



ELSEVIER

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta


Enumeration of non-crossing pairings on bit strings

 Todd Kemp^a, Karl Mahlburg^b, Amarpreet Rattan^c, Clifford Smyth^d
^a Department of Mathematics, UCSD, 9500 Gilman Drive, La Jolla, CA 92093-0112, United States

^b Department of Mathematics, Princeton University, Washington Rd, Princeton, NJ 08540, United States

^c Department of Mathematics, University of Bristol, University Walk, Bristol, BS8 1TW, UK

^d Mathematics and Statistics Department, University of North Carolina Greensboro, 116 Petty Building, Greensboro, NC 27402, United States

ARTICLE INFO

Article history:

Received 11 June 2009

Available online 14 August 2010

Keywords:

Non-crossing pairings

Free probability and random matrices

Fuss–Catalan numbers

Bijective combinatorics

ABSTRACT

A non-crossing pairing on a bit string is a matching of 1s and 0s in the string with the property that the pairing diagram has no crossings. For an arbitrary bit-string $w = 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r}$, let $\varphi(w)$ be the number of such pairings. This enumeration problem arises when calculating moments in the theory of random matrices and free probability, and we are interested in determining useful formulas and asymptotic estimates for $\varphi(w)$. Our main results include explicit formulas in the “symmetric” case where each $p_i = q_i$, as well as upper and lower bounds for $\varphi(w)$ that are uniform across all words of fixed length and fixed r . In addition, we offer more refined conjectural expressions for the upper bounds. Our proofs follow from the construction of combinatorial mappings from the set of non-crossing pairings into certain generalized “Catalan” structures that include labeled trees and lattice paths.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

A pairing π of $V = \{1, \dots, 2k\}$ is a partition of V into k pairs, $\pi = \{\{i_1, j_1\}, \dots, \{i_k, j_k\}\}$. A crossing of π is a pair of pairs $\{i_1, j_1\}, \{i_2, j_2\} \in \pi$ such that $i_1 < i_2 < j_1 < j_2$, see Fig. 1. The pairing π is called *non-crossing* if it has no crossings. Let $NC_2(2k)$ be the set of all non-crossing pairings of $\{1, \dots, 2k\}$.

A *binary word* (or bit-string) of length n is $w = w_1 w_2 \dots w_n$ with $w_i \in \{0, 1\}$ for $1 \leq i \leq n$. We write $|w|$ for the length of w . If w is a binary word of length $2k$ and $\pi \in NC_2(2k)$, we say w and π are *compatible* if, for each pair $\{i, j\} \in \pi$, $w_i \neq w_j$; that is, the letters in w that are paired by π

E-mail addresses: tkemp@math.ucsd.edu (T. Kemp), mahlburg@math.princeton.edu (K. Mahlburg), amarpreet.rattan@bristol.ac.uk (A. Rattan), [cgsmyth@uncg.edu](mailto:cdsmyth@uncg.edu) (C. Smyth).

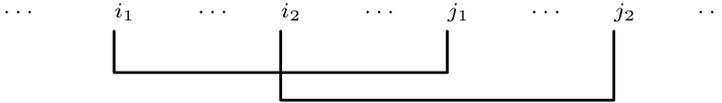


Fig. 1. A crossing in a pairing.

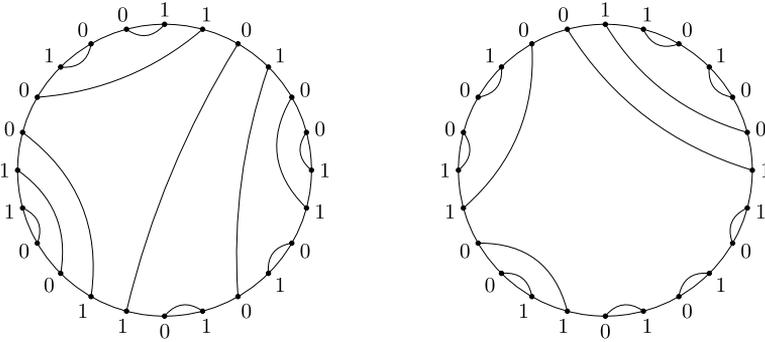


Fig. 2. Two members of $NC_2(110100110101011001100100)$.

are distinct. Note that if w is compatible with some pairing π then w is necessarily *balanced*, i.e. w contains the same number of 1s as 0s. We are interested in the set of pairings compatible with w .

Definition 1.1. Let w be a binary word with $|w| = 2k$. Then the set of *non-crossing pairings* on w is

$$NC_2(w) := \{ \pi \in NC_2(2k) : \pi \text{ and } w \text{ are compatible} \},$$

and the number of such pairings is denoted by

$$\varphi(w) := |NC_2(w)|.$$

Example 1.2. If $w = 110100$, then $NC_2(w) = \{ \pi_1, \pi_2 \}$ where $\pi_1 = \{ \{1, 6\}, \{2, 5\}, \{3, 4\} \}$ and $\pi_2 = \{ \{1, 6\}, \{2, 3\}, \{4, 5\} \}$. Thus $\varphi(w) = 2$. Similarly, $\varphi(101010) = 5$ and $\varphi(111000) = 1$.

Example 1.3. Let $w = 110100110101011001100100$. Each of the two diagrams in Fig. 2 represents a pairing compatible with w . In each diagram, w is listed clockwise around the circle, beginning with the topmost 1, while the internal arcs in the diagram represent the pairs.

The function φ arises naturally in random matrix theory. Let X_n be an $n \times n$ matrix whose entries are all independent, identically distributed complex random variables, each with mean 0 and variance $1/n$. Such a matrix is typically not normal (this is true, for example, if the matrix, viewed as an n^2 -dimensional random vector, has a continuous n^2 -dimensional density). Thus the eigenvalues, which are themselves complex random variables, are difficult to compute. Therefore, set $G_n = \frac{1}{2}(X_n + X_n^*)$, the Hermitian cousin of X_n ; the matrix G_n has been studied by physicists for over half a century. The density (histogram) of eigenvalues of G_n converges as $n \rightarrow \infty$, essentially regardless of how the matrix entries are actually distributed, to the *semicircle law* $\frac{1}{2\pi} \sqrt{4 - x^2}$ on $[-2, 2]$, cf. [3,11]. For a Hermitian matrix such as G_n , the density of eigenvalues contains the same information as the *matrix moments* $\frac{1}{n} \text{Tr}(G_n^p)$ for $p \in \mathbb{Z}_+$, where Tr denotes the ordinary trace of a matrix. For a non-Hermitian matrix like X_n , one studies instead the *mixed matrix moments* $\frac{1}{n} \text{Tr}(X_n^{p_1} X_n^{*q_1} \dots X_n^{p_r} X_n^{*q_r})$, which do not correlate as directly with eigenvalues, and in general contain vastly more data. The connection between these moments and our interests is summed up in the following proposition, whose proof can be found in [6]; see also [13,17].

Proposition 1.4. If $p, q \in (\mathbb{Z}_+)^r$ are r -tuples of positive integers then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \text{Tr}(X_n^{p_1} (X_n^*)^{q_1} \dots X_n^{p_r} (X_n^*)^{q_r}) = \varphi(1^{p_1} 0^{q_1} 1^{p_2} 0^{q_2} \dots 1^{p_r} 0^{q_r})$$

almost surely.

Definition 1.5. For all $p, q \in (\mathbb{Z}_+)^r$, the number of non-crossing pairings on r -tuples is defined as

$$\varphi(p, q) := \varphi(1^{p_1} 0^{q_1} 1^{p_2} 0^{q_2} \dots 1^{p_r} 0^{q_r}).$$

We also define the *weight* of an r -tuple of integers p to be $|p| := \sum_{i=1}^r p_i$. Note that $\varphi(p, q) = 0$ unless the underlying word is balanced, i.e. $|p| = |q|$.

Our goal, in a sense, is to calculate all asymptotic mixed matrix moments of a random matrix with independent entries. However, the enumeration of such non-crossing pairings is relevant in more general circumstances. In [12], Nica and Speicher introduced \mathcal{R} -diagonal operators, which represent the limiting eigenvalue distributions of a large class of non-Hermitian random matrices with non-independent entries (but that nevertheless have nice symmetry and invariance properties). Such ensembles of random matrices have recently played very important roles in free probability and beyond: for example, in [7], Haagerup has produced the most significant progress towards the resolution of the Invariant Subspace Conjecture in decades, and his proof is concentrated in the theory of \mathcal{R} -diagonal operators. In [8], the first author of the present paper showed that the asymptotic mixed matrix moments of \mathcal{R} -diagonal random matrices are controlled, in an appropriate sense, by the set of non-crossing pairings we consider in this paper. Indeed, the results of the present paper followed from discussions motivated by applications to \mathcal{R} -diagonal operators.

Computations with small cases illustrate that $\varphi(w)$ depends on w in a very complex way. For example, while $\varphi(1^i 0^j 1^k 0^l) = \binom{i+1}{2} + \binom{i+1}{1} \binom{j+1}{1}$ if $i \leq j \leq k$ (cf. Theorem 1.17), the general formula for $\varphi(1^{p_1} 0^{q_1} 1^{p_2} 0^{q_2} 1^{p_3} 0^{q_3})$ takes several lines to write down. Although a closed form for φ may be unobtainable, much can still be said.

We define an important parameter of a word, the number of runs, $r(w)$. For $i = 0, 1$ an i -block of w is a maximal subword of cyclically adjacent i 's in w . Let $r(w)$ be the number of 1-blocks in w (equivalently, the number of 0-blocks). Even if we restrict our attention to words with the same length and the same number of runs, the value of φ can fluctuate wildly. We seek the maximum value of φ over each such class and, ideally, the words at which the maximum occurs.

Main Problem. For all $1 \leq r \leq k$, determine

$$\text{Max}_{k,r} := \max\{\varphi(w) : w \text{ balanced, } |w| = 2k, r(w) = r\},$$

and find the w for which this maximum is attained.

We now outline our main results (the proofs are all left for Section 3). The following result is an important first step toward addressing our main problem.

Theorem 1.6 (The Symmetrization Theorem). For all $p, q \in (\mathbb{Z}_+)^r$,

$$\varphi(p, q) \leq \varphi(p, p).$$

This shows that in order to determine $\text{Max}_{k,r}$, it suffices to restrict attention to *symmetric* words, i.e. words of the form $1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r}$ with $p_i = q_i$ for $1 \leq i \leq r$.

For all $m \geq 1, r \geq 0$ the corresponding *Fuss–Catalan number* is $C_r^{(m)} := \frac{1}{mr+1} \binom{(m+1)r}{r}$. Note that $C_r^{(1)} = C_r$ is the ordinary *Catalan number*. It is well known that $|NC_2(2r)| = C_r$ (cf. [15]), and the Fuss–Catalan numbers also count certain pairings on words.

Proposition 1.7. For all $m \geq 1$, $\varphi((1^m 0^m)^r) = C_r^{(m)}$.

Remark 1.8. If $r = 1$ it is easy to see that $\varphi(1^m 0^m) = 1 = C_1^{(m)}$. If $m = 1$ it is also easy to see that any $\pi \in NC_2(2r)$ is automatically compatible with $w = 1010 \dots 10$. Thus $\varphi((10)^r) = C_r = C_r^{(1)}$. Indeed, if $\{i, j\} \in \pi$ and $w_i = w_j$ then $|i - j|$ must be even. Since an odd number of points then lie between i and j there must be another pair of π with exactly one end between i and j , contradicting the assumption that π is non-crossing. A more sophisticated version of this reasoning, together with the recurrence for the Fuss–Catalan numbers, forms a proof of Proposition 1.7 as discussed in [4]. This proposition was also proved in a more topological manner by the first author in [9], which relies on the non-crossing partition multichain enumeration results in [5] (proofs are also essentially contained in [10,14]).

Our most significant result is a near-sharp upper bound for $\varphi(w)$.

Theorem 1.9 (Main Theorem). For $1 \leq r \leq k$,

$$\text{Max}_{k,r} \leq \varphi\left(\left(1 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \uparrow \! \! \! 0 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \downarrow \! \! \! \right)^r\right) = C_r^{\left(\begin{smallmatrix} k \\ r \end{smallmatrix}\right)}.$$

Notice that when $r \mid k$, our Main Theorem is exact, and implies that

$$\text{Max}_{k,r} = \varphi\left(\left(1 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \uparrow \! \! \! 0 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \downarrow \! \! \! \right)^r\right) = C_r^{\left(\frac{k}{r}\right)},$$

i.e. the maximum occurs at a word whose 1-blocks and 0-blocks are all of the same length.

We believe the following sharp statement can be made for the other cases of r and k .

Conjecture 1.10 (Main Conjecture). If $r \nmid k$ then

$$\text{Max}_{k,r} = \varphi\left(\left(1 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \uparrow \! \! \! 0 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \downarrow \! \! \! \right)^{r'} \left(1 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \uparrow \! \! \! 0 \begin{smallmatrix} k \\ r \end{smallmatrix} \! \! \! \downarrow \! \! \! \right)^{r''}\right)$$

where r', r'' are the unique positive integers with $k = r' \lfloor \frac{k}{r} \rfloor + r'' \lceil \frac{k}{r} \rceil$ and $r = r' + r''$.

In other words, we believe $\text{Max}_{k,r}$ is attained when w is symmetric and has 1-blocks and 0-blocks that are as equal in length as possible, with all of the largest blocks grouped adjacently.

Theorem 1.17 below provides exact polynomial formulas for the symmetric case $\varphi(p, p)$. These formulas will be instrumental in our proof of Theorems 1.6 and 1.9 in Section 3. They will also play a key part in the statement of Conjecture 1.21 below, an intricate but appealing result that directly implies Conjecture 1.10. We now introduce the notation needed to enumerate non-crossing pairings in the symmetric case.

Following Stanley [15, p. 294], we recursively define a *plane tree* T to be a finite non-empty set of vertices so that (i) one specially designated vertex in T is called the *root* of T and (ii) the remaining vertices of T , excluding the root, are partitioned into an ordered list, (T_1, \dots, T_d) , of $d \geq 0$ disjoint non-empty sets T_1, \dots, T_d , each of which is a plane tree. Let $|T|$ be the number of vertices of T . If $r \geq 1$, let \mathcal{T}_r denote the set of plane trees on r vertices, i.e. the set of isomorphism classes of plane trees with $|T| = r$. It is well known that $|\mathcal{T}_{r+1}| = C_r$, the r th Catalan number (cf. [15]).

We also make use of the following standard definitions and terminology. Given a plane tree T with root u and subtrees T_i as in the definition, the *decomposition* of T is (u, T_1, \dots, T_d) . The vertices in T apart from u are all *descendants* of u , and if u_i is the root of T_i for $1 \leq i \leq d$, then the u_i are the *children* of u , and u is their *parent*. Finally, all of the u_i are *siblings*.

The *canonical ordering* (also commonly known as the *depth-first* or *clockwise* ordering) of the vertices of T is recursively defined by first putting $u < T_1 < \dots < T_d$, i.e. $u < v < w$ for all $v \in T_i, w \in T_j$, where $i < j$. Then, each T_i is canonically ordered internally. If $v_1 < \dots < v_r$ is the canonical ordering of the vertices V of T , then the *canonical vertex labeling* of T is the function $\Lambda : V \rightarrow [r]$ given by $\Lambda(v_i) = i$. Given v in T , let T_v be the subtree of T with root v . The *degree* of v in T , $d_T(v)$, is defined

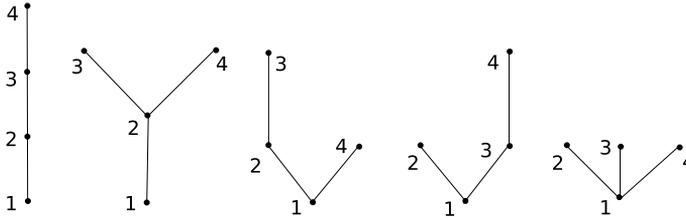


Fig. 3. The five plane trees on $V = \{1, 2, 3, 4\}$, canonically labeled. The respective degree sequences are $(1, 1, 1, 0)$, $(1, 2, 0, 0)$, $(2, 1, 0, 0)$, $(2, 0, 1, 0)$ and $(3, 0, 0, 0)$.

to be the number of children of v in T . The *degree sequence* of T is the sequence $(d_T(v) : v \in V)$ where the vertices are listed in canonical order.

Remark 1.11. As in Fig. 3, we always depict plane trees so that (i) each vertex is connected to each of its children by an edge, (ii) children are positioned above their parent, (iii) siblings are ordered from left to right in canonical order.

Remark 1.12. The notion of plane tree can be extended to define a *plane forest* F as a tuple (T_1, \dots, T_n) , where each T_j is a plane tree. The set of all such forests with r vertices is denoted by \mathcal{F}_r . Furthermore, the *canonical ordering* of F is given by setting $T_1 < \dots < T_n$, and then recursively ordering each T_j . As a final observation, a plane tree consists precisely of a root u and a plane forest (T_1, \dots, T_d) , with edges from u to the roots of each T_i .

If $\gamma : [r] \rightarrow \mathbb{Z}$ we also write γ as an r -tuple $\gamma = (\gamma(1), \dots, \gamma(r))$. If $\gamma : [r] \rightarrow \mathbb{Z}$ and $\gamma' : [r'] \rightarrow \mathbb{Z}$, the *concatenation* $\tau = \gamma\gamma'$ is the map $\tau : [r+r'] \rightarrow \mathbb{Z}$ where $\tau(i) := \gamma(i)$ for $i \in [r]$ and $\tau(r+i) := \gamma'(i)$ for $i \in [r']$.

Definition 1.13. If $T \in \mathcal{T}_r$ and $\gamma : [r] \rightarrow \mathbb{Z}$ is an injective map, then the *min-first* vertex labeling of T by γ is the following recursively defined map $\gamma_T : T \rightarrow \mathbb{Z}$. Let $1 \leq j \leq r$ be such that $\gamma_j = \gamma_{\min} := \min\{\gamma(i) : 1 \leq i \leq r\}$ and let $\gamma' = \text{Rot}_j(\gamma) := (\gamma(j), \gamma(j+1), \dots, \gamma(r), \gamma(1), \dots, \gamma(j-1))$ be the *left-rotation* of γ to its minimum element. Using the decomposition of T , written as (u, T_1, \dots, T_d) , let $\gamma'_i : [T_i] \rightarrow \mathbb{Z}$ for $1 \leq i \leq d$ be defined so that $\gamma' = \gamma_{\min}\gamma'_1 \dots \gamma'_d$. Then the labeling is recursively given by $\gamma_T := \gamma_{\min}(\gamma'_1)_{T_1} \dots (\gamma'_d)_{T_d}$.

Note that this definition gives γ_T as a permutation of γ where for each i , $\gamma_T(i)$ is the label given by γ_T to the i th vertex of T in canonical order.

Remark 1.14. Similarly, the vertex labeling of a plane forest $F = (T_1, \dots, T_n) \in \mathcal{F}_r$ by an injective map $\gamma : [r] \rightarrow \mathbb{Z}$ is the map γ_F determined by decomposing $\gamma' = \gamma'_1 \dots \gamma'_n$ and recursively labeling each T_i by γ'_i .

Example 1.15. We give an example to illuminate Definition 1.13. Let T be the third tree in Fig. 3 and $\gamma = (4, 1, 3, 2)$. T has decomposition $T = (1, T_2, T_4)$ where $T_2 = (2, T_3)$, $T_3 = (3)$ and $T_4 = (4)$. We have $\gamma' = (1, 3, 2, 4) = 1\gamma'_1\gamma'_2$ where $\gamma'_1 = (3, 2)$ and $\gamma'_2 = (4)$. Recursively, $(\gamma'_1)_{T_2} = (2, 3)$ and $(\gamma'_2)_{T_4} = (4)$ so $\gamma_T = 1(\gamma'_1)_{T_2}(\gamma'_2)_{T_4} = (1, 2, 3, 4)$. Fig. 4 shows each $T \in \mathcal{T}_4$ labeled with γ_T for $\gamma = (4, 1, 3, 2)$.

Note that in Example 1.15 the resulting labeling of T is identical to the canonical labeling that arises from labeling by $e = (1, 2, 3, 4)$; the labeling of the fourth tree in Fig. 3 by γ is different from the canonical labeling. Fig. 5 shows a less trivial example of a (non-canonical) vertex labeling of a plane tree.

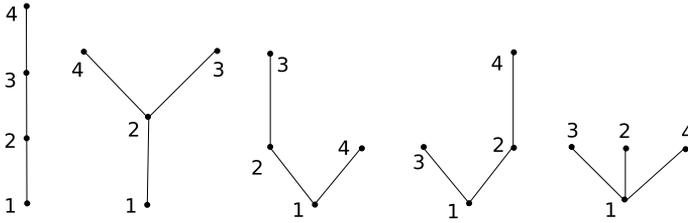


Fig. 4. Plane trees labeled by $\gamma = (4, 1, 3, 2)$. The label $\gamma_T(v)$ appears next to each vertex v .

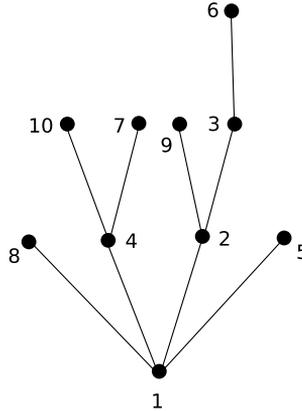


Fig. 5. A plane tree T . If $\gamma = (2, 9, 6, 3, 5, 1, 8, 7, 4, 10)$, $\gamma_T = (1, 8, 4, 10, 7, 2, 9, 3, 6, 5)$. $\gamma_T(v)$ is shown next to each vertex v .

Remark 1.16. It is clear from the definition that (i) γ_T is increasing, i.e. $(\gamma_T)(v) < (\gamma_T)(w)$ whenever w is a descendant of v , and (ii) for each v , the sub-labeling $\gamma_T(T_v)$ is a cyclically consecutive subsequence of γ . Note that if γ is increasing, i.e. $\gamma(i) < \gamma(j)$ for all $i < j$, then no rotations ever occur in the calculation of γ_T . In this case, if $v_1 < \dots < v_r$ is the canonical ordering of T then $(\gamma_T)(v_i) = \gamma(i)$ for all i .

We use the vertex labelings described above to define polynomials that arise in the enumeration of $\varphi(p, p)$. Let S_r denote the symmetric group on $[r]$, where $e = (1, 2, \dots, r)$ is the identity permutation. If $T \in \mathcal{T}_r$ and $\gamma \in S_r$, the tree-monomial in the indeterminates (x_1, \dots, x_r) is defined as

$$m_{T,\gamma}(x_1, \dots, x_r) := \prod_{v \in T} \binom{x_{\gamma_T(v)} + 1}{d_T(v)}.$$

The tree-polynomial is then the sum

$$P_\gamma(x_1, x_2, \dots, x_r) := \sum_{T \in \mathcal{T}_r} m_{T,\gamma}(x_1, \dots, x_r).$$

Finally, we construct arbitrary r -tuples of integers by beginning with a weakly increasing vector $p \in (\mathbb{Z}_+)^r$ (i.e. $p_1 \leq \dots \leq p_r$), and then permuting the entries. In particular, we define the natural group action as $p_\gamma := (p_{\gamma_1}, \dots, p_{\gamma_r})$.

Theorem 1.17 (The Polynomial Formula). If $\gamma \in S_r$ and $p \in (\mathbb{Z}_+)^r$ is weakly increasing, then

$$\varphi(p_\gamma, p_\gamma) = P_\gamma(p).$$

For notational convenience in writing tree polynomials, if $d \geq 0$ is an integer and x an indeterminate, we define $[x]^d := \binom{x+1}{d} = \frac{1}{d!} (x+1)(x) \dots (x-d+2)$, with $[x]^0 := 1$. Note that for integers $p \geq 0$, $[p]^d = 0$ when $p < d - 1$, so some of the terms in $P_\gamma(p)$ may vanish.

Example 1.18. Let $e = (1, 2, 3, 4)$ and $\gamma = (4, 1, 3, 2)$. If $p = (p_1, p_2, p_3, p_4)$ is weakly increasing, we can apply Theorem 1.17 to show

$$\begin{aligned} \varphi(p, p) &= P_e(p) = [p_1]^1 [p_2]^1 [p_3]^1 + [p_1]^1 [p_2]^2 + [p_1]^2 [p_2]^1 + [p_1]^2 [p_3]^1 + [p_1]^3, \\ \varphi(p_\gamma, p_\gamma) &= P_\gamma(p) = [p_1]^1 [p_2]^1 [p_3]^1 + [p_1]^1 [p_2]^2 + [p_1]^2 [p_2]^1 + [p_1]^2 [p_2]^1 + [p_1]^3 \end{aligned}$$

(see Remark 1.16 and Fig. 4). Note that these formulas imply

$$\varphi(p_\gamma, p_\gamma) \leq \varphi(p, p)$$

for all weakly increasing p . This is because we have $m_{T,4132} \equiv m_{T,e}$ for every $T \in \mathcal{T}_4$ except T' , the fourth tree listed in Fig. 4. But for T' , we have $m_{T',4132}(p) = [p_1]^2 [p_2]^1 \leq [p_1]^2 [p_3]^1 = m_{T',e}(p)$ since $p_2 \leq p_3$.

The polynomial inequalities in this example are indicative of a much larger pattern of such comparisons. Given a sequence $d = (d_1, \dots, d_r)$, we say that d' is a *swap* from d if there exist $1 \leq i < j \leq r$ so that $d'_i = d_j$, $d'_j = d_i$, and $d'_k = d_k$ for all $k \neq i, j$. This swap is *increasing* if $d_i > d_j$. We say that d is *below* d' , written $d \sqsubseteq d'$, if and only if d' can be obtained by applying a sequence of increasing swaps to d . Note that if $d \sqsubseteq d'$, then d and d' are equal when considered as multi-sets.

Definition 1.19. If $x = (x_1, \dots, x_r)$ and d is an r -tuple of non-negative integers, we define the *monomial* $[x]^d := \prod_{i=1}^r [x_i]^{d_i}$. We say $[x]^d$ is *below* $[x]^{d'}$, written $[x]^d \sqsubseteq [x]^{d'}$, iff $d \sqsubseteq d'$.

An easy argument on binomial coefficients shows that if $0 \leq p_1 \leq p_2$ and $0 \leq d_2 \leq d_1$ then $[p_1]^{d_1} [p_2]^{d_2} \leq [p_1]^{d_2} [p_2]^{d_1}$. Thus if p is weakly increasing and d' is an increasing swap of d , $[p]^d \leq [p]^{d'}$. These observations are succinctly stated in the following result.

Lemma 1.20. If $d \sqsubseteq d'$, then $[p]^d \leq [p]^{d'}$ for all weakly increasing sequences $p \in (\mathbb{Z}_+)^r$.

In Example 1.18, we had $m_{T,\gamma} \sqsubseteq m_{T,e}$ for all T . For example, $m_{T',4132}(x) = [x]^d$ where $d = (2, 1, 0, 0)$ and $m_{T',e}(x) = [x]^{d'}$ where $d' = (2, 0, 1, 0)$. Since $d \sqsubseteq d'$, Lemma 1.20 implies our earlier observation that $m_{T',4132} \sqsubseteq m_{T',e}$.

We believe that a similar phenomenon of termwise inequality always holds, although in general it may be necessary to further permute the vertex labeled trees.

Conjecture 1.21 (The Rearrangement Conjecture). For all $r \geq 1$ and $\gamma \in S_r$ there exists a permutation τ of \mathcal{T}_r such that $m_{T,\gamma} \sqsubseteq m_{\tau(T),e}$, for all $T \in \mathcal{T}_r$.

Given such a bijection τ , we have

$$\varphi(p_\gamma, p_\gamma) = \sum_{T \in \mathcal{T}_r} m_{T,\gamma}(p) \leq \sum_{T \in \mathcal{T}_r} m_{\tau(T),e}(p) = \varphi(p, p)$$

for all weakly increasing p . Thus the Rearrangement Conjecture immediately implies a corresponding inequality for non-crossing pairings.

Conjecture 1.22 (The Weak Rearrangement Conjecture). Let $r \geq 1$. If $\gamma \in S_r$,

$$\varphi(p_\gamma, p_\gamma) \leq \varphi(p, p)$$

for all weakly increasing $p \in (\mathbb{Z}_+)^r$.

Either version of the Rearrangement Conjecture is sufficient to prove our main desired result.

Theorem 1.23. *The Rearrangement Conjecture implies the Main Conjecture.*

We have used Mathematica to computationally verify the Rearrangement Conjecture (and hence the Main Conjecture) for every $\gamma \in S_r$ with $r \leq 7$.

We conclude with another partial result that addresses certain special cases of the Main Conjecture. A sequence (a_1, \dots, a_r) is *unimodal* if there is a $1 \leq k \leq r$ such that $a_1 \leq a_2 \leq \dots \leq a_{k-1} \leq a_k \geq a_{k+1} \geq \dots \geq a_r$. Similarly, a permutation $\gamma \in S_r$ is *unimodal* if $(\gamma(1), \dots, \gamma(r))$ is unimodal. Given any permutation $\gamma \in S_r$, we define $\mathcal{M}_\gamma := \{m_{T,\gamma} : T \in \mathcal{T}_r\}$ to be the multiset of monomials appearing in $P_\gamma(p)$.

Theorem 1.24. *The Rearrangement Conjecture holds for any γ that is unimodal or a cyclic rotation of a unimodal permutation. In those cases, there exists a permutation of τ of \mathcal{T}_r with $m_{T,\gamma} = m_{\tau(T),e}$ for all $T \in \mathcal{T}_r$. For all other γ , we have $\mathcal{M}_e \not\subseteq \mathcal{M}_\gamma$.*

The rest of this paper is organized as follows. In Section 2, we prove some basic results about φ concerning its symmetries, recurrence relation, and (non-commutative) generating function. We also give several preliminary upper and lower bounds. In Section 3, we prove all the main results of our paper as outlined above, as well as some additional enumeration results that connect the polynomials $P_e(x)$ to certain classes of Dyck paths. We conclude in Section 4 with a brief discussion comparing the present work to other generalized Catalan structures, including the enumeration of monomials in certain algebraic expressions [2].

2. Basic results

2.1. *Symmetries*

The set of non-crossing pairings exhibits both rotational and reflective symmetry. Let $\text{Refl} := (2k, 2k - 1, \dots, 2, 1)$ and, for $1 \leq l \leq 2k$, let $\text{Rot}_l = (l, l + 1, \dots, k, 1, 2, \dots, l - 1)$ be the *left-rotation* by l . These permutations of S_{2k} extend to permutations of $NC_2(2k)$. If $\pi \in NC_2(2k)$ then $\text{Refl}(\pi) := \{\{\text{Refl}(i), \text{Refl}(j)\} : \{i, j\} \in \pi\} \in NC_2(2k)$; $\text{Rot}_l(\pi)$ is defined analogously. These permutations generate the (dihedral) automorphism group of the lattice $NC_2(2k)$ (cf. [13]). We also define similar operations on words.

Definition 2.1. If $w = w_1 w_2 \dots w_{2k}$, the *reflection* of w is $\text{Refl}(w) := w_{2k} \dots w_2 w_1$, for $1 \leq l \leq 2k$, the *left-rotation* of w by l is $\text{Rot}_l(w) := w_l w_{l+1} \dots w_{2k} w_1 w_2 \dots w_{l-1}$. The *negation* of w replaces each w_i by $1 - w_i$, and is denoted by \bar{w} .

Note that if $\pi \in NC_2(w)$, then $\text{Refl}(\pi)$ is compatible with $\text{Refl}(w)$, and similarly for Rot_l . Also note that switching the roles of 1s and 0s in w does not affect whether a pairing π is compatible with w . Thus the set of non-crossing pairings is preserved under all of these simple operations.

Proposition 2.2. *If $w = w_1 \dots w_n$ is a binary word and $1 \leq l \leq n$,*

$$\varphi(w) = \varphi(\text{Rot}_l(w)) = \varphi(\text{Refl}(w)) = \varphi(\bar{w}).$$

Proposition 2.2 makes it clear that it is natural to draw non-crossing pairings of binary words around a circle as in Fig. 2. However, we will mostly use linear representations of pairings (such as those in Fig. 8) while keeping Proposition 2.2 in mind.

2.2. Recursion formula

We consider $\varphi : \{0, 1\}^* \rightarrow \mathbb{N}$ as a function defined on binary words; it is clear that $\varphi(w) = 0$ if $|w|$ is odd. We let λ denote the *empty word*, the unique word of length 0. We consider the empty set to have exactly one pairing, the *empty pairing* $\pi_0 = \emptyset$, which is vacuously compatible with λ . Thus $\varphi(\lambda) = 1$.

One of the fundamental properties of φ is that it satisfies a quadratic recurrence formula.

Theorem 2.3. *If $w = w_1 \dots w_n$ is a binary word, then*

$$\varphi(w) = \sum_{j:w_j \neq w_1} \varphi(w_2 \dots w_{j-1})\varphi(w_{j+1} \dots w_n).$$

This formula is straightforwardly derived by partitioning $NC_2(w)$ into classes according to the value j such that $\{1, j\} \in \pi$.

Theorem 2.3 can also be expressed as a functional equation for the non-commutative generating function of φ . Let $\mathbb{C}\langle\langle x_0, x_1 \rangle\rangle$ be the ring of power series in the non-commuting indeterminates x_0, x_1 . The generating function for φ is the power series $F \in \mathbb{C}\langle\langle x_0, x_1 \rangle\rangle$ given by

$$F(x_0, x_1) = \sum_w \varphi(w)x_w$$

where $x_w := \prod_{i=1}^n x_{w_i}$.

Theorem 2.4. *$F = \sum_w \varphi(w)x_w$ is the unique solution in $\mathbb{C}\langle\langle x_0, x_1 \rangle\rangle$ to the functional equation*

$$F = 1 + x_1 F x_0 F + x_0 F x_1 F.$$

This result can be proven by partitioning pairs (w, π) with $\pi \in NC_2(w)$ according to the value of w_1 and the value of j such that $\{1, j\} \in \pi$. This result is also stated in Example 16.17 in [13], but is proven there by different means.

2.3. Path representations of words

Definition 2.5. Given a binary word $w = w_1 w_2 \dots w_n$, set $Y_0 = 0$, and $Y_i := \sum_{j=1}^i (-1)^{w_j+1}$ for $1 \leq i \leq 2k$. Define the points $P_i := (i, Y_i) \in \mathbb{R}^2$, and the corresponding *lattice path* of w , denoted by $\mathcal{P}(w) \in \mathbb{R}^2$, as the piecewise linear path consisting of the union of the n line segments $P_{i-1}P_i$ for $1 \leq i \leq 2k$ (i.e., 1s in w correspond to northeast moves, $(1, 1)$, and 0s to southeast moves, $(1, -1)$).

Definition 2.6. Given $w = w_1 w_2 \dots w_n$ as above, set $m := \min\{Y_0, Y_1, \dots, Y_n\}$. Then the *height* of w_i for $1 \leq i \leq 2k$ is defined to be the integer $h_i := \frac{1}{2}(Y_{i-1} + Y_i) + \frac{1}{2} - m$. We define $h(w) = \max\{h_i : 1 \leq i \leq n\}$ to be the *height* of the path $\mathcal{P}(w)$.

Note that h_i is the y -coordinate of the midpoint of the line segment $P_{i-1}P_i$ shifted so that $\{h_i : 1 \leq i \leq n\} = \{1, 2, \dots, h(w)\}$ (see Fig. 6).

Lattice path heights give a simple necessary and sufficient condition for the existence of non-crossing pairings with a specified pair.

Lemma 2.7. *Let w be a balanced binary word and let $1 \leq i < j \leq |w|$. Then there exists a $\pi \in NC_2(w)$ with $\{i, j\} \in \pi$ if and only if $w_i \neq w_j$ and $h_i = h_j$. In particular, $NC_2(w) \neq \emptyset$ if and only if w is balanced.*

Proof. Let $1 \leq i < j \leq n$. We claim that $w' = w_{i+1} w_{i+2} \dots w_{j-1}$ and $w'' = w_i w_{i+1} \dots w_j$ are both balanced if and only if $h_i = h_j$ and $w_i \neq w_j$. Each of these sets of conditions can be written as a

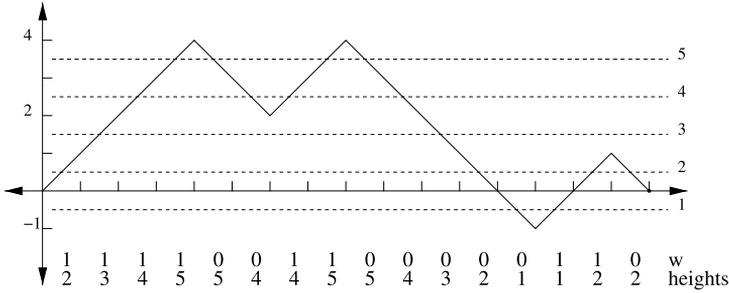


Fig. 6. The lattice path $\mathcal{P}(w)$ of the word $w = 1^4 0^2 1^2 0^5 1^2 0$ and the heights of the characters in w .

system of equations in the Y_i . The conditions that w' and w'' are balanced are equivalent to the equations $Y_i = Y_{j-1}$ and $Y_{i-1} = Y_j$, respectively. On the other hand the sets of conditions $h_i = h_j$ and $w_i \neq w_j$ are equivalent to the equations $Y_{i-1} + Y_i = Y_{j-1} + Y_j$ and $Y_i - Y_{i-1} = -(Y_j - Y_{j-1})$, respectively. It is easy to check that these two sets of equations are equivalent.

If $\{i, j\} \in \pi \in NC_2(w)$, then w' and w'' must both be balanced, and hence $h_i = h_j$ and $w_i \neq w_j$. For the other direction, we proceed by induction and use the fact that if $h_i = h_j$ and $w_i \neq w_j$, then w' and w'' are balanced. Since w and w'' are balanced, so is $w_0 = w_1 w_2 \dots w_{i-1} w_{j+1} \dots w_n$. By the inductive hypothesis, there then exist $\pi_0 \in NC_2(w_0)$ and $\pi' \in NC_2(w')$. Thus $\pi = \{\{i, j\}\} \cup \pi' \cup \pi_0 \in NC_2(w)$. (Some obvious re-labelings of the ground sets of π' and π_0 must be carried out to make this expression for π formally correct.)

For the final claim, if a non-empty word w is balanced, pick i so that $w_i \neq w_{i+1}$. Then $h_i = h_{i+1}$ and there exists $\pi \in NC_2(w)$ with $\{i, i + 1\} \in \pi$. Remove w_i, w_{i+1} from w to obtain a new, shorter balanced word w' and proceed inductively to construct the pairing. \square

Remark 2.8. Let w be a binary word of length $2k$. If for some $h \geq 1$, only two letters in w are at height h , say w_i and w_j , then every $\pi \in NC_2(w)$ must contain $\{i, j\}$.

Corollary 2.9. For any word w with a unique tallest peak (resp. lowest valley) we have $\varphi(w) = \varphi(\tilde{w})$, where \tilde{w} is the result of removing the tallest peak (resp. lowest valley) in w to level it with the second tallest peak (resp. second lowest valley).

2.4. Rough bounds

We conclude this section with a few upper and lower bounds on φ .

Proposition 2.10. Let $w = 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r}$ be a balanced word, and let i be the minimum block size, $i := \min\{p_1, q_1, \dots, p_r, q_r\} \geq 1$. Then

$$\varphi(w) \geq (1 + i)^{r-1}.$$

If $r = 1, 2$, then $\varphi(w) = (1 + i)^{r-1}$.

Proof. Clearly $\varphi(1^i 0^i) = 1$. Suppose $r \geq 2$. Without loss of generality we assume that $p_1 = i$, see Proposition 2.2. Since both $q_1, q_r \geq i$, for any $0 \leq \ell \leq i = p_1$ we may pair the last ℓ 1s in the block 1^{p_1} to the first ℓ 0s in the 0^{q_1} block and the remaining $i - \ell \leq q_r$ 1s to the last $i - \ell$ 0s in the 0^{q_r} block (see Fig. 7). The remaining word is then $0^{q_1 - \ell} 1^{p_2} \dots 0^{q_{r-1}} 1^{p_r} 0^{p_r - (i - \ell)}$, which can be rotated to

$$\tilde{w} = 1^{p_2} 0^{q_2} \dots 1^{p_r} 0^{q_1 + q_r - i}.$$

This is a balanced word with $r - 1$ runs, and with minimum run length $\tilde{i} = \min\{p_2, q_2, \dots, p_r, q_r + q_1 - i\} \geq i$. The inductive hypothesis then implies that $\varphi(\tilde{w}) \geq (1 + \tilde{i})^{r-2}$.

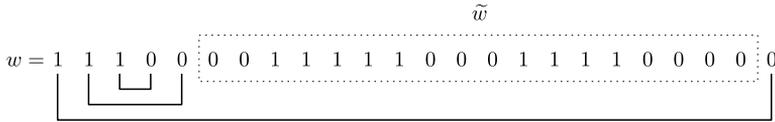


Fig. 7. A binary word whose first 1-block, 1^i , is the smallest, here $i = 3$. These 1s can be paired to the first and last blocks of 0s in exactly $i + 1$ ways. In this example, $\ell = 2$.

Thus for each choice of $0 \leq \ell \leq i$, we have at least $(1 + i)^{r-2}$ distinct pairings of w . Furthermore, pairings corresponding to different ℓ are distinct. This implies that $\varphi(w) \geq (1 + i)^{r-1}$ as claimed. When $r = 2$, the pairings counted are the only types possible, as there are only two blocks of 0s. Furthermore the remaining word is always a rotation of $\tilde{w} = 1^{p_2}0^{p_2}$, which has only one compatible pairing. Thus $\varphi(w) = 1 + i$ when $r = 2$. \square

The preceding inductive proof actually yields a somewhat larger lower bound. Let i_1, \dots, i_{r-1} be the minima defined by the inductive process in the proof of Proposition 2.10 (i.e. $i_1 = i$ is the global minimum in the proof and each i_{k+1} is the minimum of the block lengths in the leftover word after the inductive step has been applied at stage k (so $i_2 = \tilde{i}$ from the proof, and so on)). The following is a strengthening of Proposition 2.10.

Proposition 2.11. *Let w be defined as in Proposition 2.10 and i_1, \dots, i_{r-1} be defined as in the preceding paragraph. Then*

$$\varphi(w) \geq (1 + i_1) \dots (1 + i_{r-1}).$$

Remark 2.12. This bound is sharp, as demonstrated by applying Lemma 2.7 to the family of examples

$$w = 1^{a_1+1}0^{a_2}1^{a_2+1}0^{a_2+a_3}0^{a_3} \dots 1^{a_{r-1}+a_r}0^{a_r}1^{a_r+a_{r+1}}0^{a_1+a_2+\dots+a_r+a_{r+1}},$$

where the a_i are any positive integers.

In the other direction, we prove a simple upper bound (which is not sharp in general).

Proposition 2.13. *Let w be a binary word with height $h = h(w)$ and r runs. Then*

$$\varphi(w) \leq C_r^{(h)} \leq \frac{r^{r-1}}{r!} (1 + h)^{r-1}. \tag{2.1}$$

Proof. The proof relies on the following simple injection of pairings on w to pairings on $w' = (1^h 0^h)^r$. In w , the 1's in the block 1^{p_k} have successive heights $a, a + 1, \dots, a + p_k - 1$ for some a , and all heights are in the range $[1, h]$. The k th run of 1s in w' hits every height $1, \dots, h$, and thus we use the *height-preserving map* from w to w' (the situation for runs of 0s is inverted and analogous). Furthermore, we preserve the pairings of w when injecting into w' . If a run 1^{p_k} in w ends at position i with $h_i = a$, then the following run of 0s in w also begins at the same height $h_{i+1} = a$. This leaves excess bits $1^{h-a}0^{h-a}$ at the “top” of a run in w' , at heights $a + 1, \dots, h$, which we pair locally. Similarly, if a run of 0s in w ends at height b , then there will in general be excess bits $0^{b-1}1^{b-1}$ in w' that are also paired locally.

This gives the inclusion, and the first inequality then follows from Proposition 1.7. The second inequality is an elementary rough estimate of the Fuss–Catalan number, which is left to the reader. Fig. 8 demonstrates the inclusion. \square

Note that the lattice path height is the smallest h that can be used in the proof of Proposition 2.13, since all heights appearing in $\mathcal{P}(w)$ must appear in $\mathcal{P}((1^h 0^h)^r)$. Unfortunately, $h(w)$ can be quite large in comparison to the average (or even maximum) block size in w : consider the word

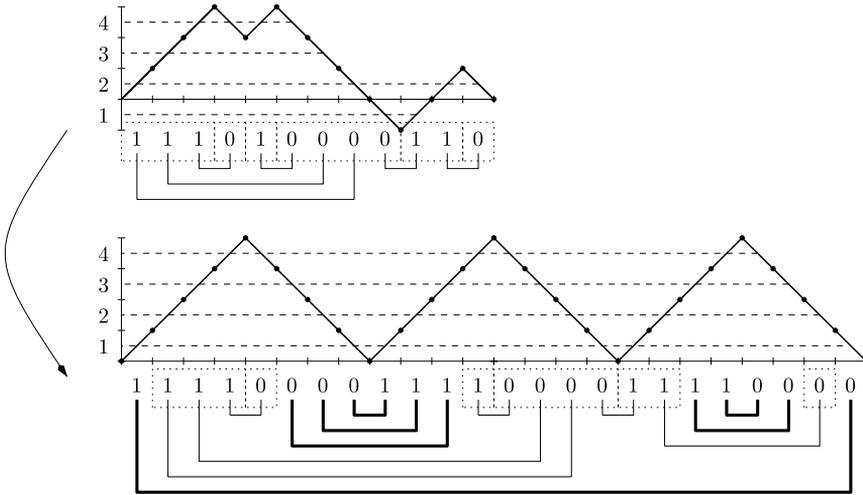


Fig. 8. w is injected into $(1^b 0^h)^r$, with extraneous labels (dark lines) paired locally.

$(1^k 0)^\ell (10^k)^\ell$. The maximum block size is k , while the lattice path height is $(k - 1)\ell + 1$. Indeed, this word has length $2(k + 1)\ell$, and the height is nearly half the total length. In general, a word of length $2k$ with r runs can have height $k - r + 1$, so the following corollary is essentially the best that can be said using height considerations.

Corollary 2.14. *Let w be a balanced word of length $2k$ with r runs. Then*

$$\varphi(w) \leq \frac{r^{r-1}}{r!} (1 + k)^{r-1}. \tag{2.2}$$

Remark 2.15. The bound in Corollary 2.14 is not tight but for fixed r is of the right order in k . Theorem 1.9 implies that k may be replaced with k/r in this bound.

3. Proofs of main results

3.1. Overview

We prove our results by constructing injective maps from the set of non-crossing pairings into sets of related combinatorial structures whose properties are easier to manage. In Section 3.2 we define an injective mapping from $NC_2(w)$ into a certain class of (edge) labeled trees. The properties of this map together with a straightforward enumeration of the trees allow us to prove Theorems 1.6 and 1.17.

In Section 3.3, we define an injective mapping from the pairings enumerated by $\varphi(p, q)$ into certain classes of lattice paths. The results we obtain on these types of paths allow us to prove Theorems 1.9 and 1.23. We conclude the section with the proof of Theorem 1.24.

3.2. Injection from $NC_2(w)$ into edge labeled trees

In Section 2.4 we saw that the number of pairings enumerated by $\varphi(p, q)$ is strongly dependent on the minimum run length (cf. Proposition 2.10), and we also have remarked repeatedly on the rotational invariance of non-crossing pairings. In this section we construct injections from non-crossing pairings into edge-labeled trees, which allow us to encode the successive minima and rotational structure very easily.

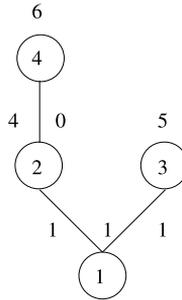


Fig. 9. T with a (p, γ) edge-labeling L where $p = (5, 3, 6, 4)$ and $\gamma = (3, 1, 4, 2)$. $L(1) = (1, 1, 1)$ has weight $W(1) = p_{\gamma^{-1}\gamma_T(1)} = p_2 = 3$. Each vertex v is depicted containing the label $\gamma_T(v)$. Note that $\gamma_T = (1, 2, 4, 3)$.

Before beginning, we must also note that our perspective and terminology here is somewhat “inverted” from our earlier presentation of Theorem 1.17, where we began with an ordered, weakly increasing multiset p and considered words p_γ that arose from permuting these run lengths (there, γ referred to such a permutation). In this section we instead begin with an arbitrary sequence p of runs and consider the permutations γ that might have led to such a sequence; i.e., such γ s that $p_{\gamma^{-1}}$ is weakly increasing. Such a permutation may not be uniquely defined, which is the reason for much of the intricate and technical notation in this section.

The labeled trees will be built from balanced words in which the runs of 1s are specified, but the 0s are arbitrarily distributed. If $r \geq 1$ and $p \in (\mathbb{Z}_+)^r$, then the set of p -words is defined to be

$$W_p := \{0^a 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r - a} : w \text{ is balanced, } q_i \geq 0 \text{ for all } 1 \leq i \leq r, \text{ and } 0 \leq a \leq q_r\}.$$

Given $w = 0^a 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r - a} \in W_p$, we define the subword 1^{p_i} in this representation to be the i th 1-block of w for $1 \leq i \leq r$. We do this even if some 1^{p_i} are adjacent in w , i.e. even if some q_i s are 0.

We next define a set of related trees that are also determined by the vector p . If $d \geq 0$ and $W > 0$ are integers, then a label of degree d and weight W is a $(d + 1)$ -tuple of integers $\ell = (\ell_0, \ell_1, \dots, \ell_d)$ such that (i) $\ell_0, \ell_d \geq 0$, (ii) $\ell_i > 0$ for all $0 < i < d$, and (iii) $|\ell| = W$. For a plane tree $T \in \mathcal{T}_r$ and permutation $\gamma \in S_r$, an edge (p, γ) -labeling of T is a function L mapping each $v \in T$ to a label $L(v)$ of degree $d_T(v)$ and weight $W(v) = p_{\gamma^{-1}(\gamma_T(v))}$ (recall the vertex labeling from Definition 1.13). Note that a tree T does not necessarily need to have an edge (p, γ) -labeling as labels on vertices of degree d must have weight $W \geq d - 1$. We will consider $v \in T$ to be additionally labeled with $\gamma_T(v)$ as before.

Definition 3.1. Adopting the preceding notation, the set of edge (p, γ) -labeled trees is

$$LT(p, \gamma) := \{(T, L) : T \in \mathcal{T}_r, L \text{ is a } (p, \gamma)\text{-labeling of } T\}.$$

Remark 3.2. Properly speaking, a (p, γ) labeling is not on the edges of a tree T , but rather can be thought of as lying in the “gaps” between edges. See Fig. 9 for an example of such a labeling and see Figs. 10–13 for the context in which these trees arise.

For a fixed $w \in W_p$ and a permutation $\gamma \in S_r$, we now define a map from non-crossing pairings to edge-labeled trees that we denote by $LT(w, \gamma, \cdot) : NC_2(w) \rightarrow LT(p, \gamma)$. For each $\pi \in NC_2(w)$, the map is defined by the following recursive construction of an edge-labeled tree $LT(w, \gamma, \pi) \in LT(p, \gamma)$. As this construction is somewhat involved we first give an overview and then the complete details immediately following. Initially, the i th 1-block of w is labeled $\gamma(i)$, which we call its γ -label. As the construction of $LT(w, \gamma, \pi)$ proceeds w is repeatedly rotated and divided into subwords. Individual 1-blocks are never divided by these processes and we view each 1-block maintaining its identity and its γ -label throughout the process. We first rotate w and π so that in the resulting w' and π' ,

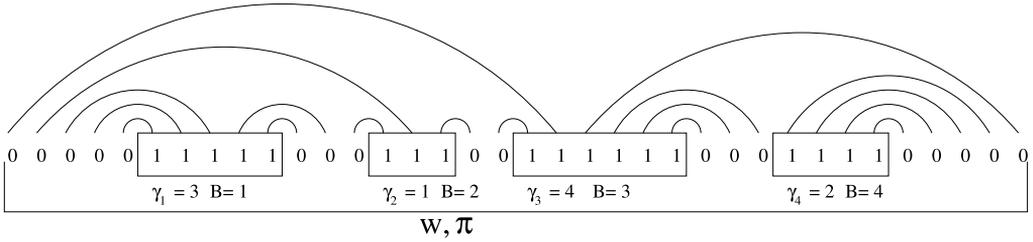


Fig. 10. $w \in W_p$ and $\pi \in NC_2(w)$, where $p = (5, 3, 6, 4)$, $\gamma = (3, 1, 4, 2)$. Here and in following figures, the numbers below each 1-block are its original block number B in w and its γ -label.

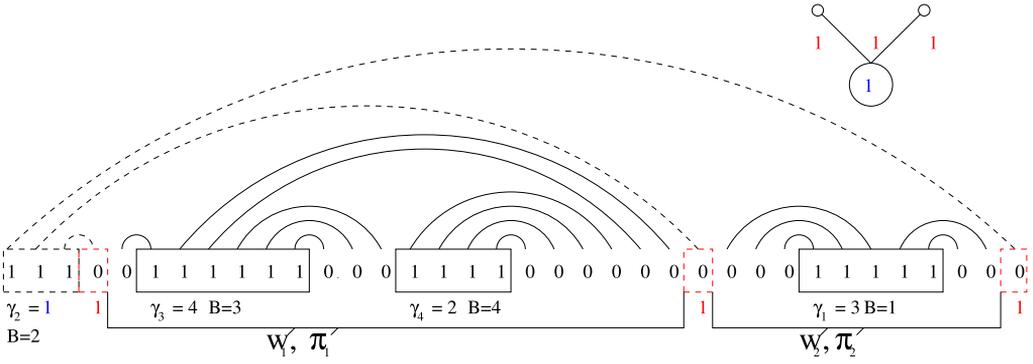


Fig. 11. $LT(w, \gamma, \pi)$ after one round of recursion. $L(1) = (1, 1, 1)$. Removed characters are shown in dotted boxes.

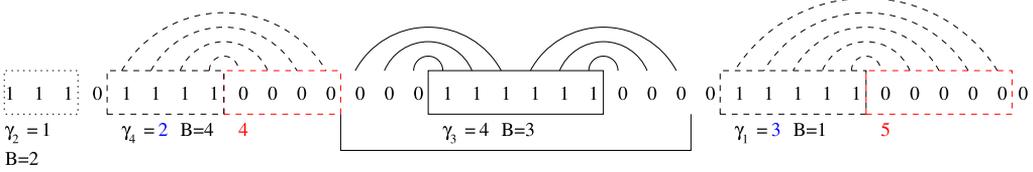
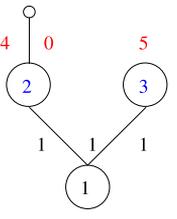


Fig. 12. $LT(w, \gamma, \pi)$ after two rounds of recursion. $L(2) = (4, 0)$ and $L(4) = (5)$.

w' starts with the j th 1-block where $\gamma(j) = \min \gamma$. We remove this block and the 0s paired to it by π' to get contiguous subwords w'_i with restricted pairings π'_i for $i = 1, \dots, d$. The 1-blocks of each w'_i inherit their γ -label from w' . Hence, $LT(w, \gamma, \pi)$ is a tree T with root u with a label $L(u)$ of weight p_j that records how the paired 0's interleave with the w'_i . Each π'_i is recursively recorded as the i th subtree pendant from u and its L -labeling. We now give the complete details.

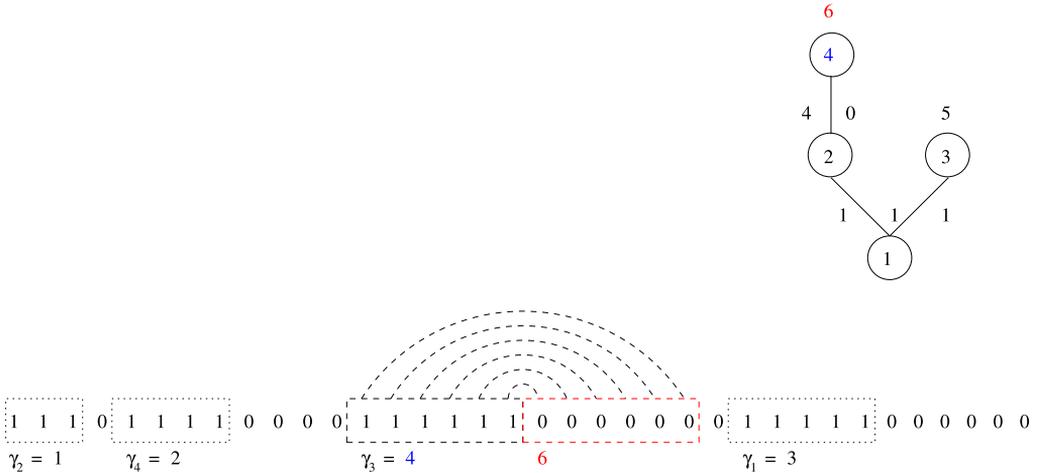


Fig. 13. $LT(w, \gamma, \pi)$, completed. $L(3) = (6)$. L is an edge (p, γ) -labeling whose weight on vertex i (in canonical order) is $W(i) = p_{\gamma^{-1}\gamma_T(i)}$.

Choose $1 \leq j \leq r$ so that $\gamma(j) = 1$, and define the rotations $\gamma' := \text{Rot}_j(\gamma)$ and $p' := \text{Rot}_j(p)$. Furthermore, define the rotation $\pi' := \text{Rot}_s(\pi)$, where s is the integer such that $w' = \text{Rot}_s(w)$ begins with the 1^{p_j} block of w .

We can now write down the decompositions of w' and π' :

$$w' = 1^{p_j} 0^{\ell_0} w'_1 0^{\ell_1} w'_2 0^{\ell_2} \dots 0^{\ell_{d-1}} w'_d 0^{\ell_d},$$

$$\pi' = \pi'_0 \cup \pi'_1 \cup \dots \cup \pi'_d,$$

where π'_0 pairs (some of) the 1s in the 1^{p_j} of w' to the 0s in 0^{ℓ_i} for $0 \leq i \leq d$ (so $\ell_0 + \dots + \ell_d = p_j$), and $\pi'_i \in NC_2(w'_i)$ for $1 \leq i \leq d$. Recall that both ℓ_0 and ℓ_d may be 0. The pairings π'_i are simply the restriction of π' to the subwords w'_i . Note that in the above procedure we have $d = 0$ if and only if $r = 1$. In this case, $w' = 1^{p_1} 0^{p_1}$, all 0s are paired to 1^{p_1} , and $\ell_0 = \ell_d = p_1$.

We now create the root vertex in $LT(w, \gamma, \pi)$ with d children and edge-label $\ell = (\ell_0, \dots, \ell_d)$. The weight of this label is $p_j = p_{\gamma^{-1}(1)} = p_{\gamma^{-1}(\gamma_T(1))}$, so this is the beginning of a valid (p, γ) labeling.

For $1 \leq i \leq d$, let $t_i > 0$ be the number of 1-blocks of w' that are contained in w'_i , and decompose γ' into corresponding components $\gamma'_i : [t_i] \rightarrow \mathbb{Z}$ so that

$$\gamma' = \gamma_{\min} \gamma'_1 \dots \gamma'_d.$$

With this definition, the inherited γ -labeling of the 1-blocks of w'_i is given by γ'_i , i.e. a 1-block in w'_i labeled according to γ'_i retains the same label it had in w . Note that when this procedure is invoked recursively we must temporarily re-index labels so that $\gamma'_i \in S_{t_i}$.

See Fig. 10 for an example of a pairing $\pi \in NC_2(w)$ and Fig. 11 for the decompositions of $\pi' \in NC_2(w')$. As can be seen in Fig. 11, the γ -labels of blocks remain invariant.

To complete the construction of the labeling, for $1 \leq i \leq d$, we recursively calculate $(T_i, L_i) = LT(w'_i, \gamma'_i, \pi'_i)$. Then $LT(w, \gamma, \pi) := (T, L)$, where T has decomposition (u, T_1, \dots, T_d) , and where

$$L(v) := \begin{cases} L(u) & \text{if } v = u, \\ L_i(v) & \text{if } v \in T_i. \end{cases}$$

Figs. 10–13 illustrate the complete calculation of $LT(w, \gamma, \pi)$ for the given example.

We say p and γ are *concordant* if and only if for all i, j , $\gamma(i) < \gamma(j)$ implies $p_i \leq p_j$. In such a case γ should be viewed as an encoding of a weakly increasing ordering of p . Note that every p has

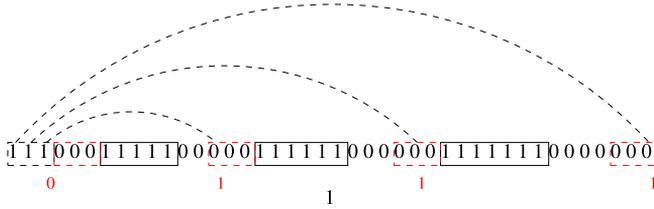


Fig. 14. A symmetric word with smallest 1-block coming first. A 1 in this block at height h can be paired to any 0 at height h . There is one such 0 in each 0-block. Once these 1s are paired respecting the non-crossing and height conditions, the resulting w'_i will again be symmetric.

at least one concordant γ , and that if p is weakly increasing, then p_γ is by definition concordant with γ .

The following result describes how concordance relates to bijective mappings.

Theorem 3.3. *Let $p \in (\mathbb{Z}_+)^r$. If $w \in W_p$ and $\gamma \in S_r$, then $LT(w, \gamma, \cdot) : NC_2(w) \rightarrow LT(p, \gamma)$ is injective. If, in addition, w is symmetric (or the rotation of a symmetric word) and γ is concordant with p , then $LT(w, \gamma, \cdot)$ is a bijection.*

Proof. It is clear that when $r = 1$, $LT(w, \gamma, \cdot)$ is bijective. In this case $w' = 1^{p_1}0^{p_1}$ and the mapping sends the only member π of $NC_2(w)$ to the only member of $LT(p, \gamma)$, a single root with label $L(1) = (p_1)$.

We now prove that $LT(w, \gamma, \cdot)$ is injective by induction on r . Suppose $\pi, \tau \in NC_2(w)$ and $LT(w, \gamma, \pi) = LT(w, \gamma, \tau) = (T, L)$. If T has the decomposition (u, T_1, \dots, T_d) and the root label is denoted $\ell = L(u)$, then we have the word decomposition $w' = 1^{p_j}0^{\ell_0}w'_10^{\ell_1}w'_20^{\ell_2} \dots 0^{\ell_{d-1}}w'_d0^{\ell_d}$, and pairing decomposition $\pi' = \pi'_0 \cup \pi'_1 \cup \dots \cup \pi'_d$, where $\gamma_j = 1$ and $\pi'_i \in NC_2(w'_i)$. Since $(T_i, L_i) = LT(w'_i, \gamma'_i, \pi'_i)$, the number of 1-blocks in w'_i is $|T_i|$. This means the w'_i are completely determined by w, γ and (T, L) . Thus $\tau' = \tau'_0 \cup \tau'_1 \cup \dots \cup \tau'_d$, where $\tau'_i \in NC_2(w'_i)$. We must have $\tau'_0 = \pi'_0$ since these pairs match the 1s in 1^{p_j} to the same set of 0s in w' . Since $LT(w'_i, \gamma'_i, \pi'_i) = LT(w'_i, \gamma'_i, \tau'_i) = (T_i, L_i)$ we have $\pi'_i = \tau'_i$ by induction and thus $\pi = \tau$.

We now prove that $LT(w, \gamma, \cdot)$ is bijective when (i) w is the rotation of a symmetric word, and (ii) γ is concordant with p . We proceed by induction on r , and note that the base case $r = 1$ has already been shown. If $\gamma(j) = 1$, then since p and γ are concordant, we have $p_j = \min\{p_i : 1 \leq i \leq r\}$, and as w is symmetric,

$$w' = 1^{p_j}0^{p_j}1^{p_{j+1}}0^{p_{j+1}} \dots 1^{p_r}0^{p_r}1^{p_1}0^{p_1} \dots 1^{p_{j-1}}0^{p_{j-1}}.$$

This means that the last p_j 0s in each 0-block of w' are at heights $p_j, p_j - 1, \dots, 2, 1$. Thus any 1 in 1^{p_j} of height h can be paired to any 0 of height h in any 0-block; in other words, any labeling ℓ can appear on the root vertex in the image of $LT(w, \gamma, \cdot)$. Consider the decomposition $w' = 1^{p_j}0^{\ell_0}w'_10^{\ell_1}w'_20^{\ell_2} \dots 0^{\ell_{d-1}}w'_d0^{\ell_d}$. If 0^{ℓ_i} occurs in the k th 0-block of w' and $0^{\ell_{i+1}}$ in the l th, (where $l > k$) then $w'_i = 0^{\ell_{i+1} + \dots + \ell_d}1^{p_{k+1}}0^{p_{k+1}} \dots 1^{p_l}0^{p_l - \ell_{i+1} - \ell_{i+2} - \dots - \ell_d}$ and thus w'_i is the rotation of a symmetric word (see Fig. 14). Since γ'_i is also concordant with p'_i , the maps $LT(w'_i, \gamma'_i, \cdot) : NC_2(w'_i) \rightarrow LT(p'_i, \gamma'_i)$ are bijective by induction.

Given $(T, L) \in LT(p, \gamma)$ we now construct π so that $LT(w, \gamma, \pi) = (T, L)$. Let T have decomposition (u, T_1, \dots, T_d) . Let $L(u) = \ell$, a label of degree d and weight p_j . Pick a partial pairing π'_0 on 1^{p_1} that achieves a decomposition $w' = 1^{p_j}0^{\ell_0}w'_10^{\ell_1}w'_20^{\ell_2} \dots 0^{\ell_{d-1}}w'_d0^{\ell_d}$ so that each w'_i contains $|T_i|$ 1-blocks. The restricted edge-labelings satisfy $(T_i, L_i) \in LT(p'_i, \gamma'_i)$, and so by induction there exists $\pi'_i \in NC_2(w'_i)$ with $LT(w'_i, \gamma'_i, \pi'_i) = (T_i, L_i)$. Thus if $\pi' = \pi'_0 \cup \pi'_1 \cup \dots \cup \pi'_d$, then $LT(w, \gamma, \pi) = (T, L)$. \square

Note that a simple counting argument shows that there are $[W]^d := \binom{W+1}{d}$ possible labels $\ell = (\ell_0, \dots, \ell_d)$ of weight W and degree d . This means that edge-labeled trees can be enumerated by the tree polynomials of Section 1.

Lemma 3.4. If $p \in (\mathbb{Z}_+)^r$ and $\gamma \in S_r$, then $|LT(p, \gamma)| = P_\gamma(p_{\gamma^{-1}})$.

Proof. Suppose that $T \in \mathcal{T}_r$. In any (p, γ) -labeling of T , the weight of the label of vertex i is by definition $p_{\gamma^{-1}(\gamma_T(i))}$, and the degree is $d_T(i)$. Thus there are $m_{T,\gamma}(p_{\gamma^{-1}}) = \prod_{v \in T} [p_{\gamma^{-1}(\gamma_T(v))}]^{d_T(v)}$ possible (p, γ) -labelings of T . Summing over all of \mathcal{T}_r gives the result. \square

Finally, Theorem 3.3 immediately enables us to translate formulas for edge-labeled trees to formulas and bounds for non-crossing pairings.

Proof of Theorem 1.6. Let γ be concordant with p . By Theorem 3.3 we have

$$\varphi(p, q) \leq |LT(p, \gamma)| = \varphi(p, p). \quad \square$$

Proof of Theorem 1.17. Suppose that p is weakly increasing and that $\gamma \in S_r$. Then by definition p_γ is concordant with γ , so Theorem 3.3 and Lemma 3.4 give

$$\varphi(p_\gamma, p_\gamma) = |LT(p_\gamma, \gamma)| = P_\gamma((p_\gamma)_{\gamma^{-1}}) = P_\gamma(p). \quad \square$$

We close with an alternative proof for Theorem 1.6 that injectively maps the non-crossing pairings counted by $\varphi(p, q)$ to those counted by $\varphi(p, p)$ without using labeled trees or any other auxiliary combinatorial structures. This proof was provided to us by one of the anonymous referees, who has graciously allowed us to reproduce it here.

Second proof of Theorem 1.6. Recall that for $p \in \mathbb{Z}_+^r$ fixed, W_p is the set of all balanced words

$$0^a 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r - a},$$

where each $q_i \geq 0$ (and $q_i = 0$ is allowed) and where $0 \leq a \leq q_r$. Let W'_p be those words in W_p with $q_i = p_i$ for all i .

The end goal is to show that if $w \in W_p$ and $w' \in W'_p$, then there is an injective map from $NC_2(w)$ to $NC_2(w')$. The proof is by induction on $n = |w|$, and is trivial for $n = 0$. Rotate w so that block 1^{p_i} comes first, where i is the first block in w with $p_i = \min p$. Rotate w' so that 1^{p_i} also comes first in w' , and rename w, w', p to be these newly rotated copies. For all $s \leq n$, let $NC_2(w)|_s$ be the set of pairings π in $NC_2(w)$ with $\{p_1, s\} \in \pi$. If $NC_2(w)|_s$ is non-empty then $w_s = 0$, and we argue that we can conclude that $w'_s = 0$ as well.

Indeed, since p_1 (formerly p_i) satisfies $p_1 = \min p$, let j be such that the second through j th 1-blocks of w are contained in $\omega_1 := w_{p_1+1} \dots w_{s-1}$. For convenience, we also extend this definition of j to the case that ω_1 is empty: if $\omega_1 = \lambda$, then set $j = 1$. Let $\omega_2 := w_1 \dots w_{p_1-1} w_{s+1} \dots w_n$. Since $\{p_1, s\} \in \pi$ for some $\pi \in NC_2(w)$, we must have ω_1, ω_2 balanced. Thus $\omega_1 \in W_{p_2, \dots, p_j}$, $\omega_2 \in W_{p_1-1, p_{j+1}, \dots, p_r}$, and $s = p_1 + 2(p_2 + \dots + p_j) + 1$. In w' the block 0^{p_j} ranges from character $2(p_1 + \dots + p_{j-1}) + p_j + 1$ to character $2(p_1 + \dots + p_j)$. Since this range contains s (due to the minimality of p_1), $w'_s = 0$ as claimed. Let $\omega'_1 := w'_{p_1+1} \dots w'_{s-1}$ and $\omega'_2 := w'_1 \dots w'_{p_1-1} w'_{s+1} \dots w'_n$. Clearly $\omega'_1 \in W'_{p_2, \dots, p_j}$. Since w' is symmetric and $p_1 = \min p$ we also have $\omega_2 \in W'_{p_1-1, p_{j+1}, \dots, p_r}$. A pairing $\pi \in NC_2(w)|_s$ restricts to $\pi_1 \in NC_2(\omega_1)$ and $\pi_2 \in NC_2(\omega_2)$. By the inductive hypothesis, π_1 and π_2 may be injectively mapped to $\pi'_1 \in NC_2(\omega'_1)$ and $\pi'_2 \in NC_2(\omega'_2)$. In turn, π'_1 and π'_2 are the restrictions of a unique element $\pi' \in NC_2(w')|_s$ and thus $NC_2(w)|_s$ maps injectively into $NC_2(w')|_s$ for all s . The proof is complete. \square

3.3. Injection from $NC_2(w)$ into lattice paths

We now define a map from non-crossing pairings to a certain class of lattice paths. The edge-labeled trees were very useful for managing the successive minima and rotations that arose in the enumeration of pairings on symmetric words, but to prove our remaining inequalities we need to be

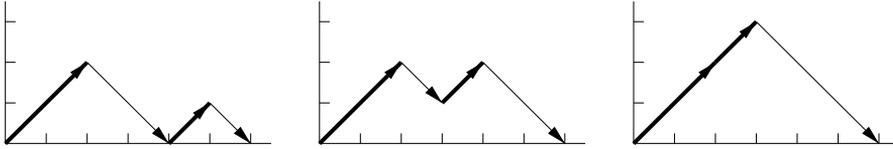


Fig. 15. The three paths in $\text{Path}(2, 1)$, namely $\mathcal{P} = (p, q)$ with $p = (2, 1)$ and $q = (2, 1), (1, 1),$ or $(0, 3)$. The northeast steps are shown in bold.

able to easily compare pairings on different words. While this is cumbersome using labeled trees, lattice paths have a natural ordering that makes such comparisons straightforward.

A Dyck path is a finite walk in \mathbb{Z}^2 taking steps of the form $(1, 1)$ or $(1, -1)$ that starts at $(0, 0)$, visits no point below the x -axis, and ends on the x -axis. If $p, p' \in \mathbb{Z}^r$ we say p is dominated by p' if $\sum_{i=1}^j p_i \leq \sum_{i=1}^j p'_i$ for all $1 \leq j \leq r$. This is denoted by $p \preceq p'$. Recall from Definition 2.5 that if w is a binary word, then $\mathcal{P}(w)$, the lattice path of w , is the walk in \mathbb{Z}^2 starting at $(0, 0)$ and with i th step equal to $(1, 1)$ if $w_i = 1$, or equal to $(1, -1)$ if $w_i = 0$.

For a fixed $p \in (\mathbb{Z}_+)^r$, we again consider a class of p -words $w = 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r}$, where q is an r -tuple of non-negative integers. It is easy to show that $\mathcal{P}(w)$ is a Dyck path if and only if $q \preceq p$ and $|q| = |p|$. In this case we say that $\mathcal{P}(w)$ is a p -path and w is a p -Dyck word (note that these are more restricted than the p -words of Section 3.2). Let

$$\text{Dyck}(p) := \{1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r} : \forall i \ q_i \geq 0, \ q \preceq p, \ |q| = |p|\}$$

be the set of p -Dyck words, and let $\text{Path}(p)$ denote the set of all p -paths. See Fig. 15. We will denote both a word $w = 1^{p_1} 0^{q_1} \dots 1^{p_r} 0^{q_r}$ and the corresponding path $\mathcal{P}(w)$ by the pair (p, q) .

We now recursively define a map $F : \text{LT}(p, e) \rightarrow \text{Dyck}(p)$ from edge-labeled trees to p -Dyck words. If $T = u$ and $L(u) = (p_1)$, then $F(T, L) = 1^{p_1} 0^{p_1}$. Otherwise, suppose $(T, L) \in \text{LT}(p, e)$, where T has the decomposition (u, T_1, \dots, T_d) and where the subtree T_i has labeling L_i . Also, write $\ell := L(u)$ for the root label. Apply the map recursively to obtain $w_i = F(T_i, L_i)$ for $1 \leq i \leq d$; then

$$F(T, L) := 1^{p_1} 0^{\ell_0} w_1 0^{\ell_1} \dots 0^{\ell_{d-1}} w_d 0^{\ell_d}.$$

Recall that by definition the weight of the root label is $|\ell| = p_1$.

Lemma 3.5. *If $p \in (\mathbb{Z}_+)^r$, then the map $F : \text{LT}(p, e) \rightarrow \text{Dyck}(p)$ is a bijection.*

Proof. This is clear when $r = 1$. The sole member of $\text{LT}(p, e)$ is a tree with root u and label $\ell = (p_1)$. Applying F results in $1^{p_1} 0^{p_1}$, the sole member of $\text{Dyck}(p)$. It is then clear that in the general case $w = F(T, L)$ is always a p -Dyck word, as the $w_i = F(T_i, L_i)$ are recursively Dyck words, and thus $w = 1^{p_1} 0^{\ell_0} w_1 0^{\ell_1} \dots 0^{\ell_{d-1}} w_d 0^{\ell_d}$ is as well. $F(T, L)$ is a p -Dyck word as its i th 1-block is $1^{p'_i}$.

To show that this is a bijection, observe that every $w \in \text{Dyck}(p)$ has a unique decomposition of the form

$$w = 1^{p_1} 0^{\ell_0} w_1 0^{\ell_1} \dots 0^{\ell_{d-1}} w_d 0^{\ell_d},$$

where ℓ is a label of degree d and weight p_1 , and each w_i is a p' -Dyck word for some subsequence p' of p . Indeed, for all $1 \leq i \leq p_1$, the i th 0 shown in the above decomposition is the first 0 in w of height $p_1 + 1 - i$. Let $1 = k_0 < k_1 < \dots < k_d = r$ be the indices such that $w_i \in \text{Dyck}(p'_i)$ where $p'_i = (p_{k_{i-1}+1}, \dots, p_{k_i})$. By induction $\text{LT}(p'_i, e)$ is in bijective correspondence with $\text{Dyck}(p'_i)$. Thus trees $(T, L) \in \text{LT}(p, e)$ of the form $T = (u, T_1, \dots, T_d)$ with label $L(u) = \ell$ are in bijective correspondence with words of the form $1^{p_1} 0^{\ell_0} w_1 0^{\ell_1} \dots 0^{\ell_{d-1}} w_d 0^{\ell_d}$; considering all possible labels ℓ of weight p_1 and of arbitrary degree gives the claim. \square

Composing this bijection with the map from Theorem 3.3 gives a map from non-crossing pairings to Dyck paths.

Theorem 3.6. Suppose that $p \in (\mathbb{Z}_+)^r$. If $w \in W_p$, then the map $P(w, \cdot) : NC_2(w) \rightarrow \text{Dyck}(p)$ given by $P(w, \pi) = F(LT(w, e, \pi))$ is an injection. If p is weakly increasing and w is symmetric, then it is a bijection.

Thus Theorem 1.17 gives an enumeration for paths.

Lemma 3.7. If $p \in (\mathbb{Z}_+)^r$ is weakly increasing, then $|\text{Path}(p)| = P_e(p)$.

Now that we have related Dyck paths to non-crossing pairings, we provide a simple comparison criterion for the number of paths associated to different vectors p .

Lemma 3.8. For all $r \geq 1$ and $p, p' \in (\mathbb{Z}_+)^r$, if $p \preceq p'$ then

$$|\text{Path}(p)| \leq |\text{Path}(p')|.$$

Proof. Since $p \preceq p'$, the difference $D := |p'| - |p| \geq 0$. We define an injection $\text{Path}(p) \rightarrow \text{Path}(p')$ by mapping $(p, q) \in \text{Path}(p)$ to (p', q') where $q' = q + (0, \dots, 0, D)$. Since $q \preceq p \preceq p'$, we have

$$\sum_{j=1}^i q'_j = \sum_{j=1}^i q_j \leq \sum_{j=1}^i p_j \leq \sum_{j=1}^i p'_j$$

for all $1 \leq i < r$. By construction $|q'| = |p'|$, so $q' \preceq p'$ and $(p', q') \in \text{Path}(p')$. This map is clearly injective, so the claimed inequality holds. \square

Note that if p' is weakly increasing, then $\varphi(p', p') = P_e(p')$ by Theorem 1.17, and thus if $p \preceq p'$, Theorem 3.6 and Lemmas 3.7 and 3.8 imply $\varphi(p, p) \leq |\text{Path}(p)| \leq |\text{Path}(p')| = \varphi(p', p')$.

Corollary 3.9. If $p' \in (\mathbb{Z}_+)^r$ is weakly increasing and $p \preceq p'$ then $\varphi(p, p) \leq \varphi(p', p')$.

Next we characterize a few key situations in which Corollary 3.9 applies.

Lemma 3.10. If $p, p' \in (\mathbb{Z}_+)^r$ and $|p| \leq |p'|$ then for some $1 \leq l \leq r$, $\text{Rot}_l(p) \preceq \text{Rot}_l(p')$.

Proof. We extend p, p' cyclically to all $i \in \mathbb{Z}$ by putting $p_i = p_j$ iff $i \equiv j \pmod{r}$. For $i \geq 1$ let $h_i = \sum_{j=1}^i (p'_j - p_j)$. Note $h_{i+r} - h_i = |p'| - |p| \geq 0$, for all $i \geq 1$. Thus there exists $0 \leq k < r$ so that $h_k = \min\{h_i : i \geq 1\}$. Thus $h_k \leq h_{k+i}$ for $1 \leq i \leq r$, or equivalently, $\sum_{j=k+1}^{k+i} (p'_j - p_j) \geq 0$ for $1 \leq i \leq r$. This means that $\text{Rot}_{k+1}(p) \preceq \text{Rot}_{k+1}(p')$. \square

Let $1 \leq r \leq k$. The Main Conjecture claims that $\text{Max}_{k,r} = \varphi(P, P)$ where P is the r -tuple $P = (m - 1, \dots, m - 1, m, \dots, m)$ with $m = \lceil \frac{k}{r} \rceil$ and $|P| = k$. The Main Theorem states that $\text{Max}_{k,r} \leq \varphi(P', P')$ where P' is the r -tuple (m, \dots, m) . If $|p| = k$, we relate p to P and P' in order to use the above results.

Lemma 3.11. Suppose $1 \leq r \leq k$. If $p \in (\mathbb{Z}_+)^r$ is weakly increasing and $|p| = k$ and P, P' are defined as in the preceding paragraph, then $p \preceq P \preceq P'$.

Proof. It suffices to show $p \preceq P$; by definition we have $|P| = |p|$. Suppose that $p \not\preceq P$ and let j be the smallest integer $1 \leq j < r$ for which $\sum_{i=1}^j p_j > \sum_{i=1}^j P_j$. Note that we must have $p_j > P_j$, and hence $p_{j+1} \geq p_j \geq P_j + 1 \geq m \geq P_i$ for all i . But then

$$|p| = \sum_{i \leq j} p_i + \sum_{i > j} p_i \geq \sum_{i \leq j} p_i + (r - j)p_{j+1} > \sum_{i \leq j} P_i + \sum_{i > j} P_j = |P|,$$

a contradiction. \square

Proof of Theorems 1.9 and 1.23. To prove the Main Theorem, let $p = (p_1, \dots, p_r)$ have weight $|p| = k$. By construction $|p| \leq |P'|$, and thus by Lemma 3.10 there is a rotation γ such that $p_\gamma \preceq P'_\gamma = P'$. Thus $\varphi(p, q) \leq \varphi(p, p) = \varphi(p_\gamma, p_\gamma) \leq \varphi(P', P')$; the first inequality is due to Theorem 1.6, the middle equality is due to the rotational invariance of φ , and the last inequality comes from Corollary 3.9. It is in this last step that it is essential to use P' rather than P (as we would like), as Corollary 3.9 requires weakly increasing sequences.

As for Theorem 1.23, suppose that the Rearrangement Conjecture is true. Then we may assume $\varphi(p, p) \leq \varphi(\tilde{p}, \tilde{p})$, where \tilde{p} is the weakly increasing rearrangement of p . But then by Theorem 1.6, Lemma 3.11, and Corollary 3.9 we have $\varphi(p, q) \leq \varphi(p, p) \leq \varphi(\tilde{p}, \tilde{p}) \leq \varphi(P, P)$, and the Main Conjecture follows. \square

For the next proof we will need the following result on identities involving multivariable polynomials.

Proposition 3.12 (The Combinatorial Nullstellensatz). (See Theorem 1.2, [1].) Let F be an arbitrary field, and let $f = f(x_1, \dots, x_n)$ be a polynomial in $F[x_1, \dots, x_n]$. Suppose the degree $\deg(f)$ of f is $\sum_{i=1}^n t_i$, where each t_i is a non-negative integer, and suppose the coefficient of $\prod_{i=1}^n x_i^{t_i}$ in f is non-zero. If S_1, \dots, S_n are subsets of F with $|S_i| > t_i$, then there is an $s \in S_1 \times \dots \times S_n$ so that $f(s) \neq 0$.

Our results give the following interesting polynomial recurrences.

Theorem 3.13. If $p = (p_1, \dots, p_r) \in (\mathbb{Z}_+)^r$ and $p' = (p_1, \dots, p_i + 1, p_{i+1} - 1, \dots, p_r)$, then

$$|\text{Path}(p')| = |\text{Path}(p)| + |\text{Path}(p_1, \dots, p_{i-1}, p_i)| \cdot |\text{Path}(p_{i+1} - 1, p_{i+2}, \dots, p_r)|.$$

Furthermore, if $x = (x_1, \dots, x_r)$ is an r -tuple of indeterminates and $x' := (x_1, \dots, x_i + 1, x_{i+1} - 1, \dots, x_r)$, then

$$P_e(x') = P_e(x) + P_e(x_1, \dots, x_{i-1}, x_i)P_e(x_{i+1} - 1, x_{i+2}, \dots, x_r).$$

Proof. Note that $p \preceq p'$. Recall the injective map $I : \text{Path}(p) \rightarrow \text{Path}(p')$ from the proof of Lemma 3.8, and consider any $(p', q') \in X := \text{Path}(p') \setminus I(\text{Path}(p))$, i.e. a p' -path that is not the image of a p -path (p, q) . This means that (p, q') is not a p -path. This can only happen if $p_1 + \dots + p_i = q'_1 + \dots + q'_i - 1$, as every other truncation of p and p' have the same sum. Thus if $a = (p_1, \dots, p_{i-1}, p_i)$ and $b = (p_{i+1} - 1, p_{i+2}, \dots, p_r)$, then $(a, (q'_1, \dots, q'_i - 1))$ is an a -path and $(b, (q'_{i+1}, \dots, q'_r))$ is a b -path. Conversely, any a -path and b -path can be recombined to get (p', q') in X . This proves the first equation.

By Lemma 3.7 we have the second equation for all $x \in (\mathbb{Z}_+)^r$ with x, x' both weakly increasing. Since $P_e(x)$ is polynomial with fixed total degree of $r - 1$, the polynomials are identical (apply the Combinatorial Nullstellensatz with $S_i = \{2ir + 1, 2ir + 2, \dots, 2ir + r\}$). \square

3.4. The unimodal case of the Rearrangement Conjecture

We conclude this section by proving Theorem 1.24. Without loss of generality (cf. Proposition 2.2), we assume $\gamma : [r] \rightarrow \mathbb{Z}_+$ is injective and unimodal. It is easy to see that any consecutive subsequence of γ is also unimodal; in particular, the smallest element of any subsequence is either the first or last element. Define $\bar{\gamma}$ to be the increasing rearrangement of γ .

We use the more general setting of plane forests while describing unimodal labelings; recall that \mathcal{F}_r is the set of plane forests on r vertices (cf. Remarks 1.12 and 1.14). For all $1 \leq t \leq r$, let $\mathcal{F}_{r,t} := \{F \in \mathcal{F}_r : F \text{ has } t \text{ trees}\}$. For any forest F on r vertices, $m_{F,\gamma} := \prod_{v \in F} [x_{\gamma(v)}]^{df(v)}$ is a polynomial in the variables $\{x_i : i \in \mathbb{Z}_+\}$.

The following result shows that the forest polynomial for a unimodal labeling γ is equal to the polynomial for $\bar{\gamma}$ on a related forest.

Proposition 3.14. *Suppose that $1 \leq t \leq r$ and $F \in \mathcal{F}_{r,t}$. If $\gamma : [r] \rightarrow \mathbb{Z}$ is injective and unimodal, then there exists an $F' \in \mathcal{F}_{r,t}$ such that $m_{F',\gamma} = m_{F,\bar{\gamma}}$.*

Proof. The proof proceeds by induction on r . If $r = 1$ the statement is trivially true. Let $r \geq 2$ and assume the theorem has been proven for $1 \leq r' < r$. Suppose that $F \in \mathcal{F}_{r,t}$, where $1 \leq t \leq r$, and let d_0 denote the degree of the root of the first tree in F , so that

$$F = ((u, T_1, \dots, T_{d_0}), T_{d_0+1}, \dots, T_{d_0+t-1}).$$

If the root u is removed, denote the remaining forest by

$$G := F - u = (T_1, \dots, T_{d_0}, T_{d_0+1}, \dots, T_{d_0+t-1}) \in \mathcal{F}_{r-1,d_0+t-1}.$$

By inductive hypothesis there is a forest $G' = (T'_1, \dots, T'_{d_0+t-1}) \in \mathcal{F}_{r-1,d_0+t-1}$ such that $m_{G',\gamma'} = m_{G,e(\gamma')}$. Here the injective, unimodal sequence $\gamma' : [r-1] \rightarrow \mathbb{Z}$ is obtained from γ by removing $\gamma_0 := \min_i \gamma(i)$. As γ is unimodal, we must have $\gamma = \gamma_0 \gamma'$ or $\gamma = \gamma' \gamma_0$.

If $\gamma = \gamma_0 \gamma'$, then set $F' := ((u', T'_1, \dots, T'_{d_0}), T'_{d_0+1}, \dots, T'_{d_0+t-1})$ so that $G' = F' - u'$. Let $T'_0 := (u', T'_1, \dots, T'_{d_0})$ be the first tree of F' and write the decomposition $\gamma = \gamma_0 \tau' \tau''$ so that $|\gamma_0 \tau'| = |T'_0|$. Then we conclude that

$$\begin{aligned} m_{F',\gamma} &= m_{T'_0,\gamma_0 \tau'} m_{F'-T'_0,\tau''} = [x_{\gamma_0}]^{d_0} m_{T'_0-u',\tau'} m_{F'-T'_0,\tau''} \\ &= [x_{\gamma_0}]^{d_0} m_{G',\gamma'} = [x_{\gamma_0}]^{d_0} m_{G,\bar{\gamma}} = m_{F,\bar{\gamma}}. \end{aligned} \tag{3.1}$$

The third equality requires the inductive hypothesis for the existence of G' ; all other steps follow simply from the definitions of forest-labelings and forest polynomials.

Finally, in the case that $\gamma = \gamma' \gamma_0$, set $F' := (T'_1, \dots, T'_{t-1}, (u', T'_t, \dots, T'_{t+d-1}))$ and denote the final tree by $T'_0 := (u', T'_t, \dots, T'_{t+d-1})$. Now decompose $\gamma = \tau' \tau'' \gamma_0$ so that $|\tau'' \gamma_0| = |T'_0|$. Again we reach the desired conclusion, as

$$\begin{aligned} m_{F',\gamma} &= m_{F'-T'_0,\tau'} m_{T'_0,\tau'' \gamma_0} = m_{F'-T'_0,\tau'} ([x_{\gamma_0}]^{d_0} m_{T'_0-u',\tau''}) \\ &= [x_{\gamma_0}]^{d_0} m_{G',\gamma'} = [x_{\gamma_0}]^{d_0} m_{G,\bar{\gamma}} = m_{F,\bar{\gamma}}. \quad \square \end{aligned} \tag{3.2}$$

Remark 3.15. A simple, recursive procedure for constructing F' from F and γ may also be recovered from the preceding inductive proof of Proposition 3.14.

Proof of Theorem 1.24. By Proposition 3.14 there exists a map $\tau : \mathcal{T}_r \rightarrow \mathcal{T}_r$ such that $m_{\tau(T),\gamma} = m_{T,e}$ for all $T \in \mathcal{T}_r$. Since $m_{T,e} = [x]^d$ where d is the degree sequence of T , τ must be a bijection.

Now suppose γ is not a rotation of a unimodal sequence. Thus there exists a $1 \leq l \leq r$ such that no consecutive subsequence in $(\gamma(1), \dots, \gamma(r))$ consists of precisely $\{l+1, l+2, \dots, r\}$. Let T be the tree in \mathcal{T}_r whose root has l children, and whose largest child has $r-l-1$ children of its own. The degree sequence d of this tree has $d_1 = l$, $d_{l+1} = r-l-1$ and $d_i = 0$ for $i \neq 1, l+1$. We claim that \mathcal{M}_γ does not contain $m_{T,e(x)} = [x]^d$, which means that $\mathcal{M}_e \not\subseteq \mathcal{M}_\gamma$.

If there were a tree $T' \in \mathcal{T}_r$ such that $m_{T',\gamma} = [x_1]^l [x_{l+1}]^{r-l-1}$, then some vertex v of degree $r-l-1$ in T' must have been labeled $l+1$ by $\gamma_{T'}$. Since $\gamma_{T'}$ is increasing (cf. Remark 1.16) this means all the vertices in T'_v must be labeled by $S := \{l+1, \dots, r\}$. Since T'_v has at least $r-l$ vertices, all of the labels in S must have been used to label T'_v . But by Remark 1.16, only sets that were originally consecutive in γ' are ever used to label a subtree, which is a contradiction. \square

Remark 3.16. Thus if γ is unimodal or the rotation of a unimodal permutation, then $P_\gamma(x) = P_e(x)$ and $\varphi(p_\gamma, p_\gamma) = \varphi(p, p)$, for all weakly increasing p .

4. Conclusion

One of the chief difficulties in evaluating φ explicitly is that the rotational invariance of Proposition 2.2 is surprisingly strong compared to other, related combinatorial structures. Although the numbers $\varphi(w)$ are a kind of generalization of a Catalan number (which are known to count unrestricted non-crossing pairings), they seemingly stand out when compared to other generalizations of Catalan structures.

For arbitrary non-negative integers p_i , define the generalized Catalan numbers by

$$C_r(p_1, \dots, p_r) := \sum_{T \in \mathcal{T}_r} \prod_{i=1}^r \binom{p_i + 1}{d_i(T)},$$

where $d_i(T)$ denotes the degree of the i th (clockwise) vertex of a plane tree T (compare with Theorem 1.17 and Lemma 3.4). Note $C_r(x) = P_e(x)$. These polynomial sums are easily seen to enumerate generalized versions of most of the structures found in [15,16], including:

- The number of words $0^{p_1}1^{q_1} \dots 0^{p_r}1^{q_r}$ such that the cumulative number of 0s is always at least as large as the cumulative number of 1s.
- The number of lattice points in certain polygonal regions defined by Ambdeberhan and Stanley [2].
- The number of “ordered” decompositions of a $(p_1 + \dots + p_r + 2)$ -gon into $(p_i + 2)$ -gons. More precisely, label the polygon’s vertices in clockwise order by $1, 2, \dots, p_1 + \dots + p_r + 2$, and consider any collection of $r - 1$ chords and the resulting polygons P_1, \dots, P_r . Construct a plane tree \mathcal{T} with vertices P_i and root P_1 by placing an edge between P_i and P_j if and only if they share a chord, and order the children of P_i from left to right according to the relative clockwise order on P_i inherited from its placement in the polygon. The generalized Catalan number $C_r(p_1, \dots, p_r)$ then counts the number of such arrangements that satisfy the following conditions:
 - (1) P_i is a $(p_i + 2)$ -gon;
 - (2) P_1 contains the edge $(1, 2)$;
 - (3) The vertices of \mathcal{T} when listed in canonical order are P_1, \dots, P_r .

For the above structures, it is not in general true that $C_r(p_1, \dots, p_r) = C_r(p_2, \dots, p_r, p_1)$, which is in sharp contrast to our Proposition 2.2. Indeed, an early version of this paper used an injection from non-crossing pairings on p to generalized Catalan structures enumerated by $C_r(p_1, \dots, p_r)$ with the express purpose of breaking this rotational symmetry. The maps that we currently use in Section 3 are somewhat less directly related to $C_r(p_1, \dots, p_r)$, but they do lead to much shorter proofs.

Acknowledgments

We thank Tewodros Ambdeberhan and Richard Stanley for informative discussions regarding generalized Catalan structures. We are also extremely grateful to the two anonymous referees, both of whom read our manuscript very carefully and provided numerous helpful comments; we are additionally grateful to one of the referees for providing us with an alternative proof of our Symmetrization Theorem.

References

- [1] N. Alon, Combinatorial Nullstellensatz, *Combin. Probab. Comput.* 8 (1999) 7–29.
- [2] T. Ambdeberhan, R. Stanley, Polynomial coefficient enumeration, preprint, <http://arxiv.org/abs/0811.3652>.
- [3] Z. Bai, Y. Yin, Convergence to the semicircle law, *Ann. Probab.* 16 (2) (1988) 863–875.
- [4] E. Chou, A. Fricano, T. Kemp, J. Poh, W. Shore, G. Whieldon, T. Wong, Y. Zhang, Convex posets in non-crossing pairings on bitstrings, preprint.
- [5] P. Edelman, Chain enumeration and non-crossing partitions, *Discrete Math.* 31 (1980) 171–180.
- [6] V. Girko, The circular law, *Teor. Veroyatn. Primen.* 29 (4) (1984) 669–679.
- [7] U. Haagerup, Random matrices, free probability and the invariant subspace problem relative to a von Neumann algebra, in: *Proceedings of the International Congress of Mathematicians, vol. I, Beijing, 2002*, pp. 273–290.

- [8] T. Kemp, \mathcal{R} -diagonal dilation semigroups, *Math. Z.* 264 (2010) 111–136.
- [9] T. Kemp, R. Speicher, Strong Haagerup inequalities for free \mathcal{R} -diagonal elements, *J. Funct. Anal.* 251 (1) (2007) 141–173.
- [10] F. Larsen, Powers of R -diagonal elements, *J. Operator Theory* 47 (1) (2002) 197–212.
- [11] M. Mehta, *Random Matrices*, third ed., Elsevier, 2004.
- [12] A. Nica, R. Speicher, \mathcal{R} -diagonal pairs—a common approach to Haar unitaries and circular elements, *Fields Inst. Commun.* 12 (1997) 149–188.
- [13] A. Nica, R. Speicher, *Lectures on the Combinatorics of Free Probability*, London Math. Soc. Lecture Note Ser., vol. 335, Cambridge University Press, 2006.
- [14] F. Oravecz, On the powers of Voiculescu's circular element, *Studia Math.* 145 (2001) 85–95.
- [15] R. Stanley, *Enumerative Combinatorics*, vol. 2, Cambridge University Press, Cambridge, MA, 1997.
- [16] R. Stanley, Catalan addendum, online <http://math.mit.edu/~rstan/ec/catadd.pdf>, accessed September 27, 2008.
- [17] T. Tao, V. Vu, Random matrices: the circular law, *Commun. Contemp. Math.* 10 (2008) 261–307.