ELSEVIER

Discrete Mathematics 204 (1999) 119–128

DISCRETE MATHEMATICS

www.elsevier.com/locate/disc

Forbidden subsequences and Chebyshev polynomials

Timothy Chow[a,*], Julian West[b]

[a] Tellabs Research Center, 1 Kendall Square, Cambridge, MA 02139, USA
[b] Department of Mathematics and Statistics, University of Victoria, BC, Canada

Received 26 June 1997; revised 10 March 1998; accepted 3 August 1998

Abstract

In (West, Discrete Math. 157 (1996) 363–374) it was shown using transfer matrices that the number $|S_n(123;3214)|$ of permutations avoiding the patterns 123 and 3214 is the Fibonacci number F_{2n} (as are also $|S_n(213;1234)|$ and $|S_n(213;4123)|$). We now find the transfer matrix for $|S_n(123;r,r-1,\dots,2,1,r+1)|$, $|S_n(213;1,2,\dots,r,r+1)|$, and $|S_n(213;r+1,1,2,\dots,r)|$, determine its characteristic polynomial in terms of the Chebyshev polynomials, and go on to determine the generating function as a quotient of modified Chebyshev polynomials. This leads to an asymptotic result for each r which collapses to the exact results 2^n when $r=2$ and F_{2n} when $r=3$ and to the Catalan number c_n as $r\to\infty$. We observe that our generating function also enumerates certain lattice paths, plane trees, and directed animals, giving hope that these areas of combinatorics can be applied to enumerating permutations with excluded subsequences.

Keywords: Permutations; Catalan numbers; Lattice paths; Plane trees; Convex directed animals

1. Introduction

Let S_n denote the symmetric group on n letters. There is a large literature on enumerating permutations with excluded subsequences, which are defined as follows.

Definition 1.1. Let $\tau\in S_k$. A permutation $\pi\in S_n$ is *τ-avoiding* if there is no sequence $i_1,i_2,\dots,i_k$ of integers such that

$$1\leqslant i_{\tau(1)}<i_{\tau(2)}<\cdots<i_{\tau(k)}\leqslant n \quad\text{and}\quad \pi(i_1)<\pi(i_2)<\cdots<\pi(i_k).$$

The subsequence $\{\pi(i_{\tau(j)})\}_{j=1}^{k}$ is said to have *type* τ. We write $S_n(\tau)$ for the set of τ-avoiding permutations of length n.

* Corresponding author. E-mail: tchow@alum.mit.edu.

0012-365X/99/$ – see front matter
PII: S0012-365X(98)00384-7

For example, the permutation $\pi = 3461752$ is not 3421-avoiding because we can take $i_1 = 7$, $i_2 = 6$, $i_3 = 3$, and $i_4 = 5$. Informally, the subsequence of π consisting of the 3rd, 5th, 6th, and 7th numbers — namely, 6752 — has the same 'shape' as 3421.

One tool in this subject is the concept of a generating tree [1,2,4].

Definition 1.2. A *generating tree* is a rooted labelled tree with the property that if v_1 and v_2 are any two nodes with the same label and ℓ is any label, then v_1 and v_2 have exactly the same number of children with the label ℓ. To specify a generating tree it therefore suffices to specify

(1) the label of the *root*, and

(2) a set of *succession rules* explaining how to derive from the label of a parent the labels of all of its children.

Example 1.3 (*The complete binary tree*).
Root: (2),
Rule: (2) $\rightarrow$ (2)(2).

Our notation for the succession rule simply means that any node with the label 2 has two children, each of which also has the label 2. We are generally interested in Σ_n, the number of nodes on level n of the tree, and sometimes in $(label)_n$, the number of nodes on level n labelled $(label)$. In this example, $\Sigma_n = (2)_n = 2^{n-1}$, if the root is considered to be level 1.

Example 1.4 (*The Fibonacci tree*).
Root: (1),
Rules: (1) $\rightarrow$ (2) (2) $\rightarrow$ (1)(2).

We could alternatively use (non-breeding pair) for (1) and (breeding pair) for (2), but the numeric labels provide a convenient record of the number of children of each node.

In this case, $\Sigma_n = (2)_{n+1}$ and

$$\begin{pmatrix} (1)_n \\ (2)_n \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}^{n-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} F_{n-2} \\ F_{n-1} \end{pmatrix}. \tag{1.1}$$

Here F_n is the nth Fibonacci number. The 2×2 matrix in this equation is called the *transfer matrix*, and $\binom{1}{0}$ is the vector representing the labels present on the root level. It should be clear how the transfer matrix is derived from the succession rules.

For a longer discussion with more examples, see [5].

The connection between τ-avoiding permutations and generating trees comes from an idea in [1]. Given τ, define a rooted tree as follows. The nodes on level n are precisely the elements of $S_n(\tau)$. The parent of a permutation $\pi = p_1, p_2, \ldots, p_n$ is the unique permutation $p_1, \ldots, p_{j-1}, p_{j+1}, \ldots, p_n$ such that $p_j = n$. We call the resulting tree $T(\tau)$. We abbreviate $S_n(\tau_1) \cap S_n(\tau_2)$ by $S_n(\tau_1;\tau_2)$ and its corresponding tree by $T(\tau_1;\tau_2)$.

Now to make $T(\tau)$ a *generating* tree, we must assign labels to the nodes. In general this can be done in many different ways, but rather than giving a universal construction here, we shall describe the labelling rules case by case during our analysis in the next section.

It is worth making one general comment about the construction of the tree. Although the tree has been defined by explaining how to find the parent of a given child, in practice we proceed by finding all the children of a given parent. Thus instead of *deleting* the largest element, n, from wherever it happens to be, we construct the tree from the root down by *inserting* n in every position where it does not create a τ-subsequence. We emphasize that the element being inserted is always the *largest* element of the resulting child permutation and we rely on this fact heavily.

Our main result enumerates classes of permutations avoiding certain pairs τ_1 and τ_2 of permutations, where τ_1 has length three and τ_2 has arbitrary length. In Section 2, we use combinatorial arguments to describe the structure of the generating tree for each class. In Section 3, we use algebraic techniques to extract information from the resulting transfer matrices. In Section 4, we place our new results in the context of previous work. Finally, in Section 5, we sketch some connections with lattice paths, plane trees, and directed animals. We feel that these connections deserve further exploration and may give new insight into enumerating permutations with excluded subsequences.

2. Combinatorics

We first study $S_n(123)$ and $S_n(213)$, reproducing results from [4].

Definition 2.1. The *i.i.s.* of a permutation is its initial increasing subsequence: take numbers from the beginning of the permutation until a descent is encountered. The *i.d.s.* of permutation is its initial decreasing subsequence.

Consider a permutation $\pi \in S_{n-1}(123)$, $\pi = p_1,\ldots,p_{n-1}$. Its potential children in the generating tree $T(123)$ are

$$\pi^j = p_1, p_2, \ldots, p_{j-1}, n, p_j, \ldots, p_{n-1},$$

as j ranges from 1 to n. Clearly, $\pi^j \notin S_n(123)$ exactly when there exist $1 \leqslant i_1 < i_2 < j$ such that $p_{i_1} < p_{i_2}$, for if this happens then p_{i_1}, p_{i_2}, n is a forbidden subsequence of type 123. Conversely, $\pi^j \in S_n(123)$ iff $p_1,\ldots,p_{j-1}$ is a decreasing subsequence. The permutation π therefore has exactly k children in $T(123)$ if its longest i.d.s. has length $k-1$. Now looking at its children, π^1 has an i.d.s. of length k while π^j $(2 \leqslant j \leqslant k)$ has an i.d.s. of length $j-1$. What this says, in summary, about the succession rules for $T(123)$ is

Rule: $(k) \to (k+1)(2)\cdots(k)$.

The transfer matrix is therefore the infinite matrix A_∞ in the equation

$$\begin{pmatrix} (2)_n \\ (3)_n \\ (4)_n \\ (5)_n \\ \vdots \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & \cdots \\ 1 & 1 & 1 & \cdots \\ 0 & 1 & 1 & \cdots \\ 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \end{pmatrix}^{n-1} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{pmatrix}. \tag{2.1}$$

We refrain from an algebraic attack on this matrix as it is well known that in this instance $\Sigma_n = (2)_{n+1} = c_n = \binom{2n}{n}/(n+1)$, the nth *Catalan number*.

Turning to $S_n(213)$, we see a permutation will now have exactly k children in the tree $T(213)$ if its longest i.i.s. has length $k-1$. These children, $\pi^1, \ldots, \pi^k$, will therefore have i.i.s.'s of lengths $1, 2, \ldots, k$ respectively, so again

$$\text{Rule: } (k) \to (2)\cdots(k)(k+1) \tag{2.2}$$

and we obtain the same transfer matrix.

Our new results will be to modify these arguments to treat

$$T(123; r, \ldots, 1, r+1), \quad T(213; r+1, 1, \ldots, r) \quad \text{and} \quad T(213; 1, \ldots, r, r+1).$$

This was performed for $r=3$ in [5].

In the first example, $T(123; r, \ldots, 1, r+1)$, the first restriction, $\alpha = 123$, constrains (as above) the insertion of a new element, n, to occur within, or immediately after, the i.d.s. of the permutation $\pi \in S_{n-1}$. The second restriction, $\beta = r, \ldots, 1, r+1$, then assures that a new largest element cannot be placed as far as r positions into this i.d.s., otherwise $p_1, p_2, \ldots, p_r, n$ would be of type β. It follows that if the i.d.s. has length $k \geqslant r$, the number of children will be r rather than $k+1$, and their labels will be $(r)(2)(3)\cdots(r)$, arguing from the length of the resulting i.d.s.'s.

Actually, the very last step in this argument relies on verifying that a permutation in $S_n(\alpha; \beta)$ having an i.d.s. of length *exactly* $r-1$ also behaves itself, namely it has r children labelled $(r), (2), \ldots, (r)$. This is easy enough; we can then summarize the succession rules as

$$\begin{aligned} &\text{Rule: } (k) \to (2)\cdots(k)(k+1) \quad (\text{if } k<r), \\ &\text{Rule: } (r) \to (2)\cdots(r)(r). \end{aligned} \tag{2.3}$$

The first rule applies if the i.d.s. has length $k-1 < r-1$, and the second rule if the i.d.s. has length $r-1$ or greater.

This gives a *finite* transfer matrix of dimension $r-1$ (labels smaller than 2 or greater than r never arising) of the form

$$A_r = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 1 & 1 & \cdots & 1 & 1 & 1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 0 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 1 & 2 \end{bmatrix} \tag{2.4}$$

which can be regarded as a truncated form of A_∞.

The other two cases will proceed analogously. First we consider $\alpha = 213$, $\beta = 1, 2, \ldots, r, r+1$. Let $\pi \in S_{n-1}(\alpha; \beta)$, and look at the potential children of π in $T(\alpha; \beta)$. The restriction α guarantees that the insertion point for n is in or immediately after the i.i.s. We again distinguish three cases depending on whether the length $k-1$ of this i.i.s. is less than, equal to, or greater than, $r-1$.

First, if $k-1 < r-1$, all k potential children are β-free, and they have i.i.s.'s of lengths $1, 2, \ldots, k$. Second, if $k-1 = r-1$, all k potential children are β-free, and their i.i.s.'s have lengths $1, 2, 3, \ldots, r$. Third, if $k-1 > r-1$, only the first r of the k potential children are β-free, with i.i.s.'s of lengths $1, 2, \ldots, r$. Since we have just shown that the number of children can be deduced directly from $k-1$, we can read off

$$\begin{array}{rccccccc} & & \text{i.i.s.} & & & \text{children} & \\ \text{short} & k-1 & \to & 1,2,\ldots,k & k & \to & 2,3,\ldots,k,k+1 \\ \text{critical} & r-1 & \to & 1,2,\ldots,r & r & \to & 2,3,\ldots,r,r \\ \text{long} & k-1 & \to & 1,2,\ldots,r & r & \to & 2,3,\ldots,r,r \end{array} \tag{2.5}$$

The succession rules, and hence the transfer matrix, are thus identical to the previous example.

Parenthetically, we should remark that in this case $k-1 > r-1$ means $k-1 = r$, since r is the maximum length permitted for an i.i.s. by β. This type of restriction was not true for the i.d.s. in the previous example; nevertheless the succession rules are identical.

In our third example, $\alpha = 213$, $\beta = r+1, 1, \ldots, r$. Again α restricts our attention to the i.i.s., of length $k-1$. Now β means that a new largest element can only be inserted either immediately after the i.i.s., or in the *last* (rightmost) $r-1$ positions of the i.i.s., a total of r possible insertions. However, in this instance it is not obvious (or true) that all r of these possibilities are valid; there might be an increasing sequence of length r lurking in the permutation somewhere beyond the i.i.s. What makes this example more delicate than the preceding two is that in those instances both α and β ended with their respective largest elements, so that we only needed to look to the left of our insertion point to be sure of avoiding new forbidden patterns. Now we must look to the right as well as to the left.

The general situation is of a permutation $\pi \in S_{n-1}$, with i.i.s. of length $k-1$, and where there are l legitimate places to insert a new largest element. These will give rise to the permutations $\pi^{k-l+1}, \pi^{k-l+2}, \ldots, \pi^k$. We know the active insertion points form such a consecutive run because if a site is excluded by pattern β, so will all sites to its left be excluded. Now select one of these sites, $k-l+1 \leqslant j \leqslant k$, insert n to form a new permutation ρ, and consider which sites in ρ now admit an insertion of element $(n+1)$. If $l<r$, the legitimate children of ρ will be $\rho^{k-l+1},\dots,\rho^{j+1}$ – ρ^{k-l} being excluded because it contains essentially the same subsequence which excludes π^{k-l}, and ρ^{j+2} being excluded because $\rho^j,\rho^{j-1},n+1$ would be of type α. Thus as j ranges from $k-l+1$ to k the number of children of ρ ranges from $(k-l+1+1)-(k-l+1)+1=2$ to $(k+1)-(k-l+1)+1=l+1$. This gives us the succession rule

Rule: $(l) \to (2)\cdots(l)(l+1)$.

It is clearly not possible to have $l>r$ as then the leftmost of these insertion points would necessarily create a forbidden β. We are left with $l=r$ and the insertion points $k-r+1,\dots,k$. This proceeds exactly as $l<r$ except that now $\rho=\pi^k$ has only l children rather than $l+1$, the reason being that ρ^{k-l+1} is excluded by reason of causing a forbidden β. We therefore exceptionally have

Rule: $(r) \to (2)\cdots(r)(r)$.

Hence the transfer matrix is again identical to the two preceding examples.

In the following section we will therefore be analyzing only the single matrix equation

$$[(i+1)_n] = A_r^{n-1} e_1^{\mathrm{T}}. \tag{2.6}$$

3. Algebra

For each r, we are interested primarily in a generating function for Σ_n, the cardinality of $S_n(\alpha;\beta)$. Conveniently, $\Sigma_n=(2)_{n+1}$, since each permutation begets exactly one child with the label (2). It is also immediate from (2.6) that $(2)_{n+1}=(A_r^n)_{1,1}$, the upper-left entry in the nth power of the transfer matrix A_r.

It is well known (see [3] Theorem 4.7.2) that this generating function is

$$f_r(x) = \sum_{n\geqslant 0} (A_r^n)_{1,1}\, x^n = \frac{\det(I-xA_r:1,1)}{\det(I-xA_r)}, \tag{3.1}$$

where $(B:j,i)$ denotes the matrix obtained by removing the jth row and ith column of B.

In our case, A_r is a matrix of dimension $(r-1\times r-1)$ and $(I-xA_r:1,1)$ is just $I-xA_{r-1}$. Therefore if we let

$$p_m(\lambda) = \det(A_{m+1} - \lambda I_{m+1})$$

be the characteristic polynomial of A_{m+1}, then our generating function becomes

$$\frac{\det(I-xA_{r-1})_{(r-2\times r-2)}}{\det(I-xA_r)_{(r-1\times r-1)}} = \frac{(-x)^{r-2}p_{r-2}(1/x)}{(-x)^{r-1}p_{r-1}(1/x)}. \tag{3.2}$$

To determine the $p_m(\lambda)$, we expand down the first column. This gives us $p_m(\lambda) = (1-\lambda)p_{m-1}(\lambda) - q_{m-1}(\lambda)$, where

$$q_m(\lambda) = \det \begin{bmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 1-\lambda & \cdots & 1 & 1 \\ & 1 & \cdots & 1 & 1 \\ & & & \vdots & \vdots \\ & & & 1-\lambda & 1 \\ & & & 1 & 2-\lambda \end{bmatrix}_{(m\times m)}. \tag{3.3}$$

Expanding this new determinant similarly gives us $q_m(\lambda) = p_{m-1}(\lambda) - q_{m-1}(\lambda)$. We may eliminate the q's to obtain the recurrence

$$p_m + \lambda p_{m-1} + \lambda p_{m-2} = 0.$$

This suggests that $p_m(\lambda)$ is related to Chebyshev polynomials; indeed, it is straightforward to use this recurrence and the well-known recurrence

$$U_{m+2}(y) = 2yU_{m+1}(y) - U_m(y) \quad (m>2) \tag{3.4}$$

for the Chebyshev polynomials to show by induction that

$$p_m(\lambda) = (-1)^m \lambda^{(m-2)/2} U_{m+2}\left(\frac{\sqrt{\lambda}}{2}\right), \tag{3.5}$$

where U_m is the Chebyshev polynomial of the second kind.

We can state our conclusion as follows.

Theorem 3.1. *We have*

$$\begin{aligned} \Sigma_n &= |S_n(123; r, r-1, \ldots, 2, 1, r+1)| \\ &= |S_n(213; 1, 2, \ldots, r, r+1)| \\ &= |S_n(213; r+1, 1, 2, \ldots, r)|, \end{aligned}$$

and the generating function for this sequence is

$$f_r(x) = \sum_{n\geq 0} \Sigma_n x^n = -\frac{p_{r-2}(1/x)}{x p_{r-1}(1/x)}, \tag{3.6}$$

where

$$p_r(\lambda) = (-1)^r \lambda^{(r-2)/2} U_{r+2}\left(\frac{\sqrt{\lambda}}{2}\right).$$

4. Context

The present result is an extension of previously known results for $r=2$ and 3. In [2] it is shown that $|S_n(123; 213)| = |S_n(213; 312)| = 2^{n-1}$, the generating function being

$(1-x)/(1-2x)=1+x+2x^2+4x^3+\cdots$. This has the expected denominator; in this instance it is not clear what numerator to expect from our result, the determinant in question being void.

In [5] we find $|S_n(123;3214)|=|S_n(213;1234)|=|S_n(213;4123)|=F_{2n}$, the alternating Fibonacci numbers. Our present result gives the generating function

$$\frac{(-x)(2-1/x)}{x^2(1-3/x+1/x^2)}=\frac{1-2x}{1-3x+x^2}=1+x+2x^2+5x^3+13x^4+\cdots\ . \tag{4.1}$$

This result was obtained in [5] with an identical combinatorial approach and the succession rules $(2)\to(2)(3)$ and $(3)\to(2)(3)(3)$ were presented there, together with other closed-form enumerations for $|S_n(\alpha;\beta)|$ where $\alpha\in S_3, \beta\in S_4$. The present paper marks the first time to our knowledge that such an explicit result has been obtained for any $|S_n(\alpha;\beta)|$ in which one of the restrictions has indeterminate length.

It is well known that the order of growth of the generating function $f_r(\lambda)$ is given by the largest zero of $p_m(\lambda)=p_{r-1}(\lambda)$ (that is, the largest eigenvalue of the transfer matrix). Thus when $r=2$, the largest zero of $2-\lambda$ is 2; when $r=3$, the largest zero of $1-3\lambda+\lambda^2$ is $(3+\sqrt{5})/2=((1+\sqrt{5})/2)^2$.

Since we have

$$p_m(\lambda)=(-1)^m\lambda^{(m-2)/2}U_{m+2}\left(\frac{\sqrt{\lambda}}{2}\right),$$

λ will be a non-zero root of $p_{r-1}(\lambda)=0$ exactly when $y=\sqrt{\lambda}/2$ is a non-zero root of $U_{r+1}(y)=0$. The zeroes of the Chebyshev polynomial $U_{r+1}(y)$ are

$$\cos\left(\frac{\pi}{r+2}\right),\cos\left(\frac{2\pi}{r+2}\right),\ldots,\cos\left(\frac{(r+1)\pi}{r+2}\right),$$

so the roots of $p_{r-1}(\lambda)$ are

$$4\cos^2\left(\frac{\pi}{r+2}\right),\ldots,4\cos^2\left(\frac{(r+1)\pi}{r+2}\right).$$

The largest of these is $4\cos^2(\pi/(r+2))$.

We thus obtain the asymptotics for each of the three cases enumerated in Theorem 3.1.

Corollary 4.1. $\Sigma_n \asymp (4\cos^2(\frac{\pi}{r+2}))^n$.

Recalling that $\lim_{r\to\infty} A_r=A_\infty$ or noting that an infinitely long restriction is no restriction at all, we see that in the limit as $r\to\infty$,

$$|S_n(123)|\asymp(4\cos^2(0))^n=4^n.$$

This makes sense as we know that $|S_n(123)|=c_n\asymp 4^n$.

5. Connections

Ratios of successive Chebyshev polynomials occur in various areas of enumerative combinatorics, and this suggests that the numbers Σ_n might count other natural combinatorial objects. This is indeed the case, as we now indicate briefly. Optimistically, these connections could shed new light on the enumeration of permutations with excluded subsequences.

The first example is a set of lattice paths. Consider lattice paths from $(0,0)$ to $(2n,0)$ that never dip below the x-axis, with each step in the path either parallel to the vector $\langle 1,1\rangle$ and called a northeast (NE) step or else parallel to $\langle 1,-1\rangle$ and called a southeast (SE) step. Each path contains the same number n of each type of step, and is considered to be of length n. It is well known that the set of all such paths is enumerated by the Catalan numbers.

To show that these lattice paths are generated by the succession rule (2.2), give each path the label k, where k is one greater than its longest terminal sequence of SE steps. The k children of this path can be formed by inserting an NE step into each of the final $k+1$ possible positions, and then appending a SE step at the very end. Clearly the labels of the n resulting paths are then, in order, $k+1, k, k-1, \ldots, 2$. Also it is clear that the parent is recoverable from the child by deleting the last NE and the last SE steps, so that each path of length $n+1$ is generated exactly once.

To obtain the modified succession rules (2.3), we simply insist that the inserted NE occurs in one of the final r possible insertion points. The resulting class, which is enumerated by (3.6), is the subclass of all lattice paths from $(0,0)$ to $(2n,0)$ which never rise above the line $y=r$.

The second example is a set of plane trees. A plane tree is recursively defined as a root together with an ordered list of plane trees. The last tree in an ordered list is the *right-hand* subtree, and its root is called the *rightmost* child. We can form the *right-hand edge* of a plane tree by recursively taking the root together with the right-hand edge of its right-hand subtree.

Label a plane tree with k, the number of nodes in its right-hand edge. To form children, take some node in the right-hand edge and add a new rightmost child to it. This can be done in k ways, one for each of the nodes in the right-hand edge. The labels of these children are, starting at the root, $2,3,\ldots,k+1$. To recover a parent from a child, simply delete the last node in the right-hand edge.

To obtain (2.3), add nodes only to the first r nodes in the right-hand edge. The resulting class, enumerated by (3.6), is the class of all plane trees of depth at most r.

We remark that at a talk given by the second author at the British Combinatorial Conference, Robin Chapman drew our attention to problem 10570 in the *American Mathematical Monthly* problems section, January 1997. This deals with plane trees of bounded depth, and (independently from our work, of course) Chapman found that quotients of Chebyshev polynomials arose in his solution.

The third example is convex directed animals, enumerated by perimeter. These are subsets of the plane, bounded above and below by lattice paths, where each boundary begins at $(0,0)$ and is constructed of steps in the directions $\langle 0,1\rangle$ and $\langle 1,0\rangle$. The animal can thus be thought of as made up of squares, with each square identified with the lattice point at its lower left corner.

If each boundary consists of n steps, the semiperimeter is n. If an animal has $k-1$ squares in its rightmost column, let its label be k. To form the k children, either add a square to the top of the rightmost column (creating a child with label $k+1$), or add a new rightmost column containing $1,2,\ldots,k-1$ squares and flush at the top with the old rightmost column (creating a child with label $2,3,\ldots,k$). It is simple to recover a parent from a child, after checking whether the two rightmost columns are flush at the top.

If we want the restricted rules (2.3), then there should never be more than r children. So we always add a new column when an old column reaches height r. Therefore, the set of convex directed animals of bounded column height r is enumerated by (3.6).

Acknowledgements

The material in the last section of this paper was worked out collaboratively by the second author and Alberto del Lungo. We are grateful to Ira Gessel and David Wagner for insight into some of our Chebyshev polynomial calculations. We are also grateful to the Mathematical Sciences Research Institute, Berkeley, where much of this work was completed. Significant financial support came from NSERC operating grant OGP0105492. The first author was supported in part by a National Science Foundation postdoctoral fellowship and did some of the work for this paper while at the University of Michigan.

References

[1] F.R.K. Chung, R.L. Graham, V.E. Hoggatt Jr, M. Kleiman, The number of baxter permutations, J. Combin. Theory Ser. A 24 (1978) 382–394.
[2] R. Simion, F.W. Schmidt, Restricted permutations, European J. Combin. 6 (1985) 383–406.
[3] R.P. Stanley, Enumerative Combinatorics, vol. 1, Cambridge University Press, New York, 1997.
[4] J. West, Generating trees and the Catalan and Schroder numbers, Discrete Math. 146 (1995) 247–262.
[5] J. West, Generating trees and forbidden subsequences, Discrete Math. 157 (1996) 363–374.