

Tree-like Properties of Cycle Factorizations

Ian Goulden

*Department of Combinatorics and Optimization, University of Waterloo,
Waterloo, Ontario N2L 3G1, Canada*
E-mail: ipgoulden@math.uwaterloo.ca

and

Alexander Yong

Department of Mathematics, University of Michigan, Ann Arbor, Michigan 48109
E-mail: ayong@umich.edu

Received June 8, 2001

We provide a bijection between the set of *factorizations*, that is, ordered $(n-1)$ -tuples of transpositions in \mathcal{S}_n whose product is $(12\dots n)$, and labelled trees on n vertices. We prove a refinement of a theorem of J. Dénes (1959, *Publ. Math. Inst. Hungar. Acad. Sci.* 4, 63–71) that establishes new tree-like properties of factorizations. In particular, we show that a certain class of transpositions of a factorization corresponds naturally under our bijection to leaf edges (incident with a vertex of degree 1) of a tree. Moreover, we give a generalization of this fact. © 2002 Elsevier

Science (USA)

1. INTRODUCTION

Let \mathcal{T}_n be the set of labelled trees on vertices $\{1, 2, \dots, n\} = [n]$, and \mathcal{F}_n be the set of $(n-1)$ -tuples of transpositions $(\sigma_1, \dots, \sigma_{n-1})$ in the symmetric group \mathcal{S}_n acting on $[n]$, whose ordered product $\sigma_1 \cdots \sigma_{n-1}$ is equal to the cycle $\mathcal{C}_n = (12\dots n)$ (in this paper, our convention is that, in an ordered product, the permutations are performed right to left). The elements of \mathcal{F}_n are called *factorizations*, and the transpositions $\sigma_1, \dots, \sigma_{n-1}$ in a factorization are called *factors*. Cayley [3] proved that $|\mathcal{F}_n| = n^{n-2}$, and Dénes [4] proved that $|\mathcal{F}_n| = |\mathcal{T}_n|$, by giving a bijection between sets of cardinality $(n-1)! |\mathcal{F}_n|$ and $(n-1)! |\mathcal{T}_n|$. Dénes posed the problem of finding a bijection between \mathcal{F}_n and \mathcal{T}_n , and subsequently two such bijections have been given, by Moszkowski [12] and Goulden and Pepper [10].

Although both of these bijections are reasonably simple, neither of them restricts nicely to natural combinatorial subsets (e.g., so that the image of a combinatorially natural subset of \mathcal{T}_n corresponds to a natural subset of \mathcal{F}_n). However, by examining the elements of \mathcal{T}_n and \mathcal{F}_n for small n , we find that there are natural combinatorial subsets of \mathcal{T}_n and \mathcal{F}_n of equal cardinality, as follows. Let $\mathcal{T}_n(k)$ be the set of trees in \mathcal{T}_n with k leaves (vertices of degree one). A transposition (st) on $[n]$ is called a *consecutive pair* if $t \equiv s+1$ modulo n , where throughout, we write n to mean 0. Let $\mathcal{F}_n(k)$ be the set of factorizations in \mathcal{F}_n with k factors that are consecutive pairs. Table I gives the cardinalities $|\mathcal{T}_n(k)|$ for $n \leq 6$, and a systematic examination of the factorizations in \mathcal{F}_n shows that $|\mathcal{F}_n(k)| = |\mathcal{T}_n(k)|$ for $3 \leq n \leq 6$. This suggests that there exists a bijection between \mathcal{F}_n and \mathcal{T}_n for arbitrary $n \geq 3$ that maps consecutive pairs to leaves, but neither of the previous bijections exhibits this property.

In this paper, we describe a bijection between \mathcal{F}_n and \mathcal{T}_n in which consecutive pairs of factorizations correspond to leaves of trees. We refer to our bijection as a *structural* bijection because of this correspondence between these combinatorial structures (consecutive pairs and leaves). But more is true; the bijection extends to generalizations of consecutive pairs and of leaves respectively, as described below.

For a tree $T \in \mathcal{T}_n$, consider removing any single edge from the tree, to get two trees T_1 and T_2 (the components of the graph that results when the edge is deleted from T). Let t_i , $i = 1, 2$, be the number of vertices in T_i (so, e.g., $t_1 + t_2 = n$), and define the *edge-deletion index* of the edge to be $\min\{t_1, t_2\}$. Define the *edge-deletion distribution* of the tree T to be $\varepsilon(T) = (a_1, a_2, \dots)$ where a_j is the number of edges in T with edge-deletion index j (so e.g., $a_1 + a_2 + \dots = n-1$). Let $\mathcal{T}_n(a_1, a_2, \dots)$ be the set of trees in \mathcal{T}_n with edge-deletion distribution (a_1, a_2, \dots) .

For a transposition (st) , $s < t$, define the *difference index* to be $\min\{t-s, n-t+s\}$. For a factorization $F \in \mathcal{F}_n$, define the *difference distribution* of F to be $\delta(F) = (d_1, d_2, \dots)$ where d_j is the number of factors in F

TABLE I

The Number of Trees on n Vertices with k Leaves for $n \leq 6$

n/k	2	3	4	5	Total
2	1				1
3	3				3
4	12	4			16
5	60	60	5		125
6	360	720	210	6	1296

with difference index j (so e.g., $d_1 + d_2 + \dots = n - 1$). Let $\overline{\mathcal{F}}_n(d_1, d_2, \dots)$ be the set of trees in \mathcal{F}_n with difference distribution (d_1, d_2, \dots) .

Our structural bijection, described in Section 3, actually gives a bijection between $\overline{\mathcal{F}}_n(c_1, c_2, \dots)$ and $\mathcal{T}_n(c_1, c_2, \dots)$ for all (c_1, c_2, \dots) , $c_i \geq 0$, $c_1 + c_2 + \dots = n - 1$, $n \geq 1$. The main result of our paper is as follows:

THEOREM 1.1. *For each $n \geq 1$, there is a bijection*

$$\phi: \overline{\mathcal{F}}_n \rightarrow \mathcal{T}_n : F \mapsto T$$

such that $\delta(F) = \varepsilon(T)$.

Note that for factorizations, a factor with difference index 1 is a consecutive pair, and for trees, an edge with edge-deletion index 1 is incident with a leaf, so this is a generalization of the consecutive pair-leaf correspondence, as promised. In the case $n = 2$, the single edge in the unique tree has difference index equal to 1, but is incident with two leaves. Of course, the single factor (12) in the unique factorization is a consecutive pair, so the bijection claimed in Theorem 1.1 holds for $n = 2$, although the consecutive pair-leaf correspondence breaks down.

The bijection is based on a geometrical interpretation of a factorization, called a *chord diagram*, whose properties are developed in Section 2. The bijection, described in Section 3, has a smooth composition with the well-known Prüfer code bijection between trees and the set $[n]^{n-2}$. Consequently one can obtain a bijection under this composition that canonically proves that $|\overline{\mathcal{F}}_n| = n^{n-2}$.

Further motivation for this paper, which gives a third bijection for the Dénes result, beyond the combinatorial benefits of exhibiting tree properties of edges as differences of factors, is provided in Section 4, where we describe recent work on more general factorization questions in the symmetric group, related to certain problems arising from algebraic geometry.

Finally, there are immediate enumerative consequences of our main result. For example, there is a nice formula for the entries in Table I, which can be obtained in various ways by counting trees with a given number of leaves. This formula is a simple multiple of a Stirling number of the second kind, and in closed form it gives

$$|\mathcal{T}_n(k)| = \binom{n}{k} \sum_{i=0}^{n-k} \binom{n-k}{i} (-1)^{n-k-i} i^{n-2},$$

for $2 \leq k \leq n$ (see, e.g., Stanton and White [14, p. 67]). Using our bijection, we therefore have established that this formula also holds for $|\overline{\mathcal{F}}_n(k)|$.

2. THE CIRCLE CHORD DIAGRAM

We begin with a detailed analysis of several aspects of factorizations in \mathcal{F}_n . First, if $\rho \in \mathcal{S}_n$ and $\tau = (st) \in \mathcal{S}_n$ is a transposition, then there are two cases that arise in determining the product $\alpha = \tau\rho$. If s, t appear on the same cycle in the disjoint cycle representation of ρ , then that cycle is cut into two different cycles (one containing s , the other t) in the disjoint cycle representation of α . Otherwise, if s, t appear on two different cycles of ρ , then these cycles are joined into one cycle (containing both s and t) of α . We call τ a *join* or a *cut* of ρ , respectively, in these cases (see, e.g., Berge [2, p. 142]).

For a factorization $F = (\sigma_1, \dots, \sigma_{n-1})$ in \mathcal{F}_n , let $f_i = \sigma_i\sigma_{i+1}\cdots\sigma_{n-1}$, $i = 1, \dots, n-1$, be the *partial products* of F , and define σ_i to be a *join* or *cut* of F when σ_i is a join or cut of f_{i+1} , respectively (when determining the product $f_i = \sigma_i f_{i+1}$), for $i = 1, \dots, n-1$ (where $f_n = e$, the identity of \mathcal{S}_n). Now f_n has n cycles (all fixed points), and $f_1 = \mathcal{C}_n$ has 1 cycle. Moreover, each join decreases the number of cycles by 1, and each cut increases the number of cycles by 1. We conclude that each of the $n-1$ factors σ_i in $f_1 = \sigma_1 \cdots \sigma_{n-1} f_n$ must be a join, since together they decrease the n cycles of f_n by $n-1$, to the single cycle of f_1 .

We say that a sequence $\alpha_1, \dots, \alpha_m$ of elements in $[n]$ is \mathcal{C}_n -ordered if the order of the elements is consistent with their circular order on the cycle $\mathcal{C}_n = (12\dots n)$. Equivalently, this means that there is a unique i with $1 \leq \alpha_i < \alpha_{i+1} < \cdots < \alpha_m < \alpha_1 < \cdots < \alpha_{i-1} \leq n$.

PROPOSITION 2.1. *For $F \in \mathcal{F}_n$, and a partial product f_i of F , any subsequence of elements on a cycle of f_i is \mathcal{C}_n -ordered.*

Proof. For $F = (\sigma_1, \dots, \sigma_{n-1})$, we have $f_i = \sigma_i \cdots \sigma_{n-1}$, where all factors $\sigma_1, \dots, \sigma_{n-1}$ are joins, from the above join-cut analysis. Thus the effect of the sequence of subsequent joins $\sigma_1, \dots, \sigma_{i-1}$ in $f_1 = \sigma_1 \cdots \sigma_{i-1} f_i$, on the elements of a cycle of f_i is to keep them together on cycles that are formed by the joins. Moreover, it is easy to check that their circular order is maintained around such cycles. But $f_1 = \mathcal{C}_n$, so we conclude that the elements on each cycle of f_i must be \mathcal{C}_n -ordered, and therefore so must all subsequences of elements on each such cycle. ■

We now consider a *circle chord diagram*. For any fixed n , this is a circle drawn in the plane with n points on it, labelled $1, 2, \dots, n$ clockwise. In addition, there are $n-1$ chords on these n points, numbered $2, \dots, n$ distinctly. For example, Fig. 1 gives a circle chord diagram with $n = 9$; the numbers on the edges are circled to distinguish them from the names of points on the circle.




FIG. 1. A circle chord diagram with $n = 9$.

There is a natural injection from factorizations to circle chord diagrams: for $F = (\sigma_1, \dots, \sigma_{n-1}) \in \mathcal{F}_n$, the factor $\sigma_i = (s_i t_i)$ corresponds to a chord numbered $i + 1$, joining points s_i and t_i , for $i = 1, \dots, n - 1$. Let $C(F)$ be the circle chord diagram associated with F in this way. For example, the chord diagram illustrated in Figure 1 is $C(F_0)$, where

$$F_0 = ((2\ 3), (4\ 5), (3\ 6), (3\ 5), (1\ 6), (6\ 8), (8\ 9), (6\ 7)). \quad (1)$$

The circle chord diagram associated with a factorization F satisfies a number of conditions, and we establish some of these in the following result. Part (i) of this result is also part of the development of Dénes [4] (see also Berge [2, p. 143]).

THEOREM 2.2. *In the circle chord diagram $C(F)$ of a factorization $F \in \overline{\mathcal{F}}_n$,*

- (i) *the chords form a tree on $[n]$, $(*_1)$*
- (ii) *the chords meet only at endpoints, $(*_2)$*
- (iii) *the edge numbers on the chords encountered when moving around a vertex clockwise across the interior of the circle, form a decreasing sequence of elements in $\{2, \dots, n\}$. $(*_3)$*

Proof. For (i), let \mathcal{G}_j be the graph on vertex-set $[n]$, whose edges are the chords corresponding to factors $\sigma_j, \dots, \sigma_{n-1}$, for $j = 1, \dots, n$ (\mathcal{G}_n has no edges). Then \mathcal{G}_n has n components (each a single vertex), and the condition,

established above, that σ_j is a join for each j implies that the chord corresponding to σ_j is incident with vertices in different components of \mathcal{G}_{j+1} , for each $j = 1, \dots, n-1$. Thus \mathcal{G}_j has one fewer components than \mathcal{G}_{j+1} for each $j = 1, \dots, n-1$, and we conclude that \mathcal{G}_1 has one component, so it is a connected graph. But the edges of G_1 are the chords of $C(F)$, so the $n-1$ chords of $C(F)$ are a connected graph on n vertices, which must therefore be a tree.

For (ii), suppose otherwise, that the chords corresponding to $\sigma_i = (st)$ and $\sigma_j = (uv)$, where $s < t, u < v$ and $i < j$, cross each other. Now the geometric crossing condition is equivalent to the condition that the sequence $stuv$ is not \mathcal{C}_n -ordered. But in f_i , the cycle containing s will include $stuv$ as a subsequence, and we have a contradiction of Proposition 2.1, which establishes that $stuv$ must be \mathcal{C}_n -ordered. We conclude that chords do not cross and can therefore meet only at endpoints.

For (iii), for each fixed $i = 1, \dots, n$, suppose the factors moving i are

$$\sigma_{l_1} = (i s_1), \dots, \sigma_{l_k} = (i s_k)$$

where $1 \leq l_1 < \dots < l_k \leq n-1, k \geq 1$.

Then f_{l_1} will include $is_k \dots s_1$ as a subsequence on the cycle containing i , and we conclude from Proposition 2.1 that $is_k \dots s_1$ is \mathcal{C}_n -ordered. But the edge corresponding to σ_{l_j} has number $l_j + 1$, and (iii) follows. ■

For example, it is straightforward to verify that the circle chord diagram $C(F_0)$ illustrated in Figure 1 does indeed satisfy conditions $(*_1)$, $(*_2)$ and $(*_3)$.

Now for circle chord diagrams satisfying condition $(*_2)$, the $n-1$ chords and the circle partition the circle and its interior into n regions. The boundary of a region consists of a collection of chords and arcs of the circle. An arc is a segment of the circle from point i to point $i+1$ modulo n .

PROPOSITION 2.3. *For circle chord diagrams satisfying conditions $(*_1)$ and $(*_2)$, each region contains precisely one arc in its boundary.*

Proof. If the boundary of any region consists entirely of chords, then these chords form a cycle in the graph of the chords (called \mathcal{G}_1 in the proof of Theorem 2.2(i)). But this graph is a tree, from Theorem 2.2(i), and therefore has no cycles. We conclude that the boundary of each of the n regions contains at least one of the n arcs. But this means that each region has exactly one arc, giving the result. ■

For example, each region of the circle chord diagram $C(F_0)$ illustrated in Fig. 1 contains precisely one arc in its boundary.

Now consider the following condition for the above regions: the numbers on the chords of the boundary increase clockwise, starting immediately after the unique arc. $(*_3)'$

Note, for example, that each of the 9 regions in Fig. 1 satisfies $(*_3)'$.

PROPOSITION 2.4. *For circle chord diagrams, conditions $(*_1)$, $(*_2)$ and $(*_3)$ are equivalent to $(*_1)$, $(*_2)$ and $(*_3)'$.*

Proof. Immediate. ■

We end this section by showing that conditions $(*_1)$, $(*_2)$ and $(*_3)'$ characterize circle chord diagrams associated with factorizations.

LEMMA 2.5. *A circle chord diagram on n points satisfying conditions $(*_1)$, $(*_2)$ and $(*_3)'$ is equal to $C(F)$ for some $F \in \mathcal{F}_n$.*

Proof. Consider a circle chord diagram satisfying conditions $(*_1)$, $(*_2)$ and $(*_3)'$. Suppose that the chord numbered i joins points a_i and b_i , and let $\sigma_{i-1} = (a_i b_i)$, for $i = 2, \dots, n$. Now consider the product of transpositions

$$\sigma = \sigma_1 \cdots \sigma_{n-1}.$$

Condition $(*_3)'$ implies that $\sigma(j) \equiv j+1$ modulo n for each $j = 1, \dots, n$, by considering the action of the transpositions on the boundary of the region containing the arc $(j, j+1)$. Thus $\sigma = \mathcal{C}_n$, and $F' = (\sigma_1, \dots, \sigma_{n-1})$ is a factorization in \mathcal{F}_n . The result follows, since we have established that the circle chord diagram is equal to $C(F')$. ■

In summary, the results in this section have established that there is a bijection between \mathcal{F}_n and circle chord diagrams satisfying $(*_1)$, $(*_2)$, and $(*_3)'$.

3. THE STRUCTURAL BIJECTION

We are now able to describe the structural bijection that proves our main theorem. Consider the circle chord diagram $C(F)$ for some factorization $F \in \mathcal{F}_n$. Form the graph $\phi(F)$ by a “planar dual” construction, as follows. For each region of $C(F)$ we have a vertex of $\phi(F)$ (say, drawn in the middle of the arc contained in the boundary of the region). Then place an edge between two vertices if the boundaries of their corresponding regions share a chord. We (temporarily) assign label i to this edge of $\phi(F)$, where i is the number of the shared chord in $C(F)$. Thus, at this stage, $\phi(F)$ has n vertices and $n-1$ edges (one edge for each edge of $C(F)$), and is connected because $C(F)$ is connected, so we conclude that $\phi(F)$ is a tree.

Now pause to observe two properties of $\phi(F)$ at this stage. (These observations are not necessary to complete our construction, but make the reversal of our bijection straightforward.) First, condition $(*_3)'$ on $C(F)$ implies that, at each vertex of $\phi(F)$, the clockwise sequence of labels on the incident edges is \mathcal{C}_n -ordered. Second, the clockwise order of the vertices of $\phi(F)$ around the circle can be described, as follows. Vertex 1 is placed on the arc $(n, 1)$. Make the unique traversal of the outside of the tree in a clockwise direction, starting from vertex 1 and proceeding along the left side of the edge with the smallest temporary label among those edges incident with vertex 1. This traversal ends when both sides of every edge have been traversed, at vertex 1. At each visit to a vertex, compare the temporary label on the edge whose side has been traversed immediately before the vertex, to the temporary label on the edge whose side has been traversed immediately after the vertex; check for the condition that the “before” label is weakly greater than the “after” label. Note that for a leaf vertex, this condition is always satisfied at the unique visit to the vertex, since the labels being compared are identical (arising from two sides of the same leaf edge). For vertex 1, this condition is never satisfied, because of the clockwise \mathcal{C}_n -ordered property above. For every other vertex, this condition is satisfied exactly once, again because of the clockwise \mathcal{C}_n -ordered property. The order of the vertices appearing on arcs $(1, 2), \dots, (n-1, n)$ is now straightforward; it is the order in which the vertices satisfy the weakly decreasing condition along the traversal.

Finally, complete the construction by labelling the vertices distinctly with the elements of $[n]$, as follows. The vertex corresponding to the region with arc $(n, 1)$ in its boundary has label 1. For each edge, find the unique path to vertex 1 from that edge, and “slide” the temporary label on the edge to the incident vertex away from vertex 1, thus labelling the other $n-1$ vertices $2, \dots, n$. The resulting tree is $\phi(F)$, and it is clear from our description above that $\phi(F) \in \mathcal{T}_n$. For example, Fig. 2 illustrates $\phi(F_0)$ where F_0 is given in (1), and $C(F_0)$ is given in Fig. 1.

We claim that $\phi: \mathcal{F}_n \rightarrow \mathcal{T}_n$ is a bijection. In section 2 we proved that there is a bijection between \mathcal{F}_n and circle chord diagrams satisfying $(*_1)$, $(*_2)$ and $(*_3)'$. Also, our “planar dual” construction above is a bijection between circle chord diagrams satisfying $(*_1)$, $(*_2)$ and $(*_3)'$ and \mathcal{T}_n , since the clockwise \mathcal{C}_n -ordered requirement at each vertex forces a unique planar embedding of a tree, and then the vertices are placed at the midpoints of the arcs of the circle uniquely as described above. Together, these bijections prove the claim. In the resulting bijection ϕ , note that a factor in F with difference index k corresponds precisely to an edge of $\phi(F)$ with edge-deletion index k , so $\delta(F) = \varepsilon(\phi(F))$ and we have proved Theorem 1.1.

To reverse the bijection, consider an arbitrary tree T . Now “slide” the label on each vertex $2, \dots, n$ to the incident edge along the unique path to




FIG. 2. Construction of the tree $\phi(F_0)$.

vertex 1. By our two observations above, it is straightforward to see that we can uniquely embed the tree in the plane so that the clockwise order of the edge labels incident with every vertex is increasing, and then traverse the tree to place the vertices at the midpoints of the arcs on a circle with vertices $1, \dots, n$. We complete the determination of $\phi^{-1}(T)$ straightforwardly by inverting our planar dualization above to find the chords.




FIG. 3. Example of inverting the structural bijection.

For example, for the tree T_1 , with $n = 8$, and edges 13, 14, 17, 23, 28, 35, 56, we find that

$$\phi^{-1}(T_1) = ((1\ 2), (1\ 3), (1\ 6), (4\ 5), (4\ 6), (6\ 7), (7\ 8)),$$

as illustrated in Fig. 3.

As a remark, we mention that our circle chord diagram construction leads directly to another bijection between \mathcal{F}_n and \mathcal{T}_n . Given a circle chord diagram associated with a factorization, simply “push” the edge labels in the unique direction away from the vertex labelled 1. This gives an element of \mathcal{T}_n and is clearly reversible. This bijection is actually the same as that of Moszkowski [12], and the description given above has appeared, independently, in Poulalhon [13]. Also, in a different but related setting, a planar dual construction that is similar to ours has been given in Deutsch and Noy [5].

4. FACTORIZATIONS AND HURWITZ NUMBERS

The factorizations that we have considered in this paper are special cases of a more general factorization problem in \mathcal{S}_n . Consider k -tuples of transpositions $(\sigma_1, \dots, \sigma_k)$ whose ordered product $\sigma_1 \cdots \sigma_k$ is equal to an arbitrary permutation π , and such that the group generated by $\sigma_1, \dots, \sigma_k$ acts transitively on $[n]$. (When $\pi = \mathcal{C}_n$, as is the case in this paper, this transitivity is forced. Note that, in general, transitivity means in combinatorial terms that the multigraph on vertices $[n]$, and edge $i\ j$ for each factor $(i\ j)$, is connected.) For each k , the number of such factorizations is clearly constant on the conjugacy class of π in \mathcal{S}_n . Moreover, if the conjugacy class has disjoint cycle distribution specified by the partition $\alpha = (\alpha_1, \dots, \alpha_m)$ of n , with m parts, then the minimum choice of k for which such factorizations exist is $k = n + m - 2$, and corresponding factorizations are called *minimal transitive factorizations*. (Note that, from Goulden and Jackson [8], there are exactly $m - 1$ cuts in these factorizations, in addition to $n - 1$ joins.) The number of these factorizations is given by

$$(n + m - 2)! n^{m-3} \prod_{j=1}^m \frac{\alpha_j^{\alpha_j}}{(\alpha_j - 1)!}, \quad (2)$$

from [8]. Such factorizations arise in the study of ramified covers of the sphere by the sphere, with branching above infinity specified by α , simple branching above other specified points, and no other branching (see, for example, [1, 9, and 11]).

In the case where the number of factors is $k = n + m - 2 + 2g$, for an arbitrary nonnegative integer g , such factorizations arise in ramified covers of the sphere by a surface of genus g . The number of such covers, equal to the number of corresponding factorizations as specified above, are called Hurwitz numbers, and are studied extensively in algebraic geometry (see, for example, [6, and 7]). Here, the expression for k is a consequence of the Riemann–Hurwitz formula.

Clearly the expression in (2) specializes correctly to give n^{n-2} in the case that $\alpha = (n)$, for which $k = n - 1$, and these are the factorizations studied in this paper. We would like to achieve a combinatorial understanding of this expression for arbitrary α . The bijection in this paper allows us to specify transposition factors according to their difference index, and in particular to identify those in which this difference equals 1 (namely, the consecutive pairs). These pairs are mapped to leaves in the corresponding tree, and the Prüfer code bijection for trees (see, e.g., Stanton and White [14, p. 66]) is based on successively removing leaves of the tree, each iteration yielding an element of $[n]$. Thus we can compose our bijection with the Prüfer bijection “smoothly”, to identify combinatorially each of the factors n in the enumeration of the factorizations of this paper. Now expression (2) contains many similar factors, and our hope is that the combinatorial decompositions of this paper can be extended to explain these factors in the general case. As a specific instance of this possible extension, consider the case $\alpha = (n-1, 1)$, where expression (2) becomes $(n-1)^n$, and the factorizations would have a single cut. The simplicity of this expression suggests that a nice combinatorial explanation of the type referred to above should be possible, but we have not yet been able to find one.

ACKNOWLEDGMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada, through a grant to I.G., and a PGSA to A.Y. This work was partially completed while A.Y. was visiting the Fields Institute in Toronto. We thank Gilles Schaeffer for helpful discussions. References [2, 5] were brought to our attention by an anonymous referee, who made other helpful suggestions.

REFERENCES

1. V. I. Arnold, Topological classification of trigonometric polynomials and combinatorics of graphs with an equal number of vertices and edges, *Funct. Anal. Appl.* **30** (1996), 1–14.
2. C. Berge, “Principles of Combinatorics,” Academic Press, New York, 1971.
3. A. Cayley, A theorem on trees, *Quart. J. Oxford* **23** (1889), 376–378.

4. J. Dénes, The representation of a permutation as the product of a minimal number of transpositions and its connection with the theory of graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* **4** (1959), 63–71.
5. E. Deutsch and M. Noy, New statistics on non-crossing trees, in “Formal Power Series and Algebraic Combinatorics, Proceedings of FPSAC’00, Moscow, Russia, June 2000,” pp. 667–676, Springer-Verlag, Berlin, 2000.
6. T. Ekedahl, S. Lando, M. Shapiro, and A. Vainshtein, On Hurwitz numbers and Hodge integrals, *C. R. Acad. Sci. Paris Sér. I* **328** (1999), 1175–1180.
7. B. Fantechi and R. Pandharipande, Stable maps and branch divisors, math. AG/9905104.
8. I. P. Goulden and D. M. Jackson, Transitive factorisations into transpositions and holomorphic mappings on the sphere, *Proc. Amer. Math. Soc.* **125** (1997), 51–60.
9. I. P. Goulden, D. M. Jackson, and R. Vakil, The Gromov–Witten potential of a point, Hurwitz numbers and Hodge integrals, *Proc. London Math. Soc. (3)* **83** (2001), 563–581.
10. I. P. Goulden and S. Pepper, Labelled trees and factorizations of a cycle into transpositions, *Discrete Math.* **113** (1993), 263–268.
11. A. Hurwitz, Ueber Riemann’sche Flächen mit gegebenen Verzweigungspunkten, *Math. Ann.* **39** (1891), 1–61.
12. P. Moszkowski, A solution to a problem of Dénes: A bijection between trees and factorizations of cyclic permutations, *Europ. J. Combin.* **10** (1989), 13–16.
13. D. Poulalhon, “Graphes et décompositions de permutations,” Mémoire de DEA, LIX, École Polytechnique, July 1997.
14. D. Stanton and D. White, “Constructive Combinatorics,” Springer-Verlag, New York, 1986.