Surprising Symmetries of Objects Counted by Catalan numbers

Miklós Bóna

Department of Mathematics
University of Florida
Gainesville FL 32611-8105

bona@ufl.edu

June 12, 2012

Pattern Containment and Avoidance

Definition

We say that the permutation $p = p_1 p_2 \dots p_n \in S_n$ *contains a q-pattern* if and only if there is a subsequence $p_{i_1} p_{i_2} \dots p_{i_k}$ of p whose elements are in the same relative order as those in q, that is,

$$p_{i_t} < p_{i_u} \text{ if and only if } q_t < q_u$$

whenever $1 \leq t, u \leq k$.

If p does not contain q, then we say that p *avoids* q.

Example

The permutation $p = 3267415$ contains the pattern $q = 231$ (consider the first, fourth and sixth entries).

On the other hand, p avoids 4321 since p does not contain a decreasing subsequence of length 4.

A family of questions

Joshua Cooper has raised the following interesting family of questions.

Question

Let r be a given permutation pattern. What can be said about the average number of occurrences of q in a randomly selected r-avoiding permutation of a given length?

Equivalently, can we determine the total number $S_{n,r}(q)$ *of all q-patterns in all r-avoiding permutations of length n?*

Earlier Results

Present author found formulae for the generating functions of the sequence $S_{132,n}(q)$ for the cases of monotone q, that is, for $q = 12 \cdots k$ and $q = k(k-1) \cdots 1$, for any k.

He also proved that if n is large enough, then for any fixed k, among all patterns q of length k, it is the monotone decreasing pattern that maximizes $S_{132,n}(q)$ and it is the monotone increasing pattern that minimizes $S_{132,n}(q)$.

Surprising Symmetries

We first noticed the following surprising fact, that is not overly difficult to prove using generating functions.

For all positive integers n, the equalities

$$S_{132,n}(231) = S_{132,n}(312) = S_{132,n}(213) \tag{1}$$

hold.

The first equality is trivial, since taking the inverse of a 132-avoiding permutation keeps that permutation 132-avoiding, and turns 231-patterns into 312-patterns.

However, the second equality is non-trivial. (The reverse or complement of a 132-avoiding permutation is not necessarily 132-avoiding.)

In particular, if $a(p)$ denotes the number of 213-copies in p, and $b(p)$ denotes the number of 231-copies in p, then the statistics $a(p)$ and $b(p)$ are *not* equidistributed over the set of all 132-avoiding permutations of length n, but their average values (or cumulative values) are equal over that set.

We will first show a bijective proof of this fact, then we show a far reaching generalization of that proof, which deals with much longer patterns, for which generating functions would have been unlikely (or at least very unpleasant) to work.

Binary Plane Trees

In our proof, we will identify a 132-avoiding permutation p with its *binary plane tree* $T(p)$ using a very well-known bijection.

The tree $T(p)$ will be a binary plane tree, that is, a rooted unlabeled tree in which each vertex has at most two children, and each child is a left child or a right child of its parent, even if it is the only child of its parent.

The root of $T(p)$ corresponds to the entry n of p, the left subtree of the root corresponds to the string of entries of p on the left of n, and the right subtree of the root corresponds to the string of entries of p on the right of n. Both subtrees are constructed recursively, by the same rule.

Note that since p is 132-avoiding, the position of the entry n of p determines the set of entries that are on the left (resp. on the right) of n.

In fact, if n is in the ith position, the set of entries on the left of n must be $\{n-i+1, n-i+2, \cdots, n-1\}$, and the set of entries on the right of n must be $\{1, 2, \cdots, n-i\}$.
See Figure 11 for an illustration.

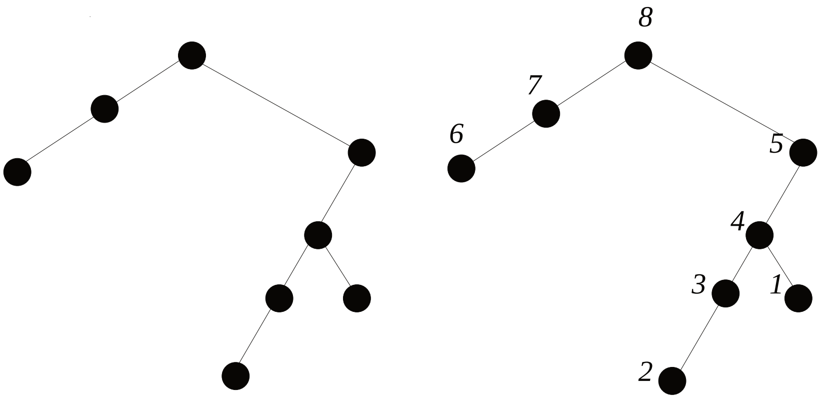

Figure: The tree $T(p)$ for $p = 67823415$, and the entries of p associated to the vertices of $T(p)$.

Let p be a 132-avoiding n-permutation, and let Q be an occurrence of the pattern 213 in p. Let Q_2, Q_1, Q_3 be the three vertices of $T(p)$ that correspond to Q, going left to right. Let us color these three entries black. There are then two possibilities.

1. Either Q_1 is a right descendant of Q_2 and Q_2 is a left descendant of Q_3, or
2. there exists a lowest left descendant Q_x of Q_3 so that Q_2 is a left descendant of Q_x and Q_1 is a right descendant of Q_x.

Let A_n be the set of all binary plane trees on n vertices in which three vertices forming a 213-pattern are colored black. Let B_n be the set of all binary plane trees on n vertices in which three vertices forming a 231-pattern are colored black.

Now we define a map $f : A_n \rightarrow B_n$. We will then prove that f is a bijection. The map f will be defined differently in the two cases described above.

Case 1. If $T \in A_n$ is in the first case, then let $f(T)$ be the pair obtained by interchanging the right subtree of Q_2 and the right subtree of Q_3. Keep all three black vertices Q_i black, even as Q_1 gets moved. See Figure 14 for an illustration.

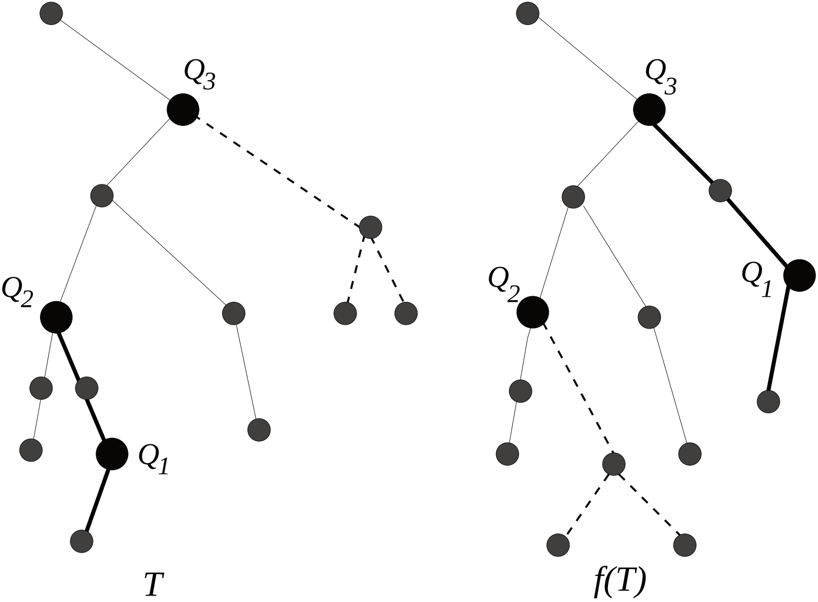

Figure: Interchanging the right subtrees of Q_2 and Q_3.

Note that in $f(T)$, in the set of black vertices, there is one that is an ancestor of the other two, namely Q_3.

Case 2. If $T \in A_n$ is in the second case, then let $f(T)$ be the tree obtained by interchanging the right subtrees of the vertices Q_x and Q_3, and coloring Q_2, Q_x and Q_1 black. See Figure 16 for an illustration.

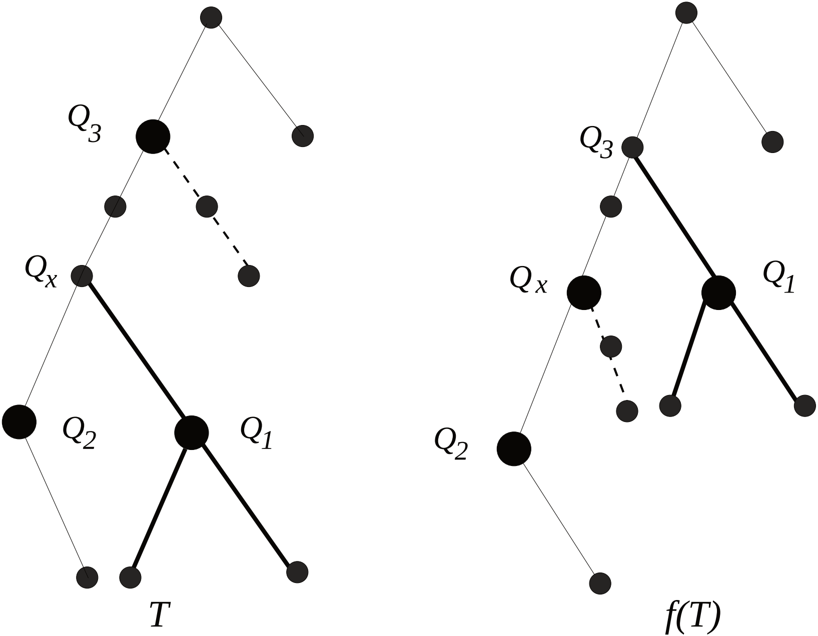

Q_3
Q_x
Q_2
Q_1
T
Q_3
Q_x
Q_1
Q_2
$f(T)$

Note that in $f(T)$, there is no black vertex that is an ancestor of the other two black vertices. Also note that in $f(T)$, the lowest common ancestor of Q_x and Q_1 is Q_3.

It is a direct consequence of our definitions that if $T \in A_n$, then $f(T) = B_n$.

Theorem

The map $f : A_n \to B_n$ defined above is a bijection.

Proof:

Let $U \in B_n$. We show that there is exactly one $T \in A_n$ so that $f(T) = U$ holds. This will show that f has an inverse.

By definition, three nodes of U are colored black, and the entries of the permutation corresponding to U form a 231-pattern.

Let K_2, K_3, and K_1 denote these three vertices, from left to right. There are two possibilities for the location of the K_i relative to each other. We show that in both cases, U has a unique preimage under f, essentially because swapping two subtrees is an involution.

1. If K_3 is an ancestor of both other black vertices, then $f(T) = U$ implies that T belongs to Case 1. In this case, the unique $T \in A_n$ satisfying $f(T) = U$ is obtained by swapping the right subtrees of K_3 and K_2, and keeping all three black vertices black, even if K_1 got moved.

2. If K_3 is not an ancestor of both other black vertices and then $f(T) = U$ implies that T belongs to Case 2. In this case, let K_x be the smallest common ancestor of U_3 and U_1. Then the unique $T \in A_n$ satisfying $f(T) = U$ is obtained by swapping the right subtrees of K_3 and K_x, and coloring K_x black instead of K_3, while keeping K_1 and K_2 black.

This completes the proof.

A Generalization

Definition

Let q be a pattern of length k and let t be a pattern of length m. Then $q \oplus t$ is the pattern of length $k+m$ defined by

$$(q \oplus t)_i = \begin{cases} q_i \text{ if } i \leq k, \\ \\ t_{i-k} + k \text{ if } i > k. \end{cases}$$

In other words, $q \oplus t$ is the concatenation of q and t so that all entries of t are increased by the size of q.

Example

If $q = 3142$ and $t = 132$, then $q \oplus t = 3142576$.

Definition

Let q be a pattern of length k and let t be a pattern of length m. Then $q \ominus t$ is the pattern of length $k + m$ defined by

$$(q \ominus t)_i = \begin{cases} q_i + m \text{ if } i \leq k, \\ \\ t_{i-k} \text{ if } i > k. \end{cases}$$

In other words, $q \ominus t$ is the concatenation of q and t so that all entries of q are increased by the size of t.

Example

If $q = 3142$ and $t = 132$, then $q \ominus t = 6475132$.

Theorem

Let q and t be any non-empty patterns that end in their largest entry. Let i_u denote the increasing pattern $12 \cdots u$. Then for all positive integers n, we have

$$S_{n,132}((q \ominus t) \oplus i_u) = S_{n,132}((q \oplus i_u) \ominus t),$$

where 1 denotes the pattern consisting of one entry.

Example

If $q = 3124$, $t = 213$, and $u = 2$, then Theorem 8 says that

$$S_{n,132}(645721389) = S_{n,132}(645789213).$$

Further direction of research: are there any non-trivial identities in which the two sides have *different* culling patterns?

A recent result of Cheyne Homberger shows the following.

Theorem

For all positive integers n, the equality

$$S_{n,132}(231) = S_{n,123}(231).$$