

Minimal Transitive Factorizations of Permutations into Cycles

John Irving

Abstract. We introduce a new approach to an enumerative problem closely linked with the geometry of branched coverings, that is, we study the number $H_\alpha(i_2, i_3, \dots)$ of ways a given permutation (with cycles described by the partition α) can be decomposed into a product of exactly i_2 2-cycles, i_3 3-cycles, etc., with certain minimality and transitivity conditions imposed on the factors. The method is to encode such factorizations as planar maps with certain *descent structure* and apply a new combinatorial decomposition to make their enumeration more manageable. We apply our technique to determine $H_\alpha(i_2, i_3, \dots)$ when α has one or two parts, extending earlier work of Goulden and Jackson. We also show how these methods are readily modified to count *inequivalent* factorizations, where equivalence is defined by permitting commutations of adjacent disjoint factors. Our technique permits us to generalize recent work of Goulden, Jackson, and Latour, while allowing for a considerable simplification of their analysis.

1 Introduction

We begin with a brief review of standard notation. We write $\alpha \vdash n$ (respectively, $\alpha \models n$) to indicate that α is a partition (composition) of n , and denote by $\ell(\alpha)$ the number of parts of α . The partition with m_i parts equal to i is denoted by $[1^{m_1} 2^{m_2} \dots]$.

The *cycle type* of a permutation π in the symmetric group \mathfrak{S}_n is the partition of n determined by the lengths of the disjoint cycles comprising π . The conjugacy class of \mathfrak{S}_n consisting of all permutations of cycle type $\lambda \vdash n$ is denoted by \mathcal{C}_λ . This notation is extended to allow λ to be a composition, in which case the ordering of the parts of λ is ignored. Members of $\mathcal{C}_{[1^{n-k} k]}$ are called *cycles of length k* , or *k -cycles*, while elements of $\mathcal{C}_{[n]}$ are *full cycles* in \mathfrak{S}_n .

For vectors $\mathbf{j} = (j_1, \dots, j_m)$ and $\mathbf{x} = (x_1, \dots, x_m)$ we use the abbreviations $\mathbf{x}^{\mathbf{j}} = x_1^{j_1} \dots x_m^{j_m}$ and $\mathbf{j}! = j_1! \dots j_m!$. Finally, if $f \in \mathbb{Q}[[\mathbf{x}]]$ is a formal power series, then we write $[\mathbf{x}^{\mathbf{j}}] f(\mathbf{x})$ for the coefficient of the monomial $\mathbf{x}^{\mathbf{j}}$ in $f(\mathbf{x})$.

1.1 Factorizations of Permutations into Cycles

A *transitive factorization* of a permutation $\pi \in \mathfrak{S}_n$ is a tuple $F = (\sigma_r, \dots, \sigma_1)$ of permutations $\sigma_i \in \mathfrak{S}_n$ such that $\pi = \sigma_r \dots \sigma_1$, and the group $\langle \sigma_1, \dots, \sigma_r \rangle$ generated by the factors acts transitively on \mathfrak{S}_n . If $\sigma_i \in \mathcal{C}_{\beta_i}$ for $1 \leq i \leq r$ and $\pi \in \mathcal{C}_\alpha$, then one

Received by the editors October 20, 2006.
AMS subject classification: Primary: 05A15; secondary: 05E10.
©Canadian Mathematical Society 2009.

can show [8] that

$$(1.1) \quad nr - \sum_{i=1}^r \ell(\beta_i) \geq n + \ell(\alpha) - 2.$$

In the case of equality above, F is said to be *minimal transitive*. For example, since

$$(1.2) \quad (1\ 2\ 3\ 4\ 5)(6\ 7\ 8)(9) = (3\ 4\ 7\ 9) \cdot (6\ 9\ 7\ 8) \cdot (1\ 2) \cdot (2\ 4\ 5) \cdot (2\ 6),$$

the tuple $((3479), (6978), (12), (245), (26))$ is a factorization of $(12345)(678)(9)$. (Fixed points have been suppressed in the factors.) This factorization is easily verified to be minimal transitive.

Minimal transitive factorizations have been very well studied, with much of this attention stemming from the fact that they serve geometers as combinatorial models for branched coverings of the sphere by the sphere. In this context, transitivity guarantees connectedness of the associated covering, while minimality implies the covering surface is the sphere. For further information on these connections see [1] and the references therein.

The focus of this paper is the class of *cycle factorizations*, by which we mean minimal transitive factorizations, such as (1.2) above, whose factors are all cycles of length at least two.¹ In particular, given a composition α and a sequence $\mathbf{i} = (i_2, i_3, \dots)$ of nonnegative integers (called the *cycle index*), we wish to determine the number $H_\alpha(\mathbf{i})$ of cycle factorizations of any fixed permutation $\pi \in \mathcal{C}_\alpha$ into exactly i_2 2-cycles, i_3 3-cycles, etc. Our results will be formulated in terms of the generating series

$$(1.3) \quad \Psi_m(\mathbf{x}, \mathbf{q}, u) := \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} H_\alpha(\mathbf{i}) \mathbf{q}^{\mathbf{i}} \frac{x_1^{\alpha_1}}{\alpha_1} \cdots \frac{x_m^{\alpha_m} u^{r(\mathbf{i})}}{\alpha_m r(\mathbf{i})!}, \quad m \geq 1.$$

Here, and throughout, $r(\mathbf{i}) := i_2 + i_3 + \dots$ denotes the total number of factors in any factorization counted by $H_\alpha(\mathbf{i})$, and $\mathbf{q} = (q_2, q_3, \dots)$ is a vector of indeterminates.

The structure of generic cycle factorizations is not well understood and, aside from explicit evaluations of $H_{[n]}(\mathbf{i})$ (see [6, 17] and Theorem 3.2 of this paper), little work has been done on their enumeration. However, a significant effort has been directed toward a natural specialization of this problem, which is to count what we call *k-cycle factorizations*. These are cycle factorizations whose factors are all k -cycles for some fixed k .

The case $k = 2$ (transposition factors) is particularly important geometrically. Counting 2-cycle factorizations of permutations is known as the *Hurwitz problem*, and dates back to Hurwitz’s original investigations into the classification of almost simple ramified coverings of the sphere by the sphere [13]. The following formula, suggested but not completely proved by Hurwitz himself, gives the number of 2-cycle factorizations of any permutation of cycle type $(\alpha_1, \dots, \alpha_m)$:

$$(1.4) \quad n^{m-3} (n + m - 2)! \prod_{i=1}^m \frac{\alpha_i^{\alpha_i+1}}{\alpha_i!}.$$

¹This condition avoids the triviality of having factors equal to the identity.

Although it has been extensively studied from various points of view, ranging from analytic to combinatorial (see [1, 4, 8], for example), no purely bijective proof of this striking enumerative formula is known. Such a proof would be of tremendous interest as it could provide further insight into the underlying geometry. This is particularly true in light of a recent celebrated result of Ekedahl, Lando, Shapiro, and Vainshtein [3] that identifies the enumeration of 2-cycle factorizations, *i.e.*, almost simple coverings, with the evaluation of certain Hodge integrals, objects of great interest in the intersection theory of the moduli space of curves. See [11] for further details on these fascinating connections. Indeed, the ultimate goal of our approach to factorization problems is a full combinatorialization of Hodge integrals.

For arbitrary $k > 2$, counting k -cycle factorizations has not been as thoroughly examined and appears quite difficult. Substantial progress was made in [9], where generating series formulations for the number of such factorizations of permutations with up to three cycles are given. In particular, letting $\Psi_{m,k}$ denote the series obtained by specializing Ψ_m at $u = q_k = 1$ and $q_i = 0$ for $i \neq k$, it is shown there that $x \frac{d}{dx} \Psi_{1,k}(x) = s(x)$ and

$$(1.5) \quad \Psi_{2,k}(x_1, x_2) = \log\left(\frac{s(x_1) - s(x_2)}{x_1 - x_2}\right) - \frac{s(x_1)^k - s(x_2)^k}{s(x_1) - s(x_2)},$$

where $s \in \mathbb{Q}[[x]]$ (which depends on k) is the unique series satisfying

$$(1.6) \quad s = xe^{s^{k-1}}.$$

A more complicated expression for $\Psi_{m,3}$ is also given in terms of s , but the calculations necessary to evaluate $\Psi_{m,k}$ for $k \geq 4$ become intractable.

1.2 Equivalence up to Commutation of Disjoint Factors

There is a natural equivalence relation on cycle factorizations, induced by permitting commutations of disjoint factors. That is, we say two cycle factorizations are *equivalent* if one can be obtained from the other by repeatedly exchanging adjacent factors that are disjoint in the sense that no symbol is moved by both. For example, the following factorizations are equivalent:

$$(3\ 4\ 7\ 9) \cdot (6\ 9\ 7\ 8) \cdot (1\ 2) \cdot (2\ 4\ 5) \cdot (2\ 6) \sim (1\ 2) \cdot (3\ 4\ 7\ 9) \cdot (2\ 4\ 5) \cdot (6\ 9\ 7\ 8) \cdot (2\ 6).$$

Let $\tilde{H}_\alpha(\mathbf{i})$ denote the number of inequivalent cycle factorizations of $\pi \in \mathcal{C}_\alpha$ with cycle index \mathbf{i} . We shall study these numbers through the series

$$(1.7) \quad \tilde{\Psi}_m(\mathbf{x}, \mathbf{q}, u) := \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\alpha = n \\ \ell(\alpha) = m}} \tilde{H}_\alpha(\mathbf{i}) \mathbf{q}^{\mathbf{i}} \frac{x_1^{\alpha_1}}{\alpha_1} \cdots \frac{x_m^{\alpha_m}}{\alpha_m} u^{r(\mathbf{i})}, \quad m \geq 1.$$

As before, let $\tilde{\Psi}_{m,k}(\mathbf{x})$ be the restricted series counting inequivalent k -cycle factorizations obtained from $\tilde{\Psi}_m$ by setting $u = q_k = 1$ and $q_i = 0$ for $i \neq k$ in $\tilde{\Psi}_m$.

The problem of counting factorizations up to commutation can apparently be traced back to Stanley, who originally posed it in the context of 2-cycle factorizations. The first result along these lines came from Eidswick [2] and Longyear [15], who proved, independently, that the number of inequivalent 2-cycle factorizations of the full cycle $(1\ 2\ \cdots\ n)$ is the generalized Catalan number

$$(1.8) \quad \frac{1}{2n-1} \binom{3n-3}{n-1}.$$

Longyear’s approach involved commutation of factorizations into a canonical form. This led to the following cubic functional equation for the generating series $h(x) = \frac{d}{dx}\Psi_{1,2}(x)$, from which (1.8) is easily deduced:

$$(1.9) \quad h(x) = 1 + xh(x)^3.$$

Springer [17] later generalized Longyear’s argument to obtain an explicit formula for $H_{[n]}(\mathbf{i})$. His result is recovered here as Theorem 4.3. Also see [7] for an alternative derivation of the number of inequivalent k -cycle factorizations of a full cycle.

More recently, Goulden, Jackson, and Latour [10] counted inequivalent 2-cycle factorizations of any permutation $\pi \in \mathcal{C}_{[n,m]}$, proving that

$$(1.10) \quad \tilde{\Psi}_{2,2}(x_1, x_2) = \log\left(1 + x_1x_2h(x_1)h(x_2)\frac{h(x_1) - h(x_2)}{x_1 - x_2}\right),$$

where h is defined by (1.9). Their method again relies on commutation to canonical form, with the additional aid of a clever combinatorial construction and an intricate inclusion-exclusion argument.

1.3 Outline of the Paper and Statement of Results

The primary goal of this paper is to introduce a new technique in the enumeration of both cycle factorizations and their equivalence classes under commutation. We believe our approach to be of interest for two principal reasons: first, it conveniently allows one to ignore the fine detail of factorizations, *i.e.*, element-wise analysis, and focus on the grander structure, and second, it clarifies the structural parallels between the enumeration of factorizations and their equivalence classes. What follows is a brief overview of the paper highlighting our main results.

In Section 2 the reader is introduced to various constructs and conventions used extensively throughout. Our analysis of cycle factorizations then begins in Section 3, where we describe a graphical representation that allows (1.3) to be viewed as a generating series for a special class of labelled planar maps. Cycle factorizations of full cycles are then seen to correspond with particularly simple maps, namely *cacti*, which are a natural generalization of trees. This leads to the recovery of a known explicit formula for the number $H_{[n]}(\mathbf{i})$ of cycle factorizations of a full cycle (see Theorem 3.2). It also marks our first encounter with the series $w = w(x, \mathbf{q}, u) \in \mathbb{Q}[\mathbf{q}, u][[x]]$, defined as the unique solution of the functional equation

$$(1.11) \quad w = xe^{uQ(w)},$$

where $Q(z) \in \mathbb{Q}[\mathbf{q}][[z]]$ is given by

$$(1.12) \quad Q(z) := \sum_{k \geq 2} q_k z^{k-1}.$$

Clearly w is a generalization of the series s given by (1.6), making it no surprise that it plays a central role in our analysis. In particular, we shall present a graphical decomposition of maps (called *pruning*) that identifies an algebraic dependence of Ψ_m on w . This is the content of Theorem 3.5, and is the centerpiece of our method. By exploiting the pruning decomposition, we deduce the following extension of (1.5) to factorizations of arbitrary cycle index.

Theorem 1.1 *Let w and Q be defined as above and, for $i = 1, 2$, set $w_i = w(x_i, \mathbf{q}, u)$. Then*

$$\Psi_2(x_1, x_2, \mathbf{q}, u) = \log\left(\frac{w_1 - w_2}{x_1 - x_2}\right) - u \frac{w_1 Q(w_1) - w_2 Q(w_2)}{w_1 - w_2}.$$

In the end, our proof of Theorem 1.1 still rests on an *ad hoc* enumeration that we have not been able to generalize. Thus a formulation of Ψ_m for arbitrary m remains out of reach. We believe that it will be more tedious than difficult to extend our methods to arrive at an expression for Ψ_3 , but this has only yet been done in a special case. We comment further on these developments in §3.6.

In Section 4 we turn to the enumeration of cycle factorizations up to the equivalence defined in §1.2. By modifying our graphical representation of cycle factorizations to allow for commutations of adjacent factors, we are led to Springer’s formula [17] for the number of inequivalent cycle factorizations of a full cycle (Theorem 4.3). We discover that the unique solution $\tilde{w} = \tilde{w}(x, \mathbf{q}, u) \in \mathbb{Q}[\mathbf{q}, u][[x]]$ of the functional equation

$$(1.13) \quad \tilde{w} = 1 + u\tilde{w}Q(x\tilde{w}^2)$$

plays a role directly analogous to that of w in the enumeration of ordered cycle factorizations. Again we develop a pruning decomposition (Theorem 4.4) from which we deduce the following generalization of (1.10) to factorizations of arbitrary cycle index.

Theorem 1.2 *Let \tilde{w} and Q be defined as above and, for $i = 1, 2$, set $\tilde{w}_i = \tilde{w}(x_i, \mathbf{q}, u)$. Then*

$$\tilde{\Psi}_2(x_1, x_2, \mathbf{q}, u) = \log\left(\frac{(x_1 \tilde{w}_1 - x_2 \tilde{w}_2)^2}{(x_1 - x_2)(x_1 \tilde{w}_1^2 - x_2 \tilde{w}_2^2)}\right).$$

Upon setting $u = q_2 = 1$ and $q_k = 0$ for $k \neq 2$, the defining equation (1.13) of \tilde{w} transforms to $\tilde{w} = 1 + x\tilde{w}^3$, thus identifying it with Longyear’s series $h(x)$ (see (1.9)). Under these same specializations, it is easy to check that Theorem 1.2 reduces to (1.10).

Again, we have been unable to extend Theorem 1.2 to give a general expression for $\tilde{\Psi}_m$. However, we have used our method to deduce a raw form of $\tilde{\Psi}_{3,2}$, thus extending the Goulden–Jackson–Latour result in a different direction. See §4.5 for further comments on this and related matters.

2 Preliminaries

The definitions and notational conventions described in this section are used extensively throughout the remainder. We caution the reader that some of these conventions are nonstandard.

2.1 Cyclic Lists

A *cyclic list* is an equivalence class under the relation identifying finite sequences that are cyclic shifts of one another. We write $(a_1, \dots, a_n)^\circ$ for the cyclic list with representative sequence (a_1, \dots, a_n) . Thus $(a_1, a_2, a_3, a_4)^\circ = (a_3, a_4, a_1, a_2)^\circ$. Some liberties will be taken with this notation; for instance, when considering the list $(a_1, a_2, a_3, a_4)^\circ$, we adopt the convention that a_5 is to be interpreted as a_1 , while $a_0 = a_4$, etc. A cyclic list $L = (a_1, \dots, a_n)^\circ$ of real numbers is *increasing* (respectively, *nondecreasing*) if one of its representative sequences is strictly increasing (non-decreasing). If $a_{i-1} \geq a_i$, then the pair (a_{i-1}, a_i) is called a *descent* of L .

2.2 Maps and Polymaps

Recall that a *planar map* (subsequently, a *map*) is a 2-dimensional cellular complex whose polyhedron is homeomorphic to the sphere. The 0-cells, 1-cells, and 2-cells of a map are its *vertices*, *edges*, and *faces*, respectively. An *isomorphism* of two maps is an orientation-preserving homeomorphism between their polyhedra which sends i -cells to i -cells and preserves incidence. We always consider isomorphic maps to be indistinguishable. If a map is labelled, then this definition of isomorphism is amplified to preserve all labels.

The *boundary walk* of a face f is the cyclic sequence $((v_0, e_0), \dots, (v_k, e_k))^\circ$ of alternating vertices and edges listed in order as they are encountered along a counterclockwise traversal of the boundary of f . (Counterclockwise here means that f is always kept to the left of the line of traversal.) A subsequence (e_{i-1}, v_i, e_i) consisting of two consecutive edges and their common incident vertex is called a *corner* of f .

A map is *2-coloured* if its faces have been painted black and white so that every edge is incident with both a black face and a white face (so no two similarly coloured faces are adjacent). We are interested in the special class of 2-coloured maps for which the boundary walk of every black face is a cycle, *i.e.*, contains no repeated vertices or edges. We call these *polymaps*, and make the following supporting definitions:

- The black faces of a polymap are called *polygons*. An m -gon is a black face of degree m .
- The white faces of a polymap are referred to simply as *faces*.
- A *corner* of a polymap always refers to a corner of a (white) face.
- The *rotator* of a vertex v in a polymap is the unique cyclic list of polygons encountered along a clockwise tour of small radius about v .

For example, a polymap with 7 polygons and 3 faces is illustrated in Figure 1A. The rotator of vertex v is $(p_1, p_2, p_3)^\circ$.




Figure 1: (A) A polymap with 7 polygons and 3 faces. (B) The descent structure of a polymap.




Figure 2: A 3-constellation on 5 vertices.

2.3 Descent Structure

Let \mathcal{M} be a polymap whose polygons are labelled with integers. We shall always regard the edges of \mathcal{M} as being labelled, with each edge inheriting its label from the unique polygon that it borders. Let f be a face of the \mathcal{M} with boundary walk $((v_0, e_0), \dots, (v_k, e_k))^\circ$. If $e_{i-1} \geq e_i$, then the pair (e_{i-1}, e_i) is a *descent* of f . We say v_i is *at* this descent, and the corner (e_{i-1}, v_i, e_i) is called a *descent corner*. The *descent set* of f is the set $\{v_i : e_{i-1} \geq e_i\}$ of all vertices at descents of f . The cyclic list obtained by listing the vertices of this set in the order in which they appear along the boundary walk of f is called the *descent cycle* of f .

Example 2.1 Consider the outer face f of the polymap in Figure 1B. The cyclic list of edge labels encountered along its boundary walk is $(1, 2, 3, 3, 2, 3, 1, 3, 4)^\circ$, so f has four descents, namely $(3, 3), (3, 2), (3, 1), (4, 1)$. The hollow vertices a, b, c , and d are at these descents, so the descent set of f is $\{a, b, c, d\}$ and its descent cycle is $(a, b, c, d)^\circ$. The descent corners of the other faces are marked with crosses.

2.4 Constellations

An r -*constellation* on n vertices is a polymap whose vertices are distinctly labelled $1, 2, \dots, n$, and whose polygons are labelled $1, \dots, r$, such that the rotator of every vertex is $(1, 2, \dots, r)^\circ$. Figure 2 illustrates a 3-constellation on five vertices.

Our interest in this special class of polymaps stems from a well-known connection between maps and factorizations. A complete description of this correspondence can

be found in [1] (albeit in dual form²), so here it will suffice to outline a version particularly suited to our needs.

Proposition 2.2 *Minimal transitive factorizations in \mathfrak{S}_n with r factors are in bijection with r -constellations on n vertices. In particular, factorizations of π correspond with constellations whose decent cycles are exactly the cycles of π .*

Sketch of the Proof Let \mathcal{C} be an r -constellation on n vertices. From \mathcal{C} define permutations $\sigma_1, \dots, \sigma_r$ as follows: to form σ_i , first obtain a collection of disjoint cycles by listing the vertices of each polygon labelled i as they appear in clockwise order around its perimeter, and then let σ_i be the product of these cycles. (See Example 2.3, below.)

Now set $\pi = \sigma_r \cdots \sigma_1$. Regarding π as the sequential product of $\sigma_1, \dots, \sigma_r$, observe that it acts on a vertex v to move it clockwise around the polygons of \mathcal{C} following edges labelled $1, 2, \dots, r$, in turn. That is, v starts at the descent corner $(r, v, 1)$ of some face and is moved by π along the boundary walk until it lands at the next descent corner $(r, v', 1)$ of that face. Since $\pi(v) = v'$, the cycles of π are the descent cycles of \mathcal{C} .

Set $F = (\sigma_r, \dots, \sigma_1)$. We claim $\mathcal{C} \leftrightarrow F$ is the desired correspondence. Clearly F is a factorization of π , and it is transitive because \mathcal{C} is connected. Note that \mathcal{C} has n vertices and nr edges. If $\pi \in \mathcal{C}_\alpha$ and $\sigma_i \in \mathcal{C}_{\beta_i}$ for all i , then \mathcal{C} has $\ell(\alpha) + \sum_i \ell(\beta_i)$ total faces (white faces + polygons). Since \mathcal{C} is planar, Euler’s polyhedral formula (#vertices – #edges + #faces = 2) gives equality in (1.1), so F is minimal transitive. ■

Example 2.3 The 3-constellation of Figure 2 corresponds to the factorization $F = (\sigma_3, \sigma_2, \sigma_1)$ of $\pi = (1\ 2\ 4)(3)(5)$, where $\sigma_1 = (1\ 5)(2\ 4\ 3)$, $\sigma_2 = (1)(2\ 5)(3)(4)$, and $\sigma_3 = (1\ 5\ 3)(2)(4)$. Note that vertices 1, 2, and 4 are at descents of the outer face, and the cycle $(1\ 2\ 4)$ of π coincides with the descent cycle of that face.

3 Cycle Factorizations

In this section we shall introduce a convenient graphical representation of cycle factorizations as decorated polymaps, and proceed to give a decomposition of these polymaps that can simplify their enumeration. The section culminates in a proof of Theorem 1.1.

3.1 Proper and α -Proper Polymaps

Let $F = (\sigma_r, \dots, \sigma_1)$ be a cycle factorization and let \mathcal{C} be the constellation corresponding to F through the bijection of Proposition 2.2. Then exactly r of the polygons of \mathcal{C} are of degree two or more, while the remainder are 1-gons bounded by loops. Let \mathcal{M}_F be the polymap obtained by removing all 1-gons. For example, Figure 3A shows \mathcal{M}_F when F is the factorization

$$(1\ 5\ 7\ 8\ 4)(2)(3\ 9\ 6) = (2\ 7\ 8\ 6) \cdot (2\ 6\ 3\ 9) \cdot (1\ 5) \cdot (4\ 5\ 8) \cdot (5\ 9).$$

²We have borrowed the term “constellation” from [1], though we caution that the term is used there for an object that is dual to those considered here.




Figure 3: (A) The polymap of a cycle factorization. (B) A (3, 1, 5)-proper polymap.

Notice that all rotators of \mathcal{M}_F are increasing, since every rotator of \mathcal{C} is $(1, \dots, r)^\circ$. In fact, the removal of loops does not affect descent structure, in the sense that \mathcal{M}_F has precisely the same descent cycles as \mathcal{C} . Moreover, \mathcal{C} is easily recovered from \mathcal{M}_F by attaching loops in the unique manner which makes each rotator $(1, \dots, r)^\circ$.

We say a loopless polymap is *proper* if its polygons are labelled with distinct integers so that all rotators are increasing. (The polygon labels are assumed to be $1, \dots, r$, for some r , unless otherwise specified.) Let the *poly-index* of a loopless polymap be the vector $\mathbf{i} = (i_2, i_3, \dots)$, where i_k is the number of k -gons it contains. With these definitions, we have the following.

Proposition 3.1 *The mapping $F \mapsto \mathcal{M}_F$ is a bijection between cycle factorizations in \mathfrak{S}_n and proper polymaps with vertices labelled $1, \dots, n$. In particular, if F is a factorization of π of cycle index \mathbf{i} , then \mathcal{M}_F is a polymap of poly-index \mathbf{i} whose descent cycles are precisely the cycles of π .*

Let $\alpha = (\alpha_1, \dots, \alpha_m)$ be a composition, and let $\mathcal{P}_\alpha(\mathbf{i})$ be the set of vertex-labelled proper polymaps of poly-index \mathbf{i} whose descent sets are $\mathbb{N}_\alpha^1, \dots, \mathbb{N}_\alpha^m$, where

$$(3.1) \quad \mathbb{N}_\alpha^j := \{k \in \mathbb{N} : \alpha_1 + \dots + \alpha_{j-1} < k \leq \alpha_1 + \dots + \alpha_j\}.$$

Under the mapping $F \mapsto \mathcal{M}_F$, members of $\mathcal{P}_\alpha(\mathbf{i})$ correspond with factorizations of permutations whose orbits are $\mathbb{N}_\alpha^1, \dots, \mathbb{N}_\alpha^m$. Since there are $\prod_j (\alpha_j - 1)!$ such permutations, each admitting the same number of factorizations of given cycle index, we conclude that $|\mathcal{P}_\alpha(\mathbf{i})| = H_\alpha(\mathbf{i}) \cdot \prod_j (\alpha_j - 1)!$.

Let us say a proper polymap is α -*proper* if its faces are labelled $1, \dots, m$ so that face j has α_j descents, for $1 \leq j \leq m$. (See Figure 3B, for example.) Notice that any α -proper polymap of poly-index \mathbf{i} can be transformed into a member of $\mathcal{P}_\alpha(\mathbf{i})$ by first labelling its vertices in any of the $\prod_j \alpha_j!$ ways that make \mathbb{N}_α^j the descent set of face j for $j = 1, \dots, m$, and then stripping face labels. In fact, no members of $\mathcal{P}_\alpha(\mathbf{i})$ are duplicated in this process when $\ell(\alpha) \geq 2$, since the face and polygon labels of α -proper polymaps preclude nontrivial automorphisms.³ So for $\ell(\alpha) \geq 2$ we have $|\mathcal{P}_\alpha(\mathbf{i})| = M_\alpha(\mathbf{i}) \cdot \prod_j \alpha_j!$, where $M_\alpha(\mathbf{i})$ is the number of α -proper polymaps of poly-index \mathbf{i} .

³This is not so when $\ell(\alpha) = 1$, since a polymap composed of a single polygon has rotational symmetry.




Figure 4: Decomposition of a rooted cactus.

Comparing the expressions above for $|\mathcal{P}_\alpha(\mathbf{i})|$ gives $H_\alpha(\mathbf{i}) = \alpha_1 \cdots \alpha_m M_\alpha(\mathbf{i})$ when $\ell(\alpha) \geq 2$. Thus (1.3) becomes

$$(3.2) \quad \Psi_m(\mathbf{x}, \mathbf{q}, u) = \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} M_\alpha(\mathbf{i}) \mathbf{q}^{\mathbf{i}} \mathbf{x}^\alpha \frac{u^{r(\mathbf{i})}}{r(\mathbf{i})!}, \quad \text{for } m \geq 2.$$

That is, $\Psi_m(\mathbf{x}, \mathbf{q}, u)$ is the generating series for α -proper polymaps, where u is an exponential marker for labelled polygons, x_j marks descents of face j (for $1 \leq j \leq m$), and q_k records the number of k -gons (for $k \geq 2$).

3.2 Proper Cacti and Factorizations of Full Cycles

A *cactus* is a polymap with only one face. Hence Proposition 3.1 implies $H_{[n]}(\mathbf{i})$ is the number of proper cacti of poly-index \mathbf{i} whose sole faces have descent cycle $(1\ 2 \cdots n)$. However, observe that the fixed descent cycle of such a cactus forces all vertex labels once the location of one is known. Thus $H_{[n]}(\mathbf{i})$ is the number of *vertex-rooted* proper cacti of poly-index \mathbf{i} . (That is, the root marks the location of a canonical label, say 1.)

Let $w = w(x, \mathbf{q}, u)$ be the generating series for vertex-rooted proper cacti with respect to total vertices (marked by x), labelled polygons (marked by u), and poly-index (marked by \mathbf{q}). That is,

$$(3.3) \quad w := \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} H_{[n]}(\mathbf{i}) \mathbf{q}^{\mathbf{i}} x^n \frac{u^{r(\mathbf{i})}}{r(\mathbf{i})!} = x \frac{d}{dx} \Psi_1(x, \mathbf{q}, u).$$

Though this may appear to be in conflict with our earlier definition (1.11) of w , we now give a combinatorial decomposition of cacti to show that the two definitions coincide.

Suppose that the root v of a rooted proper cactus C is incident with m polygons. Detaching these polygons from v results in a collection $\{C_1, \dots, C_m\}$ of rooted proper cacti, where the root of each C_i is incident with only one polygon. (See Figure 4.) The ordering of C_1, \dots, C_m around v need not be recorded as it can be de-

duced from the increasing rotator condition. Thus

$$(3.4) \quad w = x \sum_{m \geq 0} \frac{\bar{w}^m}{m!} = xe^{w\bar{w}},$$

where the series $\bar{w} = \bar{w}(x, \mathbf{q}, u)$ counts rooted proper cacti such as C_i in which x marks only *nonroot* vertices. But if the root of C_i is incident with a k -gon, then removal of this polygon leaves a $(k - 1)$ -tuple C_i^1, \dots, C_i^{k-1} of rooted proper cacti. (See Figure 4, right-hand side.) This accounts for a contribution $uq_k w^{k-1}$ to \bar{w} , and summing over k yields $\bar{w} = uQ(w)$, where Q is defined as in (1.12). Thus (3.4) gives $w = xe^{uQ(w)}$, in agreement with (1.11).

The following result was originally proved bijectively by Springer [17]. In fact, the correspondence between factorizations and cacti employed here was also developed in [17], but from a less general point of view.

Theorem 3.2 (Springer [17]) *Let $\mathbf{i} = (i_2, i_3, \dots)$ be a sequence of nonnegative integers and set $r = r(\mathbf{i}) = i_1 + i_2 + \dots$. Then the number of cycle factorizations of $(1\ 2 \ \dots \ n)$ with cycle index \mathbf{i} is*

$$H_{[n]}(\mathbf{i}) = \frac{n^{r-1} r!}{\prod_{k \geq 2} i_k!}$$

in the case that $n + r - 1 = \sum_{k \geq 2} ki_k$, and zero otherwise.

Proof From (3.3) we have $H_{[n]}(\mathbf{i}) = r! [x^n \mathbf{q}^{\mathbf{i}} u^r] w(x, \mathbf{q}, u)$. This coefficient can be obtained from (1.11) through a straightforward application of Lagrange inversion [5]. An alternative proof that relies on a Prüfer-like encoding of cacti can be found in [17]. ■

3.3 Smooth Polymaps, Cores, and Branches

Given (3.2), we now wish to count α -proper polymaps for general α with $\ell(\alpha) \geq 2$. To do so we invoke a technique we call *pruning*, whereby polymaps are simplified through the removal of cacti. In this section we lay the groundwork for this approach, which is then executed in the next.

A *leaf* of a polymap is a polygon that shares exactly one vertex with another polygon, and a polymap is *smooth* if it does not have any leaves. If \mathcal{M} is a polymap with at least two faces, then iteratively removing its leaves results in a unique smooth polymap called the *core* of \mathcal{M} and denoted by \mathcal{M}^c . Labels of \mathcal{M}^c are inherited from \mathcal{M} in the obvious way. See Figure 5, for example. If p is a polygon of \mathcal{M} sharing only one vertex v with \mathcal{M}^c , then separating p from v results in two components, one of which is a rooted cactus B whose root vertex is incident only with p . We call B a *branch* of \mathcal{M} , and refer to p as its *root polygon*. The *base corner* of B is the corner of \mathcal{M}^c in which p was attached. (See Figure 5.)

Consider now the case when \mathcal{M} is proper, with at least two faces. Let B and p be as described above, and let f be the face of \mathcal{M}^c in which B was attached. Let




Figure 5: A proper polymap, its core, and one of its branches.

$((v_0, e_0), \dots, (v_k, e_k))^\circ$ be the boundary walk of f , where the indexing has been chosen so that the sequence (e_0, \dots, e_k) is as small as possible in lexicographic order. It is not difficult to see that this condition specifies e_0, \dots, e_k uniquely, and hence also determines a unique value of b , with $0 \leq b \leq k$, such that (e_{b-1}, v_b, e_b) is the base corner of B .

The key observation here is that the increasing rotator condition on \mathcal{M} implies $(e_{b-1}, p, e_b)^\circ$ is nondecreasing, and that this puts strong restrictions on the possible values of b , as is demonstrated by the following lemma.

Lemma 3.3 *Let $L = (a_0, \dots, a_k)^\circ$ be a cyclic list of real numbers with d descents. If $a \in \mathbb{R}$ is not in L , then there are exactly d values of i with $0 \leq i \leq k$ such that $(a_{i-1}, a, a_i)^\circ$ is nondecreasing.*

Proof Let \mathcal{P} be the polygonal path in the plane connecting the points

$$(0, a_0), \dots, (k, a_k), (k + 1, a_0),$$

in that order. Let s_i be the i -th step of \mathcal{P} . Call s_i an *up step* if $a_i > a_{i-1}$ and a *down step* otherwise. Note that $(a_{i-1}, a, a_i)^\circ$ is nondecreasing if and only if either $a_{i-1} < a < a_i$, or $a > a_{i-1} \geq a_i$, or $a_{i-1} \geq a_i > a$. Plainly, one of these conditions holds if and only if either (i) s_i is an up step which the line $y = a$ crosses, or (ii) s_i is a down step which this line misses. Since \mathcal{P} begins and ends at the same y -coordinate, the numbers of up steps and down steps crossed by $y = a$ must be equal. Thus the number of indices i for which (i) or (ii) is satisfied is equal to number of down steps of \mathcal{P} . Since down steps reflect descents of L , this completes the proof. See Figure 6 for an illustration; the dashed line is $a = 4$, and steps for which $(a_{i-1}, a, a_i)^\circ$ is nondecreasing have been thickened. ■

In particular, the lemma implies $(e_{s-1}, p, e_s)^\circ$ is increasing for exactly d values of s in the range $0 \leq s \leq k$, where d is the number of descents of face f . Let these possible s -values be $s_1 < \dots < s_d$. Then since $(e_{b-1}, p, e_b)^\circ$ is nondecreasing, we have $b = s_i$ for a unique $i \in \{1, \dots, d\}$. We call i the *index* of the branch B .




Figure 6: The proof of Lemma 3.3, with $L = (2, 7, 1, 3, 6, 5, 7, 3)^\circ$ and $a = 4$.

Example 3.4 Consider branch B of \mathcal{M} shown in Figure 5, and let f be the face of \mathcal{M}^c in which B was attached. The boundary walk of f is $((v_0, e_0), \dots, (v_5, e_5))^\circ$, where we choose $(e_0, \dots, e_5) = (4, 4, 14, 8, 12, 12)$ to be lexicographically minimal.

The base corner of B is $(e_2, v_3, e_3) = (14, v_3, 8)$ and the polygon attached to its root has label 2, so in the notation used above we have $b = 3$ and $p = 2$. Note that f has $d = 4$ descents, namely $(14, 8)$, $(12, 12)$, $(12, 4)$, and $(4, 4)$. In accordance with the lemma, $(e_{s-1}, p, e_s)^\circ$ is nondecreasing for exactly $d = 4$ values of $s \in \{0, \dots, 5\}$, namely $s_1 = 0, s_2 = 1, s_3 = 3$, and $s_4 = 5$. Since $b = s_3$, the index of B is $i = 3$.

To reiterate, Lemma 3.3 guarantees that a given branch can be attached to a face with d descents in *exactly* d positions so as to maintain increasing rotators, and the index of a branch records which of these positions it occupies.

3.4 Pruning Cacti

Let $S_\theta(\mathbf{i})$ denote the number of smooth θ -proper polymaps with poly-index \mathbf{i} . For $m \geq 2$ we define the following smooth polymap analogue of the series Ψ_m :

$$(3.5) \quad \Gamma_m(\mathbf{x}, \mathbf{q}, u) := \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\theta = n \\ \ell(\theta) = m}} S_\theta(\mathbf{i}) \mathbf{q}^{\mathbf{i}} \mathbf{x}^\theta \frac{u^{r(\mathbf{i})}}{r(\mathbf{i})!}.$$

The following theorem is the centrepiece of our approach to counting cycle factorizations. At first glance it appears to be the rather transparent statement that arbitrary polymaps can be viewed as the composition of smooth polymaps with cacti. However, it is worth emphasizing that we are concerned with the decomposition of *proper* polymaps, and therefore must maintain control of descent structure through the pruning process. It is Lemma 3.3, and the resulting “index of a branch”, that enables us to deal with this subtlety.

Theorem 3.5 Fix $m \geq 2$ and set $w_i = w(x_i, \mathbf{q}, u)$ for $i = 1, \dots, m$, where w is given by (1.11). Then $\Psi_m(\mathbf{x}, \mathbf{q}, u) = \Gamma_m(\mathbf{w}, \mathbf{q}, u)$, where $\mathbf{x} = (x_1, \dots, x_m)$ and $\mathbf{w} = (w_1, \dots, w_m)$.

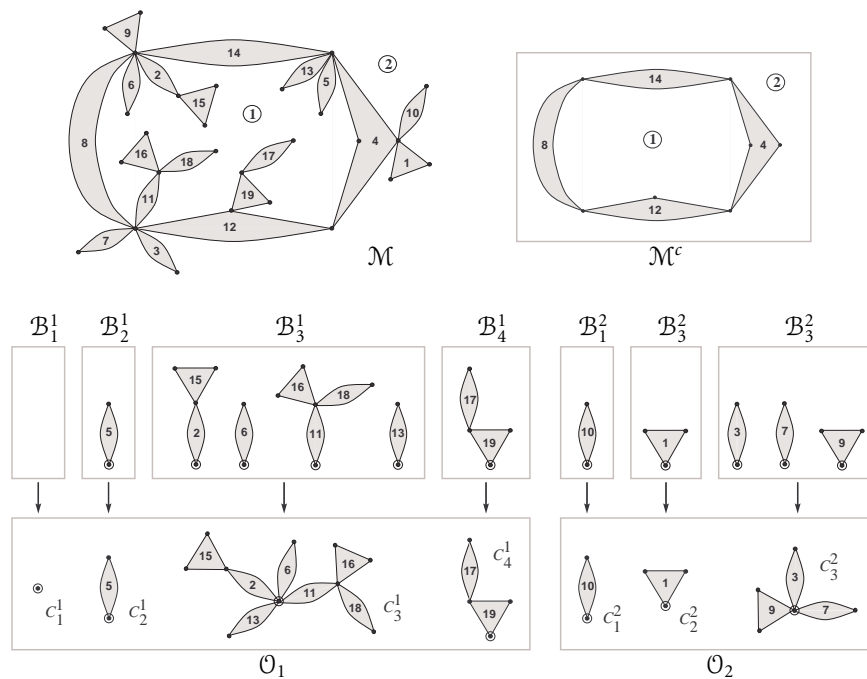


Figure 7: Pruning cacti from an α -proper polymap.

Proof Let α be an m -part composition, and let \mathcal{M} be an α -proper polymap. Suppose \mathcal{M}^c has θ_j descents in face j for $1 \leq j \leq m$. Then for all $1 \leq j \leq m$ and $1 \leq i \leq \theta_j$, let \mathcal{B}_i^j be the set of all branches of index i in face f of \mathcal{M} . Assemble all branches of \mathcal{B}_i^j into a single rooted proper cactus \mathcal{C}_i^j by identifying their root vertices, and let \mathcal{O}_j be the ordered forest $(\mathcal{C}_1^j, \dots, \mathcal{C}_{\theta_j}^j)$. See Figure 7 for an example of these constructions.

We claim that the mapping $\Phi: \mathcal{M} \mapsto (\mathcal{O}_1, \dots, \mathcal{O}_m, \mathcal{M}^c)$ is a polygon-preserving bijection between α -proper polymaps and tuples $(\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{S})$ satisfying the following properties.

- (a) \mathcal{F}_j is an ordered forest containing θ_j rooted proper cacti with a total of α_j vertices.
- (b) \mathcal{S} is a smooth θ -proper polymap, where $\theta = (\theta_1, \dots, \theta_m)$.
- (c) The polygon labels of \mathcal{S} and $\mathcal{F}_1, \dots, \mathcal{F}_m$ together partition $\{1, \dots, r\}$.

The fact that $\Phi(\mathcal{M})$ is indeed a tuple of this type follows from two simple observations: (i) a vertex is at a descent of face f of \mathcal{M} if and only if it is at a descent of face f of \mathcal{M}^c , and (ii) every non-root vertex of a branch of \mathcal{M} is at a descent of the face in which it lies.

The injectivity of Φ is immediate because two α -proper polymaps are isomorphic if and only if their cores are isomorphic and all branches in corresponding faces agree. To prove Φ is also surjective, let $(\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{S})$ satisfy (a) through (c), where




Figure 8: A smooth proper polymap.

$\mathcal{F}_j = (C_1^j, \dots, C_{\theta_j}^j)$ for $1 \leq j \leq m$. For fixed j and i with $1 \leq i \leq \theta_j$, the rooted cactus C_i^j can be viewed unambiguously as a collection of branches whose roots have been identified. Take any such branch B , and let p be the label of its root polygon. Now let $((v_0, e_0), \dots, (v_k, e_k))^\circ$ be the boundary walk of face j of \mathcal{S} , with (e_0, \dots, e_k) lexicographically minimal. Then $(e_0, \dots, e_k)^\circ$ has θ_j descents according to (b). Since (c) ensures that p is distinct from e_0, \dots, e_k , Lemma 3.3 implies $(e_{s-1}, p, e_s)^\circ$ is non-decreasing for exactly θ_j values of s , say $0 \leq s_1 < \dots < s_{\theta_j} \leq k$. Attach B to face j of \mathcal{S} at corner $(e_{s_i-1}, v_{s_i}, e_{s_i})$, doing so in the unique manner that leaves the rotator of v_{s_i} increasing. Repeat this process for all i, j , and B to iteratively build a polymap \mathcal{M} . Note that the order in which branches are attached is immaterial, so \mathcal{M} is well defined. The fact that \mathcal{M} is α -proper follows from conditions (a)–(c) and observations (i) and (ii) made above. Since $\Phi(\mathcal{M}) = (\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{S})$, we conclude by construction that Φ is surjective.

This bijection shows $M_\alpha(\mathbf{i})$ to be the number of tuples $(\mathcal{F}_1, \dots, \mathcal{F}_m, \mathcal{S})$ satisfying (a)–(c) and with poly-index \mathbf{i} . The result now follows by comparing (3.2) and (3.5), and recalling from §3.2 that w is the generating series for rooted proper cacti. ■

3.5 Factorizations of Permutations with Two Cycles

We now conclude this section with a proof of Theorem 1.1. The proof relies on Theorem 3.5 and the following well-known result, which can be found in [5]. Recall that a *circular permutation* of $\{1, \dots, n\}$ is a cyclic list $(a_1, \dots, a_n)^\circ$ such that $\{a_1, \dots, a_n\} = \{1, \dots, n\}$.

Lemma 3.6 *There are*

$$n! [x^{n-d} y^d] \log\left(\frac{x-y}{xe^y - ye^x}\right)$$

circular permutations of $\{1, \dots, n\}$ having exactly d descents.

Proof of Theorem 1.1 We first determine the series $\Gamma_2(z_1, z_2, \mathbf{q}, u)$ counting smooth α -proper polymaps with $\ell(\alpha) = 2$. Let \mathcal{S} be such a polymap, noting that this means \mathcal{S} is simply a closed chain of polygons, each incident with exactly two others. (See Figure 8. Vertex labels have been suppressed for clarity.) Let $(l_1, \dots, l_r)^\circ$ be the cyclic list of distinct polygon labels encountered along the boundary walk of face 1. Set $\gamma_i = (j_1, j_2)$ if the polygon with label l_i is a $(j_1 + j_2)$ -gon that has $j_1 - 1$ vertices

incident only with face 1 and $j_2 - 1$ incident only with face 2. Then \mathcal{S} is fully specified by $(l_1, \gamma_1, \dots, l_r, \gamma_r)^\circ$. For example, the polymap of Figure 8 corresponds with the cyclic list

$$(1, (1, 1), 3, (2, 1), 5, (3, 3), 2, (1, 2), 4, (3, 2))^\circ.$$

Let us say a vertex incident with only one face is *internal* to that face; the remaining vertices are *extremal*. Extremal vertices are hollow in Figure 8.

Clearly all vertices internal to a given face are at descents of that face. Therefore, temporarily ignoring extremal vertices, a polygon with $j_f - 1$ vertices internal to face f (for $f = 1, 2$) contributes $uq_{j_1+j_2}z_1^{j_1-1}z_2^{j_2-1}$ to $\Gamma_2(z_1, z_2, \mathbf{q}, u)$. Sum over $j_1, j_2 \geq 1$ to define

$$\delta := \sum_{j_1, j_2 \geq 1} uq_{j_1+j_2}z_1^{j_1-1}z_2^{j_2-1} = u \frac{Q(z_1) - Q(z_2)}{z_1 - z_2}.$$

Now the sole extremal vertex incident with polygons l_{i-1} and l_i is at a descent of face 1 if (l_{i-1}, l_i) is a descent of $(l_1, \dots, l_r)^\circ$, and at a descent of face 2 otherwise. Thus Lemma 3.6 gives

$$\begin{aligned} \Gamma_2(z_1, z_2, \mathbf{q}, u) &= \log\left(\frac{x - y}{xe^y - ye^x}\right) \Big|_{x=z_1\delta, y=z_2\delta} \\ &= \log\left(\frac{z_1 - z_2}{z_1 e^{-uQ(z_1)} - z_2 e^{-uQ(z_2)}}\right) - u\left(\frac{z_1 Q(z_1) - z_2 Q(z_2)}{z_1 - z_2}\right). \end{aligned}$$

The result follows immediately from Theorem 3.5 and identity (1.11). ■

3.6 Further Results

As mentioned in the introduction, we have been unable to extend the *ad hoc* enumeration applied in our proof of Theorem 1.1 to give a general formulation of Ψ_m for $m > 3$. The problem, of course, is that the structure of smooth polymaps on m faces gets far more complicated as m increases; in particular, polygons can be incident with three or more faces. The case $m = 3$ appears within reach (though tedious), and the more refined enumeration offered by Ψ_3 (as compared with Goulden and Jackson’s series $\Psi_{3,k}$) may provide some insight into the more general structure.

We have had somewhat more success in applying our techniques to analyze the Hurwitz problem. Observe that the polymaps associated with 2-cycle factorizations (through Proposition 3.1) can be simplified by collapsing 2-gons to single edges, thereby eliminating all polygons and leaving *maps* in the ordinary sense. In particular, cacti become trees, so §3.2 gives a bijection between 2-cycle factorizations of $(1\ 2 \ \dots \ n)$ and rooted trees with edges labelled $1, \dots, n - 1$. These trees, in turn, correspond with trees on n labelled vertices, as can be seen by “pushing” edge labels away from the root onto vertices and assigning label n to the root. Since there are n^{n-2} trees on n labelled vertices, we have an elegant bijective proof of Hurwitz’s formula (1.4) in the case $m = 1$. This bijection is equivalent to that given by Moszkowski [16].

In fact, the argument used in §3.2 to associate factorizations of full cycles with rooted cacti applies more generally to give a bijection between cycle factorizations of

any $\pi \in \mathcal{C}_\alpha$ and α -proper polymaps with one distinguished descent in each face. We have used this correspondence (restricted to 2-cycle factorizations) together with the pruning bijection established in Theorem 3.5 to give the first bijective proofs of (1.4) when $m = 2$ and $m = 3$. (Details can be found in [14].) Thus the simplicity of Hurwitz's formula in these cases is explained by a corresponding simplicity in the structure of certain smooth maps with two and three faces. We have as yet been unable to extend our methods to $m \geq 4$, but we hope that the shift in viewpoint afforded by pruning can be further exploited in this direction.

We should mention that the techniques employed here apply equally well to transitive factorizations $F = (\sigma_r, \dots, \sigma_1)$ not constrained by the minimality condition. If $\sigma_i \in \mathcal{C}_{\beta_i}$ and $\sigma_r \cdots \sigma_1 \in \mathcal{C}_\alpha$, then a parity argument applied to (1.1) shows $nr - \sum_{i=1}^r \ell(\beta_i) = n + \ell(\alpha) - 2 + 2g$ for a unique integer $g \geq 0$ called the *genus* of F . Thus minimal transitive factorizations are of genus 0. Unsurprisingly, genus g factorizations correspond with genus g polymaps, and many of the results described here (in particular, Theorem 3.5 and Theorem 4.4 of the next section) have obvious higher genus analogues. Our focus on minimal transitive factorizations reflects our belief that understanding the combinatorics in genus 0 is key to an understanding in all genera.

The graphical approach can also be used to study what we call β -factorizations, which are minimal transitive factorizations $(\rho, \tau_r, \dots, \tau_1)$ such that $\rho \in \mathcal{C}_\beta$ and the factors τ_1, \dots, τ_r are transpositions. Thus $[1^n]$ -factorizations are synonymous with 2-cycle factorizations. Determining the number of β -factorizations of a given $\pi \in \mathcal{C}_\alpha$ is known as the *double Hurwitz problem*. It has been the object of recent attention because of known and conjectural links with geometry, in particular intersection theory; further information can be found in [12]. Though we have recovered a number of the results of [12] through combinatorial methods, we have not made further headway.

4 Inequivalent Cycle Factorizations

We now turn to the enumeration of equivalence classes of cycle factorizations under the commutation relation introduced in §1.2. The graphical methods developed in the previous section will be reconsidered in this new context, and an appropriate pruning mechanism will be developed. The section ends with a proof of Theorem 1.2.

4.1 Marked and α -Marked Polymaps

Our starting point is the observation that two cycle factorizations are equivalent if and only if they have precisely the same factors, and the factors moving any given symbol appear in the same order in both factorizations. Interpreting this through the lens of our graphical correspondence $F \mapsto \mathcal{M}_F$ (Proposition 3.1), we see that permissible commutations of a cycle factorization F correspond with relabellings of the polygons of \mathcal{M}_F that preserve the relative order of the polygons incident with any given vertex.

In particular, *two factorizations are equivalent if and only if their corresponding polymaps have the same descent structure*. Thus the equivalence class $[F]$ containing F




Figure 9: Polymaps corresponding to equivalent cycle factorizations.

is naturally represented by the polymap $\tilde{\mathcal{M}}_F$ that results from stripping the polygon labels of \mathcal{M}_F and instead recording only the location of descents.

Example 4.1 Consider these equivalent cycle factorizations of $(1\ 2\ 3)(4\ 5)(6\ 7\ 8)$:

$$F = ((2\ 8\ 6), (3\ 5\ 7), (4\ 6), (5\ 6), (1\ 4\ 3), (2\ 7)),$$

$$G = ((3\ 5\ 7), (2\ 8\ 6), (2\ 7), (4\ 6), (1\ 4\ 3), (5\ 6)).$$

Factors $(5\ 6)$, $(4\ 6)$, and $(2\ 8\ 6)$ move symbol 6, and they appear in this same right-to-left order in both factorizations. The corresponding polymaps \mathcal{M}_F and \mathcal{M}_G are shown in Figures 9A and 9B, respectively, where vertex labels have been placed in descent corners. Note that the descent structure of these polymaps is identical. The class $[F]$ (equivalently, $[G]$) is represented by the polymap $\tilde{\mathcal{M}}_F$ (equivalently, $\tilde{\mathcal{M}}_G$) shown in Figure 9C, where again the location of descents is recorded by the placement of vertex labels.

Define a *marked* polymap to be a loopless polymap \mathcal{M} in which certain corners have been distinguished subject to the requirement that every vertex is at exactly one distinguished corner. A marked polymap \mathcal{M} is *properly marked* if its distinguished corners coincide with the descent corners under some polygon labelling of \mathcal{M} . That is, properly marked polymaps are exactly those objects that result from marking descents and deleting polygon labels from proper polymaps. (Note that not every marked polymap is properly marked.) We refer to the distinguished corners of a properly marked polymap as *descent corners*, and define terms such as *descent cycle* in the obvious way.

See Figure 10A for an example of a properly marked polymap. (Two polygon labellings which validate the indicated descent corners are shown in Figure 9.) It will be convenient to redefine the *rotator* of a vertex v in a properly marked polymap to be the tuple (P_1, \dots, P_m) of polygons incident with v , listed in order as they are encountered along a clockwise tour about v beginning in the unique descent corner containing v . In Figure 10B we have indicated polygons p_1 , p_2 , and p_3 such that the rotator of vertex v is (p_1, p_2, p_3) .

Properly marked polymaps will play the same role in the enumeration of inequivalent factorizations to follow as proper polymaps did in our approach to counting




Figure 10: (A,B) A properly marked polymap. (B) The rotator of v is (p_1, p_2, p_3) . (C) A $(3, 2, 3)$ -marked polymap.

ordered factorizations. Our observations thus far are summarized by the following analogue of Proposition 3.1.

Proposition 4.2 *The mapping $[F] \mapsto \tilde{\mathcal{M}}_F$ is a bijection between equivalence classes of cycle factorizations and vertex-labelled properly marked polymaps. If F is a factorization of π of cycle index \mathbf{i} , then $\tilde{\mathcal{M}}_F$ is of poly-index \mathbf{i} with descent cycles equal to the cycles of π .*

Let $\alpha = (\alpha_1, \dots, \alpha_m) \models n$. An α -marked polymap is a vertex-labelled properly marked polymap whose faces are labelled $1, \dots, m$ so that face j has descent set \mathbb{N}_α^j , where \mathbb{N}_α^j is defined as in (3.1). For example, under the convention that the locations of vertex labels indicate descent corners, the polymap of Figure 10C is seen to be $(3, 2, 3)$ -marked. Of course, the face labels of α -marked polymaps are superfluous, since they are determined by the descent sets. They are included in the definition only as a matter of convenience.⁴

Let $\tilde{M}_\alpha(\mathbf{i})$ denote the number of α -marked polymaps with poly-index \mathbf{i} . Under Proposition 4.2, α -marked polymaps correspond with factorizations of permutations whose orbits are $\mathbb{N}_\alpha^1, \dots, \mathbb{N}_\alpha^m$. There are $\prod_j (\alpha_j - 1)!$ such permutations, so $\tilde{M}_\alpha(\mathbf{i}) = H_\alpha(\mathbf{i}) \cdot \prod_j (\alpha_j - 1)!$. Thus (1.7) becomes

$$(4.1) \quad \tilde{\Psi}_m(\mathbf{x}, \mathbf{q}, u) = \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} \tilde{H}_\alpha(\mathbf{i}) \mathbf{q}^{\mathbf{i}} \frac{\mathbf{x}^\alpha}{\alpha!} u^{r(\mathbf{i})}.$$

That is, $\tilde{\Psi}_m(\mathbf{x}, \mathbf{q}, u)$ is the generating series for α -marked polymaps, where u marks polygons, \mathbf{q} records poly-index, and x_j is an exponential marker for labelled vertices at descents of face j .

4.2 Marked Cacti and Inequivalent Factorizations of Full Cycles

Observe that every marked cactus is properly marked. The argument used in §3.2 to show that $H_{[n]}(\mathbf{i})$ is the number of vertex-rooted proper cacti of poly-index \mathbf{i} therefore

⁴Note the contrast between this definition of α -marked polymaps and that of α -proper polymaps in §3.1. In the latter case, *only* face labels were used. Here we do not have that luxury, because face-labelled marked polymaps can admit nontrivial automorphisms, even when two or more faces are present.




Figure 11: Decomposition of a descent-marked cactus.

applies *mutatis mutandis* to identify $\tilde{H}_{[n]}(\mathbf{i})$ as the number of vertex-rooted marked cacti of poly-index \mathbf{i} . Thus the series $\tilde{w} = \tilde{w}(x, \mathbf{p}, u)$ defined by

$$(4.2) \quad \tilde{w} := \sum_{n \geq 1} \sum_{\mathbf{i} \geq 0} \tilde{H}_{[n]}(\mathbf{i}) x^{n-1} \mathbf{q}^{\mathbf{i}} u^{r(\mathbf{i})} = \frac{d}{dx} \tilde{\Psi}_1(x, \mathbf{q}, u)$$

counts such cacti with respect to the number of *non-root* vertices (marked by x), polygons (marked by u), and poly-index (marked by \mathbf{q}). As in §3.2, we now consolidate this definition of \tilde{w} with the one given in the introduction by describing a decomposition of marked cacti that proves their generating series satisfies (1.13).

Let C be a rooted marked cactus, and suppose its root vertex has rotator (p_1, \dots, p_m) . Detach p_1 from C to obtain a rooted marked cactus C' whose root has rotator (p_2, \dots, p_m) , as shown in Figure 11. (If $m = 1$, then C' consists of a single vertex.) Now focus on polygon p_1 . Suppose p_1 is a k -gon, and let v be one of its $k - 1$ non-root vertices. Let $(p_v^1, \dots, p_v^r, p_1, p_v^{r+1}, \dots, p_v^s)$ be the rotator of v , where the degenerate cases $r = 0$ and $s = r$ are possible. Detach the polygons p_v^i from v to form two cacti, C_v^1 and C_v^2 , whose roots have rotators (p_v^1, \dots, p_v^r) and $(p_v^{r+1}, \dots, p_v^s)$, respectively. See Figure 11, far right.

Thus C decomposes into C' together with k -gon p_1 and an ordered list of triples (v, C_v^1, C_v^2) , one for each of the $k - 1$ non-root vertices of p_1 . Allowing for all possible k gives

$$\tilde{w} = 1 + \sum_{k \geq 2} \tilde{w} \cdot u q_k \cdot (x \tilde{w}^2)^{k-1} = 1 + u \tilde{w} Q(x \tilde{w}^2),$$

where the summand 1 accounts for the case in which C consists of a single vertex, and Q is defined as in (1.12). This is in agreement with (1.13), as claimed.

The following result first appeared in [17]. The bijective proof given there relied on the commutation of cycle factorizations into canonical forms, together with a correspondence between these canonical forms and a class of trees for which enumerative formulae are known. In hindsight, our use of marked cacti can be viewed as a high level graphical interpretation of this bijection that bypasses the element-wise analysis necessary to establish the canonical forms.

Theorem 4.3 (Springer [17]) Let $\mathbf{i} = (i_2, i_3, \dots)$ be a sequence of nonnegative integers, not all zero, and set $r = r(\mathbf{i}) = i_2 + i_3 + \dots$. Then the number of inequivalent cycle factorizations of $(1\ 2\ \dots\ n)$ of cycle index \mathbf{i} is

$$\tilde{H}_{[n]}(\mathbf{i}) = \frac{(2n + r - 2)!}{(2n - 1)! \prod_{k \geq 2} i_k!}$$

in the case that $n + r - 1 = \sum_{k \geq 2} ki_k$, and zero otherwise.

Proof Set $v = \tilde{w} - 1$. Then (4.2) and (1.13) give $\tilde{H}_{[n]}(\mathbf{i}) = [x^{n-1}u^r\mathbf{q}^{\mathbf{i}}] v$ with $v = u(1 + v)Q(x(1 + v)^2)$. By Lagrange inversion [5] we have

$$\begin{aligned} [x^{n-1}u^r\mathbf{q}^{\mathbf{i}}] v &= [x^{n-1}\mathbf{q}^{\mathbf{i}}] \frac{1}{r} [\lambda^{r-1}] (1 + \lambda)^r Q(x(1 + \lambda)^2)^r \\ &= \frac{1}{r} [\lambda^{r-1}] (1 + \lambda)^{r+2n-2} \binom{r}{i_2, i_3, i_4, \dots}, \end{aligned}$$

and the result follows upon simplification. ■

Specializing the previous theorem to count inequivalent k -cycle factorizations entails setting $i_k = r$ and $i_j = 0$ for $j \neq k$. Doing so, we find that the number of inequivalent k -cycle factorizations of $(1\ 2\ \dots\ n)$ is

$$\frac{1}{2n - 1} \binom{2n + r - 2}{r},$$

when $n = 1 + r(k - 1)$ for some integer r . This formula first appeared in [7]. Setting $k = 2$ yields the result (1.8) of Longyear mentioned in the introduction.

4.3 Pruning Cacti

Let $\alpha = (\alpha_1, \dots, \alpha_m)$ be a composition. Define the *face degree sequence* of an α -marked polypmap to be (d_1, \dots, d_m) , where d_j is the degree of face j , for $1 \leq j \leq m$. Alternatively, d_j is the number of corners in face j . Let $\tilde{S}_\alpha(\mathbf{i}; \mathbf{d})$ denote the number of smooth α -marked polypmaps with poly-index \mathbf{i} and face degree sequence \mathbf{d} . Then, for $m \geq 1$, define

$$\tilde{\Gamma}_m(\mathbf{x}, \mathbf{t}, \mathbf{q}, u) := \sum_{n \geq 1} \sum_{\mathbf{i}, \mathbf{d} \geq \mathbf{0}} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} \tilde{S}_\alpha(\mathbf{i}; \mathbf{d}) \frac{\mathbf{x}^\alpha}{\alpha!} \mathbf{q}^{\mathbf{i}} \mathbf{t}^{\mathbf{d}} u^{r(\mathbf{i})}.$$

With these definitions we have the following marked polypmap analogue of Theorem 3.5.

Theorem 4.4 Let $m \geq 1$ and set $\tilde{w}_i = \tilde{w}(x_i, \mathbf{q}, u)$ for $1 \leq i \leq m$, where \tilde{w} is given by (1.13). Then $\tilde{\Psi}_m(\mathbf{x}, \mathbf{q}, u) = \tilde{\Gamma}_m(\mathbf{x} \circ \tilde{\mathbf{w}}, \tilde{\mathbf{w}}, \mathbf{q}, u)$, where $\mathbf{x} = (x_1, \dots, x_m)$, $\tilde{\mathbf{w}} = (\tilde{w}_1, \dots, \tilde{w}_m)$, and $\mathbf{x} \circ \tilde{\mathbf{w}} = (x_1\tilde{w}_1, \dots, x_m\tilde{w}_m)$.

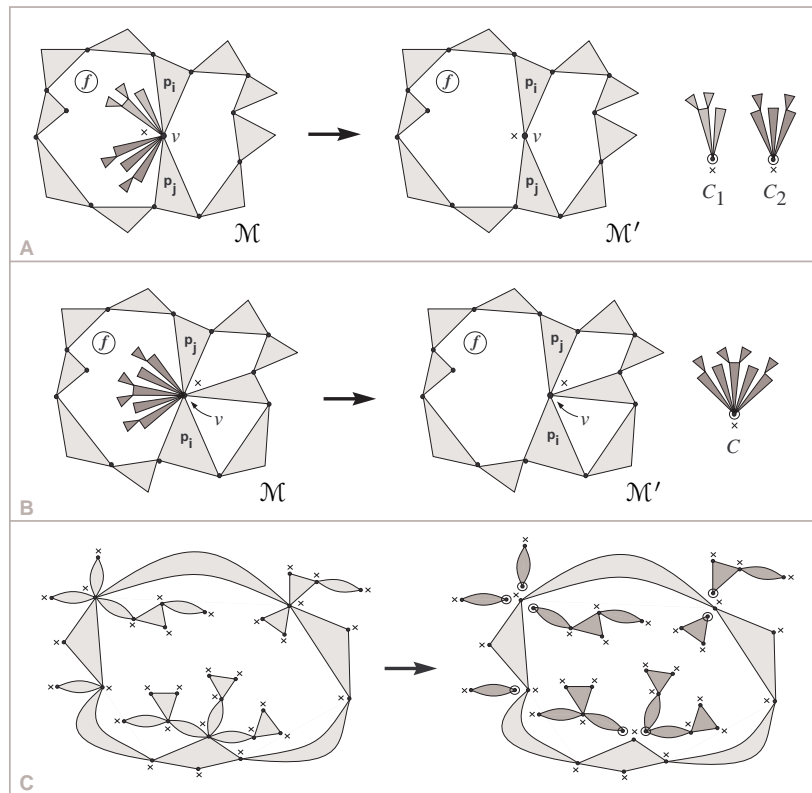


Figure 12: Pruning cacti from a marked polymap.

Proof Let \mathcal{M} be a properly marked polymap with at least two faces and let v be a vertex of \mathcal{M}^c incident with face f . Let (p_1, \dots, p_k) be the rotator of v . We wish to prune from \mathcal{M} all branches in f attached at v .

Suppose first that v is at a descent of f . Then there exist i, j with $1 \leq i \leq j \leq k$ such that $p_1, \dots, p_{i-1}, p_{j+1}, \dots, p_k$ are the root polygons of the branches attached to v in face f . (See Figure 12A. For clarity, most descent corners have not been indicated.) Detach these branches from \mathcal{M} to form a new polymap \mathcal{M}' and two rooted marked cacti, C_1 and C_2 , whose roots have rotators (p_1, \dots, p_{i-1}) and (p_{j+1}, \dots, p_k) , respectively. Note that cases $i = 1$ and $j = k$ yield trivial cacti C_1 and C_2 , respectively. The descent corner of f at v is inherited by \mathcal{M}' , and it is easy to check that this makes \mathcal{M}' properly marked. Figure 12A illustrates this pruning mechanism.

Now suppose v is not at a descent of f . Then there exist i, j with $1 \leq i < j \leq k$ such that p_{i+1}, \dots, p_{j-1} are the root polygons of the branches attached to v in f . (See Figure 12B.) Detach these branches from \mathcal{M} to form a properly marked polymap \mathcal{M}' and a rooted marked cactus C whose root has rotator $(p_{i+1}, \dots, p_{j-1})$. When $i = j$,

there are no branches attached to v in f , and C is trivial.

Repeating this pruning process for each incident face-vertex pair (v, f) of \mathcal{M} (that is, for each corner of \mathcal{M}) results in a smooth properly marked polymap \mathcal{S} and a collection of rooted marked cacti, each of which is naturally associated with a particular face of \mathcal{S} . In fact, if a face of \mathcal{S} is of degree D and has d descents, then it has $D - d$ non-descent corners and thus exactly $2d + (D - d) = D + d$ associated pruned cacti. Each non-root vertex of these cacti contributes one descent to the corresponding face of \mathcal{M} . The pruning of all cacti from a properly marked polymap is illustrated in Figure 12C.

To complete the proof, observe that if \mathcal{M} is α -marked, where $m = \ell(\alpha) \geq 2$, then its corners are distinguishable and thus the pruning process described above is reversible. Recall that $\tilde{w}(x, \mathbf{q}, u)$ counts rooted marked cacti, with x recording non-root vertices. Since rooting eliminates nontrivial automorphisms, these vertices can be freely labelled. From (4.1) it follows that $\tilde{\Psi}_m(\mathbf{x}, \mathbf{q}, u)$ is obtained from $\tilde{\Gamma}_m(\mathbf{z}, \mathbf{t}, \mathbf{q}, u)$ through the substitutions $z_j \mapsto x_j \tilde{w}_j$ and $t_j \mapsto \tilde{w}_j$, for $1 \leq j \leq m$. ■

4.4 Inequivalent Factorizations of Permutations with Two Cycles

We now use the pruning of marked polymaps to prove Theorem 1.2. As mentioned in the introduction, this theorem generalizes the main result of [10], and it is worth noting that the intricate inclusion-exclusion argument used there is avoided entirely by our method. In particular, the following lemma (which is a routine exercise in exponential generating series) is seen to be a more natural source of the logarithm appearing in (1.10).

Lemma 4.5 *For $m, n \geq 1$, let $c_{n,m}$ be the number of distinct necklaces made of n labelled white beads and m (independently) labelled black beads, counted up to rotational symmetry. Then*

$$(4.3) \quad \sum_{n,m \geq 1} c_{n,m} \frac{x^n}{n!} \frac{y^m}{m!} z^{n+m} = \log \left(1 + \frac{xyz^2}{1 - z(x + y)} \right).$$

Proof of Theorem 1.2 We begin by determining $\tilde{\Gamma}_2(z_1, z_2, t_1, t_2, \mathbf{q}, u)$. Let \mathcal{S} be a smooth α -marked polymap, where $\ell(\alpha) = 2$. Then \mathcal{S} is a closed chain of polygons, each incident with exactly two others. As in the proof of Theorem 1.1, we say a vertex incident with two polygons is *extremal*. Since \mathcal{S} is properly marked, at least one extremal vertex is at a descent of each face.

Let $L = (v_1, \dots, v_r)^\circ$ be the cyclic list of extremal vertices encountered along the boundary walk of face 1 of \mathcal{S} . Regard those v_i that are at descents of face 1 as white beads and those at descents of face 2 as black beads, so L corresponds with a necklace of the type counted by Lemma 4.5. (See Figure 13 for an illustration, where extremal vertices are indicated in grey, but vertex and bead labels are not shown.)

Let M be the monomial of $\tilde{\Gamma}_2(z_1, z_2, t_1, t_2, \mathbf{q}, u)$ corresponding to \mathcal{S} . An extremal vertex of \mathcal{S} contributes the factor $z_1 t_1 t_2$ to M if it is at a descent of face 1, and contributes $z_2 t_1 t_2$ otherwise. A polygon of \mathcal{S} with $j_f - 1$ vertices incident only with




Figure 13: A smooth marked polymap and its associated necklace.

face f , for $f = 1, 2$, further contributes the factor $uq_{j_1+j_2}(t_1z_1)^{j_1-1}(t_2z_2)^{j_2-1}$. Thus $\tilde{\Gamma}_2(z_1, z_2, t_1, t_2, \mathbf{q}, u)$ is obtained from (4.3) by performing the substitutions

$$x \mapsto z_1t_1t_2, \quad y \mapsto z_2t_1t_2,$$

$$z \mapsto \sum_{j_1, j_2 \geq 1} uq_{j_1+j_2}(t_1z_1)^{j_1-1}(t_2z_2)^{j_2-1} = u \frac{Q(t_1z_1) - Q(t_2z_2)}{t_1z_1 - t_2z_2}.$$

But “Theorem 4.4 implies $\tilde{\Psi}_2(x_1, x_2, \mathbf{q}, u)$ ” is obtained from $\tilde{\Gamma}_2(z_1, z_2, t_1, t_2, \mathbf{q}, u)$ through the substitutions $z_i \mapsto x_i\tilde{w}_i$ and $t_i \mapsto \tilde{w}_i$. Performing this chain of substitutions leaves

$$\tilde{\Psi}_2(x_1, x_2, \mathbf{q}, u) = \log \left(1 + \frac{\delta^2 x_1 x_2 \tilde{w}_1 \tilde{w}_2}{1 - \delta(x_1 \tilde{w}_1 + x_2 \tilde{w}_2)} \right),$$

where

$$\delta := u\tilde{w}_1\tilde{w}_2 \frac{Q(x_1\tilde{w}_1^2) - Q(x_2\tilde{w}_2^2)}{x_1\tilde{w}_1^2 - x_2\tilde{w}_2^2}.$$

Theorem 1.2 now follows upon simplification, using (1.13) to write $uQ(x_i\tilde{w}_i^2) = 1 - \tilde{w}_i^{-1}$. ■

4.5 Further Results

We have been unable to generalize the proof of Theorem 1.2 to obtain expressions for $\tilde{\Psi}_m$ for any $m > 2$. However, Theorem 4.4 has been used [14] to give a generating series formulation for the number of inequivalent 2-cycle factorizations of permutations composed of three cycles. The expression for $\tilde{\Psi}_{3,2}$ thus obtained is far more complicated than the very succinct (1.10) and we have been unable to simplify it in a meaningful way. Finding a natural expression for $\tilde{\Psi}_{3,2}$ remains an intriguing open problem, as it would grant far greater insight into the structure of inequivalent factorizations in general. We note that an approach from the point of view of [10] involves significant technical hurdles that do not surface in the simpler known cases.

Inspired by the definition of β -factorizations (see §3.6) we can extend Theorem 4.3 in the following way. Let $\beta = [1^{j_1}2^{j_2}\dots] \vdash n$ and consider the set of minimal transitive factorizations $(\rho, \sigma_r, \dots, \sigma_1)$ of the full cycle $(1\ 2\ \dots\ n)$ such that $\rho \in \mathcal{C}_\beta$ and the σ_i are all cycles of length at least two. Say two such factorizations

are equivalent if one can be obtained from the other by commutation of adjacent disjoint *cycle factors* $\sigma_1, \dots, \sigma_m$. That is, the position of ρ is fixed. Then the number of inequivalent factorizations of this type with exactly i_k cycle factors being k -cycles for $k \geq 2$ is

$$\frac{n(r-1)! (\ell(\beta) - 1)!}{\prod_{k \geq 2} i_k! \cdot \prod_{k \geq 1} j_k!} \binom{r + \ell(\beta) + n - 2}{r - 1},$$

where $r = i_2 + i_3 + \dots$, subject to the necessary condition $\sum k i_k = r + \ell(\beta) - 1$. This result is readily proved through a refinement of the decomposition of marked cacti given in §4.2. Note that Theorem 4.3 is recovered by setting $\beta = [1^n]$.

We conclude by noting that our definition of equivalence of factorizations may not seem the most natural. For instance, one could allow all pairs of adjacent commuting factors to be interchanged, whether or not they are disjoint. A general statement of the corresponding enumerative problem might then be the following: for a multiset B of partitions of n , find the number of inequivalent minimal transitive factorizations of a given $\pi \in \mathfrak{S}_n$ whose factors have cycle types specified by B . We have not yet investigated this problem in any detail.

Acknowledgements The bulk of the research for this article was completed during the author's doctoral studies under the supervision of David Jackson, whose support and guidance has been greatly appreciated. Thanks also go to Ian Goulden for many helpful discussions.

References

- [1] M. Bousquet-Mélou and G. Schaeffer, *Enumeration of planar constellations*. Adv. in Appl. Math. **24**(2000), no. 4, 337–368.
- [2] J. A. Eidswick, *Short factorizations of permutations into transpositions*. Discrete Math. **73**(1989), no. 3, 293–243.
- [3] T. Ekedahl, S. Lando, M. Shapiro, and A. Vainshtein, *On Hurwitz numbers and Hodge integrals*. C. R. Acad. Sci. Paris Sér. I Math. **328**(1999), no. 12, 1175–1180.
- [4] V. Goryunov and S. Lando, *On enumeration of meromorphic functions on the line*. In: The Arnoldfest. Fields Inst. Commun. 24, American Mathematical Society, Providence, RI, 1999, pp. 209–223.
- [5] I. P. Goulden and D. M. Jackson, *Combinatorial Enumeration*. Wiley Interscience, New York, 1983.
- [6] ———, *The combinatorial relationship between trees, cacti and certain connection coefficients for the symmetric group*. European J. Combin. **13**(1992), no. 5, 357–365.
- [7] ———, *Symmetric functions and Macdonald's result for top connexion coefficients in the symmetric group*. J. Algebra **166**(1994), no. 2, 364–378.
- [8] ———, *Transitive factorisations into transpositions and holomorphic mappings on the sphere*. Proc. Amer. Math. Soc. **125**(1997), no. 1, 51–60.
- [9] ———, *Transitive factorizations in the symmetric group, and combinatorial aspects of singularity theory*. Europ. J. Combinatorics **21**(2000), no. 8, 1001–1016.
- [10] I. P. Goulden, D. M. Jackson, and F. G. Latour, *Inequivalent transitive factorizations into transpositions*. Canad. J. Math **53**(2001), no. 4, 758–779.
- [11] 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. **83**(2001), no. 3, 563–581.
- [12] ———, *Towards the geometry of double Hurwitz numbers*. Adv. Math. **198**(2005), no. 1, 43–92.
- [13] A. Hurwitz, *Über Riemann'sche Flächen mit gegebenen Verzweigungspunkten*. Math. Ann. **39**(1891), no. 1, 1–60.
- [14] J. C. Irving, *Combinatorial Constructions for Transitive Factorizations in the Symmetric Group*. Ph.D. thesis, University of Waterloo, 2004.
- [15] J. Q. Longyear, *A peculiar partition identity*. Discrete Math. **78**(1989), no. 1–2, 115–118.

- [16] P. Moszkowski, *A solution to a problem of Dénes: a bijection between trees and factorizations of cyclic permutations*. *European J. Combin.* **10**(1989), no. 1, 13–16.
- [17] C. M. Springer, *Factorizations, trees, and cacti*, Eighth International Conference on Formal Power Series and Algebraic Combinatorics, University of Minnesota, June 25–29, 1996, pp. 427–438.

Department of Mathematics and Computing Science, Saint Mary's University, Halifax, NS, B3H 3C3
e-mail: john.irving@smu.ca