\mu(U)\overline{s_\nu(U)}] = \frac{1}{n!} \int_{U \in U(l)} s_\mu(U)\overline{s_\nu(U)}\, dU.$$

They have the properties that $s_\mu = 0$ if $|\mu| > l$ and $\langle s_\mu, s_\nu \rangle = \delta_{\mu\nu}$ [**4**, p. 52]. Frobenius proved that the Schur functions form a basis for the power-sum symmetric polynomials:

$$P_\lambda = \sum_{\mu \vdash n} \chi^\mu(\lambda) s_\mu \tag{5-2}$$

where $\chi^\mu(\lambda)$ is the character of the Specht module S^μ of the symmetric group S_n evaluated on the conjugacy class of λ. This important result allows us to relate the characters of the symmetric and unitary groups.

Using (5-2) we now express the expectation in (5-1) in terms of Schur functions and characters of the symmetric group:

$$E_{U \in U(l)}[P_\lambda(U)\overline{P_\mu(U)}] = \frac{1}{n!} \sum_{\nu,\kappa \vdash n} \chi^\nu(\lambda)\chi^\kappa(\mu) E_{U\in U(l)}(s_\nu(U)s_\kappa(U)). \tag{5-3}$$

We are now ready to prove our the main result of this paper.

Theorem 5-1. *The cumulative distribution function of the random variable L_n is given by the following equation:*

$$P(L_n \leq l) = \frac{1}{n!} \int_{U(l)} |Tr(U)^n|^2 \, dU. \tag{5-4}$$

where dU is the normalized Haar measure on the group of $l \times l$ unitary matrices.

Proof: The power-sum symmetric function $P_1(U)$ is simply the trace of U. If $\lambda = 1^n$, then $P_\lambda(U) = Tr(U)^n$. Hence the integral on the right-hand side of Equation (5-4) is the expected value of the product of $P_{1^n}(U)$ with itself. We expand this product using Equation (5-3) with $\lambda = \mu = 1^n$ to obtain the expression:

$$\frac{1}{n!} \int_{U(l)} |Tr(U)^n|^2 \, dU = \frac{1}{n!} \sum_{\nu,\kappa \vdash n} (\chi^\nu(1^n)\chi^\kappa(1^n)) E_{U\in U(l)}[s_\nu(U)s_\kappa(U)]. \tag{5-5}$$

We now apply the properties of Schur functions discussed earlier. Therefore all terms with $\nu \neq \kappa$, and $|\nu| > l$ in Equation (5-5) vanish, and we are left with

$$\frac{1}{n!} \int_{U(l)} |Tr(U)^n|^2 \, dU = \frac{1}{n!} \sum_{\nu \vdash n, \; |\nu| \leq l} (\chi^\nu(1^n))^2.$$

But this is the same expression as in Proposition 4-2, since the partition 1^n corresponds to the identity permutation e. The result follows immediately. $\square$

APPENDIX. COMPUTING $l(\pi)$ WITH PATIENCE SORTING

Below is code in MATLAB used to compute a histogram of $l(\pi)$.

```
%% patienceSort.m
%%
%% Performs greedy patience sorting on decks of length n
%% for given number of samples.  The output histogram is
%% saved in v.
```

```
n = 10;
samples = 1e7;

v = zeros(n,1);

for loop1 = 1:samples
  ord = randperm(n);
  piles = n;

  for ind1=1:length(ord)
    curr = ord(ind1);

    [val, pos] = find(piles > curr);

    if (val ~= [])
      piles(min(pos)) = curr;
    else
      piles = [piles curr];
    end
  end

  v(length(piles)) = v(length(piles)) + 1;
end
```

References

[**1**] Aldous, D., and Diaconis, P., *Longest increasing subsequences: from patience sorting to the Baik–Deift–Johansson theorem*, Bulletin of the AMS **36** (1999), no. 4, 413–432.

[**2**] Artin, M., "Algebra", Prentice Hall, Upper Saddle River, NJ, 1991.

[**3**] Diaconis, P., "Group Representations in Probability and Statistics," Institute of Mathematical Statistics, Hayward, California, 1998.

[**4**] Diaconis, P., and Shahshahani, M., *On the eigenvalues of random matrices*, J. Appl. Prob. **31** (1994), 49–61.

[**5**] Fulton, W., "Young Tableaux: With Applications to Representation Theory and Geometry," Cambridge University Press, Cambridge, 1997.

[**6**] Odlyzko, A. M., and Rains, E. M., *On longest increasing subsequences in random permutations*, Tech. report, AT&T Labs, 1998.

[**7**] Rains, E. M., *Increasing subsequences and the classical groups*, Elec. J. Comb. **5** (1998), no. 1, R12.

[**8**] Sagan, B. E., *The Ubiquitous Young Tableau*, in "Invariant Theory and Tableaux," Stanton, D., (ed.), Springer-Verlag, New York, 1990, pp. 262–298.

[**9**] Sagan, B. E., "The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions," Springer-Verlag, New York, 2001.

[**10**] Stanton, D., and White, D., "Constructive Combinatorics," Springer-Verlag, New York, 1986.