Subsequence containment by involutions

Aaron D. Jaggard*
Department of Mathematics
Tulane University
New Orleans, LA 70118 USA
adj@math.tulane.edu

Submitted: Mar 2, 2002; Accepted: Mar 4, 2005; Published: Mar 14, 2005
MR Subject Classifications: 05A05, 05A15, 05E10

Abstract

Inspired by work of McKay, Morse, and Wilf, we give an exact count of the involutions in $\mathcal{S}_n$ which contain a given permutation $\tau \in \mathcal{S}_k$ as a subsequence; this number depends on the patterns of the first j values of τ for $1 \leq j \leq k$. We then use this to define a partition of $\mathcal{S}_k$, analogous to Wilf-classes in the study of pattern avoidance, and examine properties of this equivalence. In the process, we show that a permutation $\tau_1 \dots \tau_k$ is layered iff, for $1 \leq j \leq k$, the pattern of $\tau_1 \dots \tau_j$ is an involution. We also obtain a result of Sagan and Stanley counting the standard Young tableaux of size n which contain a fixed tableau of size k as a subtableau.

1 Introduction

Given a permutation $\pi = \pi_n \dots \pi_n$ in the symmetric group $\mathcal{S}_n$ and a word $\sigma = \sigma_1 \dots \sigma_k$ of k distinct letters, say that π *contains* σ *as a subsequence* if $\pi_{i_1} = \sigma_1, \dots, \pi_{i_k} = \sigma_k$ for some $i_1 < \dots < i_k$. Recent work of McKay, Morse, and Wilf [MMW02] implies that the probability that an involution in $\mathcal{S}_n$ contains any fixed $\tau \in \mathcal{S}_k$ as a subsequence is $1/k! + o(1)$ as $n \to \infty$. We sharpen this aspect of their work with the following theorem.

Theorem 2.6. *For a fixed permutation $\tau = \tau_1 \tau_2 \dots \tau_k \in \mathcal{S}_k$ and $n \geq k$, the number $I_n(\tau)$ of involutions in $\mathcal{S}_n$ which contain τ as a subsequence equals*

$$\sum_j{}' \binom{n-k}{k-j} t_{n-2k+j}$$

*Part of this work is drawn from the author's Ph.D. dissertation which was written under the direction of Herbert S. Wilf at the University of Pennsylvania and partially supported by ONR grant N00014–01–1–0431 and the DoD University Research Initiative (URI) program administered by the ONR under grant N00014–01–1–0795. The part of this work carried out at Tulane University was partially supported by NSF grant DMS–0239996. The presentation of part of this work at the Permutation Patterns 2003 conference was partially supported by Penn GAPSA and the New Zealand Institute for Mathematics and its Applications.

where the sum is taken over $j = 0$ and those $j \in [k]$ such that the pattern of $\tau_1 \dots \tau_j$ is an involution in $\mathcal{S}_j$.

Here, and throughout this paper, we use t_n to denote the number of involutions in $\mathcal{S}_n$. We use $[k]$ for $\{1, \dots, k\}$; recall that the *pattern* of a word $\sigma = \sigma_1 \dots \sigma_k$ of j distinct letters is the order-preserving relabeling of the letters of σ with $[j]$. We will refer to the pattern of $\sigma_1 \dots \sigma_j$ as the *length j initial pattern of σ*.

We define σ and τ to be equivalent iff, for every n, $I_n(\sigma) = I_n(\tau)$; this leads to an apparently new classification of permutations. Defining $\mathcal{J}(\tau)$ to be the set of indices j over which the sum in Theorem 2.6 is taken, σ and τ are equivalent iff $\mathcal{J}(\sigma) = \mathcal{J}(\tau)$ and $\sigma, \tau \in \mathcal{S}_k$ for some k. We then examine some enumerative results relating to the sets $\mathcal{J}(\tau)$ in general; this leads us to an apparently new characterization of layered permutations.

Proposition 3.11. *A permutation $\tau \in \mathcal{S}_k$ is a layered permutation if and only if $\mathcal{J}(\tau) = \{0, 1, \dots, k\}$,* i.e.*, if and only if the pattern of $\tau_1 \dots \tau_j$ is an involution for every $j \in [k]$.*

We also prove the following theorem about $\mathcal{J}(\tau)$ for those τ corresponding to a given Young tableau.

Lemma 3.15. *Fix k and $\lambda \vdash k$, and let T be a standard Young tableau of shape λ. Let $\mathcal{T}$ be the set of permutations which correspond to (T, Q) for some Q, let $a_0 = f^\lambda$, and for $1 \le j \le k$ define a_j to be the number of $\tau \in \mathcal{T}$ for which $j \in \mathcal{J}(\tau)$. Then (with $f^\emptyset = 1$)*

$$a_j = \sum_{\mu \vdash j} f^{\lambda/\mu} \quad (j = 0, 1, \dots, k),$$

where $f^{\lambda/\mu}$ is the number of standard Young tableaux of skew shape λ/μ.

This leads to a result of Sagan and Stanley [SS90] counting the Young tableaux of size $n \ge k$ in which the entries $1, \dots, k$ form a specified subtableau.

Section 2 contains Theorem 2.6 and related material. Section 3 covers enumerative questions related to the sets $\mathcal{J}(\sigma)$ and connects these ideas to previous work on Young tableaux.

2 P-quasirandomness

2.1 Definition

Quasirandom permutations were introduced by McKay, Morse, and Wilf [MMW02] and are defined as follows.

Definition 2.1. Let $\mathcal{P}_n \subseteq \mathcal{S}_n$ be a non-empty set of permutations for infinitely many values of n, and let $\mathcal{P} = \cup_n \mathcal{P}_n$. For a word σ of k distinct letters from $[n]$, let $h(n, \sigma)$ be the number of permutations in $\mathcal{P}_n$ which contain σ as a subsequence. If $\mathcal{P}_n \neq \emptyset$, define

$g(n,\sigma) = h(n,\sigma)/|\mathcal{P}_n|$, the probability that $\pi \in \mathcal{P}_n$ contains σ as a subsequence. $\mathcal{P}$ is *quasirandom* (or a *quasirandom family of permutations*) if, for every $k \geq 1$, we have

$$\lim_{n\to\infty} \max_{\sigma} \left| g(n,\sigma) - \frac{1}{k!} \right| = 0, \tag{1}$$

where the maximum is over all sequences σ of k distinct elements of $[n]$ and the limit is over those n such that $\mathcal{P}_n \neq \emptyset$.

McKay, Morse, and Wilf used the quasirandomness of the set of involutions to prove theorems about the entries of Young tableaux. Their results only need Equation 1 to hold for $\sigma \in \mathcal{S}_k$ and not necessarily for arbitrary words σ of k distinct integers; this leads us to define the strictly weaker notion of *p-quasirandom* (*permutation-quasirandom*) permutations, whose definition repeats that of quasirandom permutations except that the word σ is replaced by the permutation $\tau \in \mathcal{S}_k$.

Definition 2.2. Let $\mathcal{P}_n \subseteq \mathcal{S}_n$ be non-empty for infinitely many values of n, and let $\mathcal{P} = \cup_n \mathcal{P}_n$. For $\tau \in \mathcal{S}_k$, let $h(n,\tau)$ be the number of permutations in $\mathcal{P}_n$ which contain τ as a subsequence. Let $f(n,\tau) = h(n,\tau)/|\mathcal{P}_n|$, the probability that a permutation in $\mathcal{P}_n$ contains τ, if $\mathcal{P}_n \neq \emptyset$. $\mathcal{P}$ is a *p-quasirandom* (*permutation-quasirandom*) *family of permutations* if, for all $k \geq 1$

$$\lim_{n\to\infty} \max_{\tau} \left| f(n,\tau) - \frac{1}{k!} \right| = 0, \tag{2}$$

where the maximum is now taken over all $\tau \in \mathcal{S}_k$ and the limit is again restricted to be over values of n for which $\mathcal{P}_n \neq \emptyset$.

It is clear that if $\mathcal{P}$ is quasirandom, then $\mathcal{P}$ is p-quasirandom. The following example shows that the converse is not true, so p-quasirandomness is indeed strictly weaker than quasirandomness.

Example 2.3. Define

$$\mathcal{P}_n = \begin{cases} \{\pi \in \mathcal{S}_n | \pi \text{ fixes } [n] \setminus [n/2]\}, & n \text{ even} \\ \emptyset, & n \text{ odd} \end{cases}.$$

For fixed k, for every even $n \geq 2k$ we see that every $\pi \in \mathcal{P}_n$ contains the k-element list $\sigma_n = (n-k+1)(n-k+2)\cdots(n-1)n$ as a subsequence. Thus $g(n,\sigma_n) = 1$, and we have that the limit in Equation 1 is equal to $1 - 1/k!$ and $\{\mathcal{P}_n\}$ is not a quasirandom family.

However, for fixed k and every even $n \geq 2k$, the probability that $\pi \in \mathcal{P}_n$ contains $\tau \in \mathcal{S}_k$ as a subsequence is just the probability that π' chosen uniformly at random from $\mathcal{S}_{n/2}$ contains τ as a subsequence, *i.e.*, $1/k!$. Thus $f(n,\tau) = 1/k!$ for every $\tau \in \mathcal{S}_k$ and even $n \geq 2k$, so the limit in Equation 2 equals 0 and $\{\mathcal{P}_n\}$ is a p-quasirandom family.

As illustrated by this example, p-quasirandomness is weaker than quasirandomness because the contained permutation τ is fixed before taking a limit. The following proposition shows that the fixed word τ need not be a permutation.

Proposition 2.4. *If $\{\mathcal{P}_n\}$ is a p-quasirandom family of permutations, then for any fixed word $\sigma = \sigma_1 \dots \sigma_k$ of k distinct letters the limiting probability that $\pi \in \mathcal{P}_n$ contains σ is $1/k!$ as $n \to \infty$.*

Proof. Let m be the largest letter which appears in the word σ. The probability that $\pi \in \mathcal{P}_n$ contains π is $1/m! + o(1)$ as $n \to \infty$ for each of the $\binom{m}{k}(m-k)! = m!/k!$ permutations which contain σ as a subsequence. □

2.2 The p-quasirandomness of involutions

Because the set of involutions is p-quasirandom, the probability that an n-involution contains a given k-permutation as a subsequence is $1/k! + o(1)$ as $n \to \infty$. We now sharpen this result by obtaining an exact count of the n-involutions which contain a given k-permutation as a subsequence, starting with the following lemma.

Lemma 2.5. *For $\tau = \tau_1 \dots \tau_k \in \mathcal{S}_k$ and $0 \le j \le k \le n$, the number of n-involutions π which contain τ as a subsequence and which map exactly j elements of $[k]$ into $[k]$ is $\binom{n-k}{k-j} t_{n-2k+j}$, if either $j = 0$ or the pattern of $\tau_1 \dots \tau_j$ is a j-involution, and 0 otherwise.*

Proof. Fix $\tau \in \mathcal{S}_k$ and an n-involution π containing τ as a subsequence. Let $A = \{a_1, \dots, a_j\}$ be the elements of $[k]$ which are mapped by π into $[k]$, with $a_1 < a_2 < \dots < a_j$ if $A \neq \emptyset$.

Now assume that $j \neq 0$. Because $\pi(a_i) \in [k]$ by definition and $\pi(\pi(a_i)) = a_i \in [k]$, $\pi(a_i) \in A$ and the restriction of π to A is an involution in the group of permutations of A. Because π contains τ as a subsequence and a_1 is the smallest element of $[k]$ which is mapped by π into $[k] = \{\tau_1, \dots, \tau_k\}$, we have $\pi(a_1) = \tau_1$; in general, $\pi(a_i) = \tau_i \in A$ for each $i \in [j]$. Combining this with the fact that $a_1 < a_2 < \dots < a_j$ is the ordering of $A = \{\tau_1, \dots, \tau_j\}$ in increasing order and the fact that the restriction of π to A is an involution shows that the pattern $x_1 \dots x_j$ of $\tau_1 \dots \tau_j$ is an involution in $\mathcal{S}_j$.

If $j = 0$, or if $1 \le j \le k$ and the length j initial pattern of τ is a j-involution, we may construct the permutations π which contain τ as a subsequence and which map j elements of $[k]$ into $[k]$. The requirement $\pi(a_i) = \tau_i$ noted above defines π on A. For $\tau_i \notin A$, $\pi^{-1}(\tau_i) = \pi(\tau_i) \notin [k]$, so for the preimages under π of the elements of $[k] \setminus A$ we choose $k - j$ elements $b_1 < \dots < b_{k-j}$ from $[n] \setminus [k]$. Because π must contain the τ_i in the order $\tau_1, \dots, \tau_k$, we have $\pi(b_i) = \tau_{i+j}$, with the involuting nature of π forcing $\pi(\tau_{i+j}) = b_i$. This completes the definition of π on the $2k - j$ elements of $[k] \cup \{b_i\}_{i \in [k-j]}$. We may define π on the remaining elements of $[n]$ by choosing one of the t_{n-2k+j} involutions in the group of permutations of $[n] \setminus \{1, \dots, k, b_1, \dots, b_{k-j}\}$. □

From this lemma, we may immediately count the n-involutions which contain a fixed k-permutation τ.

Theorem 2.6. *For a fixed permutation $\tau = \tau_1 \tau_2 \dots \tau_k \in \mathcal{S}_k$ and $n \ge k$, the number $I_n(\tau)$ of involutions in $\mathcal{S}_n$ which contain τ as a subsequence equals*

$$\sum_j{}' \binom{n-k}{k-j} t_{n-2k+j} \tag{3}$$

where the sum is taken over $j = 0$ and those $j \in [k]$ such that the pattern of $\tau_1 \dots \tau_j$ is an involution in $\mathcal{S}_j$. □

This leads to the following corollaries.

Corollary 2.7. *The probability that an n-involution contains $\tau \in \mathcal{S}_k$ as a subsequence equals*

$$\sum_j{}' \binom{n-k}{k-j} \frac{t_{n-2k+j}}{t_n}, \tag{4}$$

where the sum is taken over $j = 0$ and those $j \in [k]$ such that the pattern of $\tau_1 \dots \tau_j$ is an involution in $\mathcal{S}_j$. □

Remark 2.8. For every τ in $\mathcal{S}_k$, $k \geq 2$, the sum in Equation 4 is taken over at least $j = 0$, 1, and 2.

Corollary 2.9. *The probability that an n-involution contains the subsequence 1 equals 1 for every positive value of n. For $n \geq 2$ and $\tau \in \mathcal{S}_2$, the probability that an n-involution contains τ as a subsequence is exactly $1/2$.*

Corollary 2.10. *For $k > 2$ and $\tau \in \mathcal{S}_k$, a lower bound for the probability that an n-involution contains τ as a subsequence is given by*

$$\binom{n-k}{k} \frac{t_{n-2k}}{t_n} + \binom{n-k}{k-1} \frac{t_{n-2k+1}}{t_n} + \binom{n-k}{k-2} \frac{t_{n-2k+2}}{t_n} \leq f_{inv}(n, \tau). \tag{5}$$

Furthermore, it is possible for equality to hold.

Proof. The bound follows from Remark 2.8 above. The permutation $k12\dots(k-1)$ has length j initial pattern $j12\dots(j-1)$, which is not an involution in $\mathcal{S}_j$ for $j > 2$; equality holds for this permutation. □

Corollary 2.11. *For $k > 2$ and $\tau \in \mathcal{S}_k$, an upper bound for the probability that an n-involution contains τ as a subsequence is given by*

$$f_{inv}(n, \tau) \leq \sum_{j=0}^{k} \binom{n-k}{k-j} \frac{t_{n-2k+j}}{t_n}. \tag{6}$$

Furthermore, it is possible for equality to hold.

Proof. The bound follows immediately from Equation 4, with equality holding for, *e.g.*, $\tau = 12\dots k$. □

Remark 2.12. Proposition 3.5 below shows that there are exactly 2^{k-1} permutations $\tau \in \mathcal{S}_k$ for which equality holds in Equation 6.

Asymptotically expanding the terms in Equations 5 and 6, we obtain the following result on the asymptotic probability that an n-involution contains a specified subsequence. This refines the value $1/k! + o(1)$ that is implied by quasirandomness.

Proposition 2.13. *For $k > 2$, $\tau \in \mathcal{S}_k$, the probability as $n \to \infty$ that an n-involution π contains τ as a subsequence is*

$$\frac{1}{k!} - \frac{2}{3(k-3)!}n^{-3/2} + O(n^{-2}) \tag{7}$$

if the pattern of $\tau_1\tau_2\tau_3$ is not an involution in $\mathcal{S}_3$ and

$$\frac{1}{k!} + \frac{1}{3(k-3)!}n^{-3/2} + O(n^{-2}) \tag{8}$$

if the pattern of $\tau_1\tau_2\tau_3$ is an involution in $\mathcal{S}_3$.

Remark 2.14. Theorem 2.6 and its corollaries have natural analogues for fixed-point-free involutions [Jag03].

3 Applications

3.1 Classifying permutations

We now consider when permutations are equally restrictive with respect to subsequence containment by involutions. This has strong parallels to the notion of *Wilf-equivalence* from the study of pattern-avoiding permutations.

Definition 3.1. We say that two permutations σ and τ are *equivalent with respect to subsequence containment by involutions* iff for every n, the number of n-involutions which contain σ as a subsequence equals the number which contain τ as a subsequence.

Note that replacing "the number of n-involutions" with "the number of n-permutations" in this definition leads to a trivial equivalence.

Our classification of permutations using this equivalence will make use of Proposition 3.4 below, for which we need the following definition.

Definition 3.2. For $\tau \in \mathcal{S}_k$, the *j-set of τ*, denoted $\mathcal{J}(\tau)$, is the set containing 0 and exactly those $i \in [k]$ such that the length i initial pattern of τ is an involution. This is the set of indices j over which the sum in Equation 3 is taken when counting the n-involutions which contain τ as a subsequence.

Example 3.3. For $\sigma = 351264$ and $\tau = 524163$, $\mathcal{J}(\sigma) = \mathcal{J}(\tau) = \{0, 1, 2, 4, 5\}$. To count the n-involutions containing σ we use the sum in Equation 3, taken over $j = 0, 1, 2, 4, 5$; as the same indices are used when counting the n-involutions containing τ, σ and τ are equivalent in the sense of Definition 3.1.

Proposition 3.4. *Two permutations are equivalent with respect to subsequence containment by involutions iff they are of the same length and their j-sets are identical.*

$\mathcal{J}(\tau)$	$\lvert\{\tau\}\rvert$	$I_3(\tau)$	$I_4(\tau)$	$I_5(\tau)$	$I_6(\tau)$	$I_7(\tau)$	$I_8(\tau)$	$I_9(\tau)$	$I_{10}(\tau)$
$\{0,1,2\}$	2	0	1	3	10	32	110	386	1428
$\{0,1,2,3\}$	4	1	2	5	14	42	136	462	1660

Table 1: Classifying $\mathcal{S}_3$ by subsequence containment by involutions.

$\mathcal{J}(\tau)$	$\lvert\{\tau\}\rvert$	$I_4(\tau)$	$I_5(\tau)$	$I_6(\tau)$	$I_7(\tau)$	$I_8(\tau)$	$I_9(\tau)$	$I_{10}(\tau)$
$\{0,1,2\}$	6	0	0	1	4	17	65	260
$\{0,1,2,3\}$	8	0	1	3	10	33	115	416
$\{0,1,2,4\}$	2	1	1	3	8	27	91	336
$\{0,1,2,3,4\}$	8	1	2	5	14	43	141	492

Table 2: Classifying $\mathcal{S}_4$ by subsequence containment by involutions.

Proof. Assume that for two distinct permutations $\sigma \in \mathcal{S}_k$ and $\tau \in \mathcal{S}_{k'}$ the corresponding sums in Equation 3 are equal for every value of n. Because the limiting value (as $n \to \infty$) of these sums divided by t_n equals $1/k!$ and $1/k'!$, respectively, we must have $k' = k$. As no two terms in the sum in Equation 3 have the same asymptotic growth rate, the sums for σ and τ must be taken over the same values of j from $\{0, 1, \ldots, k\}$. □

Proposition 3.4 allows us to classify permutations based on subsequence containment, *i.e.*, to determine the equivalence classes with respect to Definition 3.1, by simply determining which permutations have the same j-sets. Table 1 lists the 2 possible j-sets for permutations in $\mathcal{S}_3$, the number of permutations which have each of those j-sets, and the number of n-involutions which contain the permutations from these classes as subsequences for $3 \leq n \leq 10$. (Recall that $I_n(\tau)$ denotes the number of n-involutions that contain τ as a subsequence.) Table 2 does the same for permutations in $\mathcal{S}_4$; in this case, there are 4 possible j-sets and the number of containing n-involutions is given for $4 \leq n \leq 10$.

Tables 3–5 present similar data for permutations in $\mathcal{S}_5$, $\mathcal{S}_6$, and $\mathcal{S}_7$. Note that for $\mathcal{S}_7$, not every possible j-set is realized as $\mathcal{J}(\tau)$ for some τ; no $\tau \in \mathcal{S}_7$ has $\mathcal{J}(\tau)$ equal to either $\{0,1,2,5,7\}$ or $\{0,1,2,4,5,7\}$. Additionally, no permutation $\tau \in \mathcal{S}_8$ has

$$\mathcal{J}(\tau) \in \{\{0,1,2,5,7\},\{0,1,2,6,8\},\{0,1,2,3,6,8\},\{0,1,2,4,5,7\},\\ \{0,1,2,5,6,8\},\{0,1,2,5,7,8\},\{0,1,2,3,5,6,8\},\{0,1,2,4,5,7,8\}\}.$$

This suggests that it may be interesting to determine how many j-sets actually occur.

Question. What is the sequence

$$\{|\mathcal{J}(\mathcal{S}_k)|\}_{k\geq 3} = 2, 4, 8, 16, 30, 56, 102, \ldots?$$

I.e., for $k \geq 3$, how may of the 2^{k-2} possible j-sets are actually realized by some permutation in $\mathcal{S}_k$?

A more general question is the enumeration of the k-permutations which have a particular j-set (as opposed to simply determining when this count is nonzero).

Question. Given a set E, $\{0,1,2\} \subseteq E \subseteq \{0,1,\ldots,k\}$, how many k-permutations τ have $\mathcal{J}(\tau) = E$?

$\mathcal{J}(\tau)$	$\lvert\{\tau\}\rvert$	$I_5(\tau)$	$I_6(\tau)$	$I_7(\tau)$	$I_8(\tau)$	$I_9(\tau)$	$I_{10}(\tau)$
$\{0,1,2\}$	26	0	0	0	1	5	26
$\{0,1,2,3\}$	36	0	0	1	4	17	66
$\{0,1,2,4\}$	8	0	1	2	7	21	76
$\{0,1,2,5\}$	4	1	1	2	5	15	52
$\{0,1,2,3,4\}$	24	0	1	3	10	33	116
$\{0,1,2,3,5\}$	4	1	1	3	8	27	92
$\{0,1,2,4,5\}$	2	1	2	4	11	31	102
$\{0,1,2,3,4,5\}$	16	1	2	5	14	43	142

Table 3: Classifying $\mathcal{S}_5$ by subsequence containment by involutions.

$\mathcal{J}(\tau)$	$\lvert\{\tau\}\rvert$	$I_6(\tau)$	$I_7(\tau)$	$I_8(\tau)$	$I_9(\tau)$	$I_{10}(\tau)$
$\{0,1,2\}$	146	0	0	0	0	1
$\{0,1,2,3\}$	204	0	0	0	1	5
$\{0,1,2,4\}$	46	0	0	1	3	13
$\{0,1,2,5\}$	20	0	1	2	6	17
$\{0,1,2,6\}$	10	1	1	2	4	11
$\{0,1,2,3,4\}$	136	0	0	1	4	17
$\{0,1,2,3,5\}$	20	0	1	2	7	21
$\{0,1,2,3,6\}$	12	1	1	2	5	15
$\{0,1,2,4,5\}$	8	0	1	3	9	29
$\{0,1,2,4,6\}$	2	1	1	3	7	23
$\{0,1,2,5,6\}$	4	1	2	4	10	27
$\{0,1,2,3,4,5\}$	64	0	1	3	10	33
$\{0,1,2,3,4,6\}$	8	1	1	3	8	27
$\{0,1,2,3,5,6\}$	4	1	2	4	11	31
$\{0,1,2,4,5,6\}$	4	1	2	5	13	39
$\{0,1,2,3,4,5,6\}$	32	1	2	5	14	43

Table 4: Classifying $\mathcal{S}_6$ by subsequence containment by involutions.

As special cases of this question, we have the following question and proposition. These are of particular interest because they give the number of k-permutations τ which achieve (for large enough n) the smallest and largest values of $I_n(\tau)$.

Question. How may k-permutations have j-set equal to $\{0,1,2\}$?

Proposition 3.5. *The number of $\tau \in \mathcal{S}_k$ for which $\mathcal{J}(\tau) = \{0,1,\ldots,k\}$ equals 2^{k-1}.*

$\mathcal{J}(\tau)$	$\lvert\{\tau\}\rvert$	$I_7(\tau)$	$I_8(\tau)$	$I_9(\tau)$	$I_{10}(\tau)$
$\{0,1,2\}$	992	0	0	0	0
$\{0,1,2,3\}$	1396	0	0	0	0
$\{0,1,2,4\}$	316	0	0	0	1
$\{0,1,2,5\}$	140	0	0	1	3
$\{0,1,2,6\}$	60	0	1	2	6
$\{0,1,2,7\}$	30	1	1	2	4
$\{0,1,2,3,4\}$	928	0	0	0	1
$\{0,1,2,3,5\}$	136	0	0	1	3
$\{0,1,2,3,6\}$	72	0	1	2	6
$\{0,1,2,3,7\}$	32	1	1	2	4
$\{0,1,2,4,5\}$	56	0	0	1	4
$\{0,1,2,4,6\}$	12	0	1	2	7
$\{0,1,2,4,7\}$	6	1	1	2	5
$\{0,1,2,5,6\}$	20	0	1	3	9
$\{0,1,2,5,7\}$	0	1	1	3	7
$\{0,1,2,6,7\}$	10	1	2	4	10
$\{0,1,2,3,4,5\}$	432	0	0	1	4
$\{0,1,2,3,4,6\}$	48	0	1	2	7
$\{0,1,2,3,4,7\}$	24	1	1	2	5
$\{0,1,2,3,5,6\}$	20	0	1	3	9
$\{0,1,2,3,5,7\}$	4	1	1	3	7
$\{0,1,2,3,6,7\}$	12	1	2	4	10
$\{0,1,2,4,5,6\}$	20	0	1	3	10
$\{0,1,2,4,5,7\}$	0	1	1	3	8
$\{0,1,2,4,6,7\}$	2	1	2	4	11
$\{0,1,2,5,6,7\}$	8	1	2	5	13
$\{0,1,2,3,4,5,6\}$	160	0	1	3	10
$\{0,1,2,3,4,5,7\}$	16	1	1	3	8
$\{0,1,2,3,4,6,7\}$	8	1	2	4	11
$\{0,1,2,3,5,6,7\}$	8	1	2	5	13
$\{0,1,2,4,5,6,7\}$	8	1	2	5	14
$\{0,1,2,3,4,5,6,7\}$	64	1	2	5	14

Table 5: Classifying $\mathcal{S}_7$ by subsequence containment by involutions.

Before proving this proposition, we define two operations for extending a permutation $\tau \in \mathcal{S}_k$ to a permutation in $\mathcal{S}_{k+1}$ whose length k initial pattern equals τ. It will be helpful to use the *graph of a permutation* in doing so; the graph of $\tau \in \mathcal{S}_k$ is a $k \times k$ grid, which we will co-ordinatize from the bottom left corner, with dots in exactly the boxes $\{(i, \tau_i)\}_{i \in [k]}$.

Definition 3.6. Given $\tau \in \mathcal{S}_k$, $F(\tau)$ is the permutation in $\mathcal{S}_{k+1}$ that fixes $k+1$ and permutes $[k]$ as τ does. The graph of $F(\tau)$ is obtained from the graph of τ by adding a dot in the box $(k+1, k+1)$ (as well as the appropriate additional empty boxes). The left part of Figure 1 illustrates the construction of the graph of $F(\tau)$ from the graph of τ; the white area indicates the boxes forming the graph of τ while the shaded areas and dot are added to obtain the graph of $F(\tau)$.

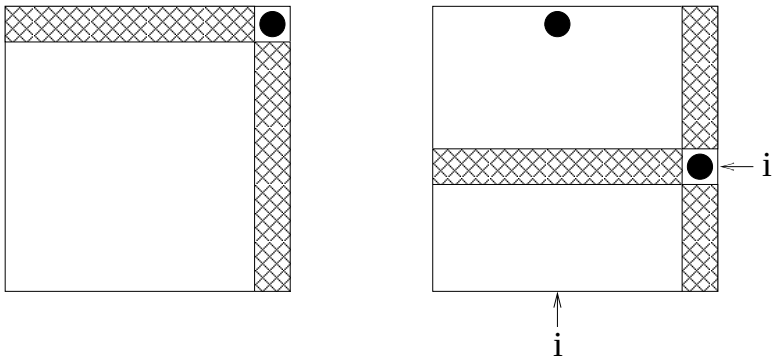

Figure 1: Constructing the graphs of $F(\tau)$, left, and $G(\tau)$, right, from the graph of τ.

Remark 3.7. By the construction of $F(\tau)$, we see that for $\tau \in \mathcal{S}_k$:

1. the length k initial pattern of $F(\tau)$ equals τ and
2. $F(\tau)$ is an involution iff τ is an involution.

Definition 3.8. Given $\tau \in \mathcal{S}_k$ with $\tau_i = k$, $G(\tau)$ is the permutation obtained by adding 1 to every value in τ that is at least i and then appending the value i. The graph of $G(\tau)$ is obtained from the graph of τ by inserting an empty row at height i (moving the rows originally at heights $i, \ldots, k$ to be at heights $i+1, \ldots, k+1$) and then adding a dot in the box $(k+1, i)$ (as well as the appropriate additional empty boxes). The right part of Figure 1 illustrates the construction of the graph of $G(\tau)$ from the graph of τ when $\tau_i = k$ is the maximum value of $\tau \in \mathcal{S}_k$. The white area again indicates the boxes forming the graph of τ, now split into the bottom $i-1$ rows and the remaining $k-i+1$ rows, while the shaded boxes and dot are added to obtain the graph of $G(\tau)$.

Remark 3.9. By the construction of $G(\tau)$, we see that for $\tau \in \mathcal{S}_k$:

1. the length k initial pattern of $G(\tau)$ equals τ and

2. $G(\tau)$ is an involution iff the length $k-1$ pattern of τ^{-1} is an involution (equivalently, iff the permutation obtained by deleting the largest value k from τ is an involution).

Note that if an involution in $\mathcal{S}_{k+1}$ has length k initial pattern equal to $\tau \in \mathcal{S}_k$, then that involution must equal $F(\tau)$ (if it fixes $k+1$) or $G(\tau)$ (if it does not).

Proof of proposition 3.5. This is true for $k = 1, 2$; assume that the proposition holds for k and pick $\tau \in \mathcal{S}_k$ such that $\mathcal{J}(\tau) = \{0, \ldots, k\}$. As τ is an involution, $F(\tau)$ is as well and $\mathcal{J}(F(\tau)) = \{0, \ldots, k, k+1\}$. The length $k-1$ initial pattern of $\tau^{-1} = \tau$ is also an involution by assumption, so $G(\tau)$ is an involution and $\mathcal{J}(G(\tau)) = \{0, \ldots, k, k+1\}$. Because $F(\tau) \neq G(\tau)$ and any involution in $\mathcal{S}_{k+1}$ whose length k initial pattern equals τ must be either $F(\tau)$ or $G(\tau)$, the proposition holds for $k+1$. □

Definition 3.10. A permutation $\tau \in \mathcal{S}_k$ is said to be *layered* if, for some composition $(a_1, a_2, \ldots, a_j)$ of k (with each $a_i \geq 1$), τ consists of the first a_1 positive integers arranged in decreasing order, then the next a_2 positive integers in decreasing order, *etc.* For example, $32146578 \in \mathcal{S}_8$ is the layered permutation corresponding to the composition $(3, 1, 2, 1, 1)$ of 8.

Proposition 3.11. *A permutation $\tau \in \mathcal{S}_k$ is a layered permutation if and only if $\mathcal{J}(\tau) = \{0, 1, \ldots, k\}$,* i.e.*, if and only if the pattern of $\tau_1 \ldots \tau_j$ is an involution for every $j \in [k]$.*

Proof. If $\sigma \in \mathcal{S}_{k'}$ is layered, let the last layer of σ be $\sigma_i = k, \ldots, \sigma_k = i$. $G(\sigma)$ agrees with σ in its first $i-1$ values; these are followed in $G(\sigma)$ by $(k+1)\ldots(i+1)i$ and so $G(\sigma)$ is also layered. It is clear that if $\sigma \in \mathcal{S}_{k'}$ is layered then $F(\sigma)$ is also layered. The proof of Proposition 3.5 shows that every $\tau \in \mathcal{S}_k$ with $\mathcal{J}(\tau) = \{0, \ldots, k\}$ may be constructed by starting with a permutation in $\mathcal{S}_2$, both of which are layered, and then applying some sequence of operations $F(\cdot)$ and $G(\cdot)$, so all 2^{k-1} of these $\tau \in \mathcal{S}_k$ are layered. There are 2^{k-1} layered permutations in $\mathcal{S}_k$, a well-known result that follows from the natural bijection between layered permutations and compositions of k. □

Lemma 3.12. *Given a layered permutation $\tau \in \mathcal{S}_k$, $k \geq 2$, there is a unique involution $\tilde{\tau} \in \mathcal{S}_{k+2}$ whose length k initial pattern equals τ and whose length $k+1$ initial pattern is not an involution.*

Proof. If such an involution $\tilde{\tau} \in \mathcal{S}_{k+2}$ exists, then it must equal $G(\sigma)$, where σ is the length $k+1$ initial pattern of $\tilde{\tau}$; $\tilde{\tau}$ cannot equal $F(\sigma)$ as $F(\cdot)$ preserves the (non-)involutive nature of permutations. If $\alpha \in \mathcal{S}_k$ is obtained from σ by deleting its largest value, then $G(\sigma)$ is an involution iff α is. Because τ is layered, the length $k-1$ initial pattern of α must be layered; in order for α to be an involution, α must be obtained by applying $F(\cdot)$ or $G(\cdot)$ to this layered permutation.

Note that the largest value of σ cannot be σ_{k+1}, otherwise σ would be layered (because τ is). If the largest value of σ is σ_k (*i.e.*, if the last layer of τ has size 1), then α cannot be the result of applying $F(\cdot)$ to a layered permutation as σ would then be a layered permutation. In this case the length $k-1$ initial pattern of α is $\tau_1 \ldots \tau_{k-1}$; we let

$\alpha = G(\tau_1 \dots \tau_{k-1})$ and obtain σ by inserting $k+1$ just before the last value of α. $\tilde{\tau} = G(\sigma)$ is then as required.

If the largest value of σ is not σ_k (*i.e.*, if the last layer of τ has size greater than 1), then α cannot be the result of applying $G(\cdot)$ to a layered permutation as σ would then be a layered permutation. In this case, if m is the index of the largest value of τ then the length $k-1$ initial pattern of α equals the pattern of $\tau_1 \dots \tau_{m-1}\tau_{m+1} \dots \tau_k$; we may construct α by applying $F(\cdot)$ to this pattern and then obtain σ by inserting $k+1$ into α just before α_m. $\tilde{\tau} = G(\sigma)$ is then as required.

As $\tilde{\tau}$ must equal $G(\sigma)$ and there is a unique way to extend any layered $\tau \in \mathcal{S}_k$ to σ such that $\tilde{\tau}$ has the stated properties (the manner of doing so depending upon whether the last layer of τ has size 1), the lemma is proved. □

Lemma 3.13. *Given $\tau \in \mathcal{S}_k$ with $\mathcal{J}(\tau) = \{0, 1, \dots, k\}$, there is a unique $\overline{\tau} \in \mathcal{S}_{k+3}$ with $\mathcal{J}(\overline{\tau}) = \{0, 1, \dots, k, k+2, k+3\}$ and length k initial pattern equal to τ.*

Proof. Construct the $\tilde{\tau} \in \mathcal{S}_{k+2}$ as in Lemma 3.12; this is the only possible length $k+2$ initial pattern of $\overline{\tau}$. As $\tilde{\tau}$ is an involution, $F(\tilde{\tau}) \in \mathcal{S}_{k+3}$ is as well. By construction, the length $k+1$ initial pattern of $\tilde{\tau}^{-1} = \tilde{\tau}$ is not an involution, so $G(\tilde{\tau})$ is not an involution. $G(\tilde{\tau})$ is the only other way that $\tilde{\tau}$ could have been extended to an involution in $\mathcal{S}_{k+3}$, so $\overline{\tau} = F(\tilde{\tau})$ is unique as claimed. □

Tables 1–5 suggest the following proposition, which we may prove using these lemmas.

Proposition 3.14. *Assume $\{0, 1, 2, k\} \subseteq E \subset \{0, 1, \dots, k\}$ and $|E| = k \geq 5$. Then*

$$|j^{-1}(E) \bigcap \mathcal{S}_k| = \begin{cases} 2^{k-3} & k-1 \notin E \\ 2^{k-4} & k-1 \in E \end{cases}.$$

Proof. Let $\{i\} = \{0, 1, \dots, k\} \setminus E$. The case $i = k-1$ is covered by Proposition 3.5 and Lemma 3.12, while the case $i = k-2$ is covered by these and Lemma 3.13.

For $i \leq k-3$, construct the unique $\overline{\tau} \in \mathcal{S}_{i+2}$ such that $\mathcal{J}(\overline{\tau}) = \{0, \dots, i-1, i+1, i+2\}$. Because $\overline{\tau}$ is an involution, the length $i+1$ initial pattern of $\overline{\tau}^{-1}$ equals the length $i+1$ initial pattern of $\overline{\tau}$ itself; by assumption, this is an involution. Thus both $F(\overline{\tau})$ and $G(\overline{\tau})$ are involutions in $\mathcal{S}_{i+3}$ whose j-sets equal $\{0, \dots, i-1, i+1, i+2, i+3\}$, and the proposition holds for $k = i+3$. The same argument shows that these permutations may be extended by any sequence of $F(\cdot)$ and $G(\cdot)$ operations until length k involutions are produced with the desired properties, and the theorem is proved. □

Finally, review of Tables 1–5 also suggests the following question.

Question. For $k > 4$ and any set $E \neq \{0, 1, 2, 3\}$, is the number of $\sigma \in \mathcal{S}_k$ for which $\mathcal{J}(\sigma) = E$ always strictly less than the number of $\sigma \in \mathcal{S}_k$ for which $\mathcal{J}(\sigma) = \{0, 1, 2, 3\}$? Whether or not this holds, are there any other statements that can be made about the frequency of particular sets appearing as $\mathcal{J}(\sigma)$?

3.2 Tableau containment

We might ask similar questions for certain classes of permutations instead of for all permutations in $\mathcal{S}_n$. As an example, fix a (standard) Young tableau T of size k and consider the set $\mathcal{T}$ of k-permutations which correspond under the Robinson–Schensted-Knuth (RSK) algorithm to (T, Q) for some tableau Q of the same shape as T. Lemma 3.15 counts the permutations $\sigma \in \mathcal{T}$ such that $\mathcal{J}(\sigma)$ contains a specified value; we first briefly recall the RSK algorithm and some of its properties.

The RSK algorithm bijectively associates to every n-permutation π a pair (P, Q) of standard Young tableaux, where P and Q have as their common shape some partition of n. Following [Sta99], we refer to P as the *insertion tableau* and to Q as the *recording tableau*. In the case that π is a general word of n distinct letters, the RSK algorithm produces a pair (P, Q), where P is a tableau whose entries are the letters appearing in π and Q is a standard Young tableau of the same shape as P. Applying the RSK algorithm to the pattern of π yields (P', Q'), where P' is the order preserving relabeling of P with the elements of $[n]$ and $Q' = Q$. Finally, for a permutation $\pi \in \mathcal{S}_n$ with corresponding pair of tableaux (P, Q), $P = Q$ iff π is an involution.

Lemma 3.15. *Fix k and $\lambda \vdash k$, and let T be a standard Young tableau of shape λ. Let $\mathcal{T}$ be the set of permutations which correspond to (T, Q) for some Q, let $a_0 = f^\lambda$, and for $1 \le j \le k$ let a_j be the number of $\tau \in \mathcal{T}$ for which $j \in \mathcal{J}(\tau)$. Then (with $f^\emptyset = 1$)*

$$a_j = \sum_{\mu \vdash j} f^{\lambda/\mu} \quad (j = 0, 1, \ldots, k), \tag{9}$$

where $f^{\lambda/\mu}$ is the number of standard Young tableaux of skew shape λ/μ.

Proof. This is immediate for $j = 0$. For $j \in [k]$ and a tableau Q of size k, let Q_j denote the subtableau of Q formed by the elements of $[j]$. Observe that for $j \in [k]$, if $\tau = \tau_1 \ldots \tau_k$ is a k-permutation which corresponds under RSK to the pair (P, Q) of tableaux, then the following three conditions are equivalent:

(i) The length j initial pattern of τ is a j-involution.

(ii) The length j initial pattern of τ corresponds under RSK to (Q_j, Q_j).

(iii) $\tau_1 \ldots \tau_j$ corresponds to (P', Q') under RSK, where $Q' = Q_j$ is the order preserving relabeling of P' with the elements of $[j]$.

Each of the f^λ permutations in $\mathcal{T}$ corresponds under RSK to (T, Q) for some recording tableau Q of shape λ. There are $\sum_{\mu \vdash j} f^{\lambda/\mu}$ different ways that a recording tableau can contain $\{j+1, \ldots, k\}$. For a fixed such arrangement ST, a skew tableau, let $\mathcal{ST}$ denote the set of (recording) tableaux which contain $\{j+1, \ldots, k\}$ in the arrangement ST. If μ is the partition for which the shape of ST is λ/μ, then there are f^μ tableaux in $\mathcal{ST}$. For a fixed set $\mathcal{ST}$ and $Q \in \mathcal{ST}$, the removal of $k - j$ elements from T by applying the inverse RSK algorithm to (T, Q) yields a pair (T'_{ST}, Q') of tableaux of shape μ, where

$Q' = Q_j$ and T'_{ST} depends only on ST. Of the f^μ tableaux Q' so obtained by letting Q range over ST, exactly one is the order preserving relabeling of T'_{ST} with the elements of $[j]$. By the observation above, this choice of Q', together with the arrangement ST, gives a recording tableaux $Q \in \mathcal{ST}$ such that (T, Q) corresponds under RSK to a permutation whose length j initial pattern is a j-involution; this is the only choice of $Q \in \mathcal{ST}$ for which this is true. Considering all possible arrangements ST of $\{j+1, \ldots, k\}$, we see that there are $\sum_{\mu \vdash j} f^{\lambda/\mu}$ k-permutations in $\mathcal{T}$ whose length j initial pattern in a j-involution. □

Remark 3.16. The values of a_j in the lemma depend only on the shape λ of T. However, the number of permutations whose length j and $j' \neq j$ initial patterns are both involutions need not be the same for different tableaux of the same shape.

As a corollary, this allows us to count the standard Young tableaux which contain a given tableau. Theorem 3.17 was first proved by Sagan and Stanley, following from Corollary 3.5 in [SS90]. The position of certain elements in a tableau was the motivating question in the work of McKay, Morse, and Wilf on quasirandomness [MMW02]. More recently, Stanley revisited the subtableau question and developed related asymptotic work [Sta03]. Finally, Grabiner has used a random-walk approach to study the asymptotics for the problem and a generalization to up-down tableaux [Gra04].

Theorem 3.17 (Sagan and Stanley [SS90]). *Let T be a standard Young tableau of shape $\lambda \vdash k$. Then the number of tableaux of size $n \geq k$ which contain T as a subtableau equals*

$$\sum_{j=0}^{k} \sum_{\mu \vdash j} f^{\lambda/\mu} \binom{n-k}{k-j} t_{n-2k+j}. \tag{10}$$

Proof. The number of tableaux containing T equals the sum of Equation 3 over all permutations τ which correspond to (T, Q) for some Q. For $j = 0, 1, \ldots, k$, there are exactly $a_j = \sum_{\mu \vdash j} f^{\lambda/\mu}$ permutations $\tau \in \mathcal{T}$ for which the range of the sum in Equation 3 includes j, from which the theorem follows. □

Acknowledgments

We are grateful to Herb Wilf for suggesting this problem to us and for helpful discussions, the referee for information about related work and other useful suggestions, Alkes Price for comments about known results on layered patterns, and Andre Scedrov for arranging research support.

References

[Gra04] David Grabiner, *Asymptotics for the distributions of subtableaux in Young and up-down tableaux*, Preprint, 2004.

[Jag03] Aaron D. Jaggard, *Involutions in the symmetric group: Containment properties and parallels to general permutations*, Ph.D. thesis, University of Pennsylvania, 2003.

[MMW02] Brendan D. McKay, Jennifer Morse, and Herbert S. Wilf, *The distributions of the entries of Young tableaux*, J. Combin. Theory Ser. A **97** (2002), no. 1, 117–128.

[SS90] Bruce E. Sagan and Richard P. Stanley, *Robinson-Schensted algorithms for skew tableaux*, J. Combin. Theory Ser. A **55** (1990), no. 2, 161–193.

[Sta99] Richard P. Stanley, *Enumerative combinatorics. Vol. 2*, Cambridge University Press, Cambridge, 1999.

[Sta03] ———, *On the enumeration of skew Young tableaux*, Adv. in Appl. Math. **30** (2003), 283–294.