Combinatorial Constructions for Transitive Factorizations in the Symmetric Group

by

John Irving

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2004

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as required by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

We consider the problem of counting *transitive factorizations* of permutations; that is, we study tuples $(\sigma_r, \dots, \sigma_1)$ of permutations on $\{1, \dots, n\}$ such that (1) the product $\sigma_r \cdots \sigma_1$ is equal to a given target permutation π, and (2) the group generated by the factors σ_i acts transitively on $\{1, \dots, n\}$. This problem is widely known as the *Hurwitz Enumeration Problem*, since an encoding due to Hurwitz shows it to be equivalent to the enumeration of connected branched coverings of the sphere by a surface of given genus with specified branching.

Much of our work concerns the enumeration of transitive factorizations of permutations into a minimal number of transposition factors. This problem has received considerable attention, and a formula for the number $c(\pi)$ of such factorizations of an arbitrary permutation π has been derived through various means. The formula is remarkably simple, being a product of well-known combinatorial numbers, but no bijective proof of it is known except in the special case where π is a full cycle. A major goal of this thesis is to provide further combinatorial rationale for this formula.

We begin by introducing an encoding of factorizations (into transpositions) as edge-labelled maps. Our central result is a bijection that allows trees to be "pruned" from such maps. This is shown to explain the appearance of factors of the form k^k in the aforementioned formula for $c(\pi)$. It also has the effect of shifting focus to the combinatorics of smooth maps (*i.e.* maps without vertices of degree one). By providing decompositions for certain smooth planar maps, we are able to give combinatorial evaluations of $c(\pi)$ when π is composed of up to three cycles.

Many of these results are generalized to factorizations in which the factors are cycles of any length. We also investigate the *Double Hurwitz Problem*, which calls for the enumeration of factorizations whose leftmost factor is of specified cycle type, and whose remaining factors are transpositions. Finally, we extend our methods to the enumeration of factorizations up to an equivalence relation induced by possible commutations between adjacent factors.

Acknowledgements

I owe a tremendous debt of gratitude to my thesis supervisor, David Jackson, for his unwavering support — academic, financial, and personal — throughout my years of graduate study. It seemed at times that his commitment and enthusiasm towards this project exceeded my own, and for his patience in this regard, and for the warmth and generosity he has always shown towards me, I am forever thankful. I feel very fortunate to have studied for so many years under the supervision of a friend.

I have also benefitted from the extensive support and generosity of Ian Goulden, who has always managed to find enough time in his busy schedule to lend mathematical insight and provide encouragement and guidance. His contributions and friendship have been greatly appreciated.

My early years as a graduate student were funded partly through a scholarship granted by the Natural Sciences and Engineering Research Council of Canada. I am grateful for their assistance.

I must acknowledge my many friends in the department, both faculty and fellow students, for creating a truly enjoyable social atmosphere in which to work. It seems everyone has shared in the laughs over the years. Special thanks should be payed to Sue, Kim, Fiona, and Marg for their exceptional kindness and for ensuring that my scholastic career did not come to an end as the result of a bureaucratic technicality.

Completing this thesis would not have been possible without the endless love and support of my family. Each of them, in their own way, has inspired me constantly for as long as I can remember. My good friends, too, deserve credit for keeping me sane throughout the process. Amps, Harry and Dave merit special mention; Amps because he is every bit as good a travelling partner as he is a mathematician, Harry because acknowledging him here might serve as a down-payment on my outstanding dinner tab, and Dave because he has been such a great friend that I would have completed my degree years ago had we not been so busy laughing.

My final words of thanks go to Maggie, who continues to be my best friend and coconspirator. She has been behind me at every turn as a source of limitless encouragement. I am unable to express my deep gratitude for her infinite patience and for the various sacrifices she has made to keep us so close while I pursued this goal.

Contents

1 Introduction **1**

1.1 Preamble . . . 1

1.2 Main Results . . . 2

1.3 Background Material and Notational Conventions . . . 3

1.3.1 Sets, Compositions and Partitions . . . 4

1.3.2 Generating Series and Lagrange Inversion . . . 4

1.3.3 Complete Symmetric Functions and Umbral Composition . . . 6

1.3.4 Cyclic Lists . . . 7

1.3.5 The Symmetric Group . . . 7

1.3.6 The Group Algebra of the Symmetric Group . . . 8

1.3.7 Graphs . . . 9

1.3.8 Maps . . . 10

2 Factorizations into Transpositions **13**

2.1 Introduction . . . 13

2.2 The Length of a Factorization . . . 14

2.2.1 Cuts and Joins . . . 14

2.2.2 Minimal Factorizations . . . 15

2.2.3 Components . . . 16

2.2.4 Transitive Factorizations . . . 19

2.2.5 Additional Notes . . . 20

2.3 Enumeration of Factorizations and Hurwitz's Problem . . . 20

2.3.1 Factorizations of a Prescribed Length . . . 21

2.3.2 Transitive Factorizations and Hurwitz Numbers . . . 22

2.3.3 Lagrangian Structure in the Hurwitz Series . . . 24

2.3.4 Labelled Trees 25
2.3.5 The Symmetrized Hurwitz Series 26
2.3.6 Geometry and Hurwitz Numbers 28
2.3.7 Additional Notes 29
2.4 Graphical Representation of Factorizations 31
2.4.1 Counting Minimal Factorizations 31
2.4.2 The Graph of a Factorization 33
2.4.3 Carriers and Orbits 34
2.4.4 The Map of a Factorization 36
2.4.5 Edge-Labelled Maps and Descent Structure 38
2.4.6 A Bijection Between Factorizations and Maps 40
2.4.7 Genus 0 Factorizations of Full Cycles 42
2.4.8 Genus 1 Factorizations of Full Cycles 43
2.4.9 Face-Labelled Maps 47
2.4.10 Properly Labelled Maps 49
2.4.11 Comments on Labelling 51
2.4.12 Additional Notes 52
2.5 Differential Equations for Labelled Maps 53
2.5.1 Decomposition of Planar Maps 53
2.5.2 A Change of Variables 56
2.6 Smooth Maps and Pruning Trees 59
2.6.1 Cores and Branches 59
2.6.2 Normally Indexed Boundary Walks 61
2.6.3 The Index of a Branch 62
2.6.4 Pruning Trees 64
2.7 Combinatorial Constructions for Smooth Maps 68
2.7.1 Two-Face Smooth Planar Maps 68
2.7.2 Attaching Edges to a Map 70
2.7.3 Two-Face Smooth Planar Maps Revisited 73
2.7.4 Three-Face Smooth Planar Maps 77
2.7.5 A Differential Decomposition for Smooth Planar Maps 84
2.8 Bijections Between Factorizations and Trees 88
2.8.1 Preliminaries 88

2.8.2 Factorizations of Class (n_1, n_2) 90
2.8.3 Factorizations of Class (n_1, n_2, n_3) 92

3 Generalizations **97**
3.1 Introduction 97
3.2 Graphical Representation of General Factorizations 98
3.2.1 Polymaps 98
3.2.2 Constellations 100
3.2.3 A Bijection Between Factorizations and Polymaps 103
3.2.4 Minimal Transitive Factorizations of Full Cycles 104
3.2.5 Suppression of Loops 106
3.2.6 Factorizations into Transpositions 108
3.2.7 Additional Notes 109
3.3 Cycle Factorizations 110
3.3.1 Preliminaries 110
3.3.2 Properly Labelled Polymaps 111
3.3.3 Minimal Transitive Cycle Factorizations of Full Cycles 114
3.3.4 Differential Equations for Planar Polymaps 116
3.3.5 Pruning Cacti 119
3.3.6 Two-Face Smooth Planar Polymaps 122
3.3.7 Attaching Digons to a Polymap 124
3.3.8 Additional Notes 128
3.4 The Double Hurwitz Problem 129
3.4.1 β-Factorizations 129
3.4.2 Hybrid Maps 130
3.4.3 Hybrid Cacti 132
3.4.4 Pruning Cacti 133
3.4.5 Combinatorial Constructions for Smooth Hybrid Maps 135
3.4.6 A Final Bijection 139
3.4.7 Additional Notes 143

4 Inequivalent Factorizations **145**
4.1 Introduction 145
4.2 Inequivalent Factorizations into Transpositions 146

4.2.1 Minimal Factorizations 146
4.2.2 Factorizations of a Prescribed Length 148
4.2.3 Transitive Factorizations 150
4.3 Inequivalent Cycle Factorizations 154
4.3.1 Factorizations of Full Cycles 154
4.3.2 Graphical Representation of Equivalence Classes 156
4.3.3 Descent-Marked Cacti 158
4.3.4 Pruning Cacti 161
4.3.5 Factorizations of Class (n_1, n_2) 163
4.3.6 Factorizations of Class (n_1, n_2, n_3) 166
4.4 Inequivalent β-Factorizations 172

A Canonical Forms for Inequivalent Factorizations of Class (n_1, n_2, n_3) 177

B Future Work 179

Bibliography 181

List of Notation 187

Index 189

Chapter 1

Introduction

1.1 Preamble

Broadly speaking, this thesis is concerned with counting *factorizations of permutations*. That is, we are interested in finding the number of decompositions of a given permutation π as a product $\pi = \sigma_r \cdots \sigma_1$ of permutation factors σ_i satisfying various conditions. Specifically, we shall focus on *transitive* factorizations, which are defined by the condition that the groups generated by their factors act transitively on the underlying set of symbols.

Transitive factorizations, in general, bear an important relation to geometry through a correspondence between them and certain branched coverings of the sphere. This discrete encoding of branched covers is due to Hurwitz [44], and will be described briefly in §2.3.6. While it is has been the primary reason for much recent interest in transitive factorizations, we emphasize that this geometric connection is peripheral here. We treat factorizations as purely combinatorial structures, and no understanding of the associated geometry is assumed or required of the reader.

Throughout, factorizations will be studied exclusively through their graphical representation as specially labelled maps. The particular correspondence exploited here between factorizations and maps is not altogether new. Rather, the novelty of our approach lies in a detailed investigation of the *descent structure* of these maps. Of particular note is the ability to simplify maps by *pruning trees*. This allows for a shift in focus from transitive factorizations to the combinatorics of *smooth maps*.

We begin, in Chapter 2, with a thorough analysis of transitive factorizations whose factors are all *transpositions*. This is the most widely studied class of transitive factorizations, and structurally the simplest. We have therefore chosen to introduce our methods in this context, despite the fact that they also apply in more general settings. After the basic approach has been established, these generalizations are then surveyed in Chapter 3. Finally, Chapter 4 treats the problem of counting transitive factorizations up to an equivalence relation defined in terms of commutations of adjacent factors. As will be seen there, the methods of Chapters 2 and 3 extend naturally to be applicable to this modified problem.

Supplementary comments and references have been collected at the end of several major sections under the heading *Additional Notes*. Appendix A contains technical material related to §4.2, and suggestions for future work are summarized in Appendix B.

We caution the reader that, in order to minimize redundancy in terminology, conventions are occasionally adopted in the text that are to be understood in a restricted context. These conventions typically have the effect of augmenting previously stated definitions. In particular, the definition of a *map* is modified for the remainder of Chapter 2 by the conventions listed on page 38, while the definition of a *polymap* is altered on page 99 for the remainder of Chapter 3, and again on page 110 for the duration of §3.3.

In a similar vein, we warn that our usage of certain symbols is context sensitive. (For instance, the pervasive symbol w is first met on page 26, and then redefined on pages 114, 132, 158, and 173.) This has been done in a deliberate effort to emphasize the similarities between a variety of different, but strongly related, problems. An index of frequently used notation is provided on page 187. Symbols are listed there in order of their first appearance in a new context.

1.2 Main Results

As mentioned above, this thesis is concerned with the analysis of factorizations through correspondences between them and labelled maps. The general link between factorizations and maps is well-known, but the particular bijections utilized here (Theorems 2.4.11, 2.4.21 and their relatives in later chapters) are significant, as they have not, to our knowledge, previously been exploited in tackling enumerative problems.

The crux of our analysis is a new method, called "tree pruning", that effectively simplifies the maps associated with factorizations. Theorems 2.6.7 and 2.6.10 in Chapter 2 describe the tree pruning bijection and its primary enumerative consequence, namely that a generating series for transitive factorizations into transpositions can be expressed as the composition of a series counting certain smooth maps with the series counting rooted, labelled trees. This algebraic dependence on the tree series has previously been observed by other authors, but tree pruning offers the first combinatorial explanation of its presence. Theorems 2.7.11 and 2.7.14 provide bijections which allow for the straightforward enumeration of smooth maps with two and three faces, respectively. When combined with the tree pruning bijection, these results lead to new bijective derivations (Theorems 2.8.4 and 2.8.7) of two special cases of Hurwitz's formula. Moreover, we give a combinatorial proof of a recursion that shows the generating series for minimal transitive factorizations to be rational when written in terms of the tree series; this is Theorem 2.7.17, which was first established algebraically in [33, 36].

A host of extensions of these ideas to more general factorizations follow in Chapter 3, where the pruning of cacti is paramount (Theorem 3.3.13). Much of the chapter is devoted to new combinatorial proofs of known results, including progress on the *double Hurwitz problem* (Corollaries 3.4.7 and 3.4.9). Some of these results appear in an amplified form. In particular, we draw attention to Corollary 3.3.15, which extends an earlier result [31] concerning the number of minimal transitive factorizations of permutations into k-cycles. See also Corollary 3.4.14, which is related to a bijection of Goulden and Yong [39].

A new graphical model of equivalence classes of factorizations is described in Chapter 4, as is the application of pruning techniques to the enumeration of these classes (Theorem 4.3.9). The model itself quickly leads to a derivation of Springer's formula [65] for the number of inequivalent minimal factorizations of a full cycle into cycles of arbitrary lengths (Theorem 4.3.6). Pruning cacti then allows for a straightforward treatment of inequivalent factorizations of permutations that are a product of two cycles. The main result along these lines is Corollary 4.3.12, which generalizes a counting series for these objects found by Goulden, Jackson, and Latour [32]. This work is extended in Corollary 4.3.14 to give an admittedly unrefined first expression for a generating series for inequivalent factorizations of permutations composed of three cycles. The thesis concludes with Theorem 4.4.2, which represents an initial step towards introducing the notion of equivalence into the double Hurwitz problem.

1.3 Background Material and Notational Conventions

It is assumed that the reader is familiar with the material summarized below. Although we have used standard notation when possible, we caution that certain nonstandard terminology has been adopted for the convenience it provides. A thorough scan of this section is therefore strongly suggested for every reader.

1.3.1 Sets, Compositions and Partitions

If $n \in \mathbb{N}$ then we frequently write $[n]$ for the set $\{1, 2, \ldots, n\}$. As usual, $|S|$ denotes the cardinality of the finite set S.

A **composition** of $n \in \mathbb{N}$ is a tuple $\alpha = (\alpha_1, \ldots, \alpha_m) \in \mathbb{N}^m$ such that $\alpha_1 + \cdots + \alpha_m = n$. The integers α_i are called the **parts** of α and the number of parts in α is known as its **length**. We write $\alpha \models n$ to indicate that α is a composition of n, and $\ell(\alpha)$ denotes the length of α.

A **partition** is a composition having weakly decreasing parts: that is, $(\alpha_1, \ldots, \alpha_m) \models n$ is a partition of n if $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$. We write $\alpha \vdash n$ to signify that α is a partition of n. If the partition α has m_i parts equal to i, then we write $\alpha = [1^{m_1}\, 2^{m_2} \cdots]$, suppressing entries with $m_i = 0$. Any ambiguity between this definition of $[n]$ and the previously mentioned $[n] = \{1, \ldots, n\}$ should be easily resolved from context. We also define $|\operatorname{Aut}(\alpha)| = \prod_i m_i!$, which is the number of automorphisms of α. For example, $\alpha = (4, 4, 2, 2, 2, 1) = [1\, 2^3\, 4^2]$ is a partition of 15 having $\ell(\alpha) = 6$ and $|\operatorname{Aut}(\alpha)| = 12$.

1.3.2 Generating Series and Lagrange Inversion

Let R be a commutative ring with a unit. Recall that $R[S]$ and $R[[S]]$ are, respectively, the rings of **polynomials** and **formal power series** in the set S of algebraically independent and commuting indeterminates with coefficients from R. All generating series appearing in this thesis belong to $R[[S]]$, where the coefficient ring R is invariably a subring of $\mathbb{Q}[[T]]$ for some set T.

If $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{i} = (i_1, \ldots, i_n)$, then we define $\mathbf{x}^{\mathbf{i}}$ to be the monomial $x_1^{i_1} \cdots x_n^{i_n}$. The coefficient of $\mathbf{x}^{\mathbf{i}}$ in the series $f(\mathbf{x}) \in R[[x_1, \ldots, x_n]]$ is denoted by $[\mathbf{x}^{\mathbf{i}}]\, f(\mathbf{x})$ It is also convenient to define $\mathbf{i}! = i_1! i_2! \cdots i_n!$ so that, for instance, $[\mathbf{x}^{\mathbf{i}}/\mathbf{i}!]\, f(\mathbf{x}) = i_1! \cdots i_n!\, [\mathbf{x}^{\mathbf{i}}]\, f(\mathbf{x})$. We write $f(\mathbf{0})$ for the constant term of the series f. The **formal derivative** of the series $f(x) = \sum_n a_n x^n \in R[[x]]$ is defined to be the series $\frac{df}{dx} = \sum_n n a_n x^{n-1}$, and the **formal integral** of f is given by $\int f(x)\, dx = \sum_n \frac{1}{n+1} a_n x^{n+1}$. Note that $f'(x) = g(x)$ and $f(0) = 0$ implies $f(x) = \int g(x)\, dx$.

The next result, known as the **Lagrange implicit function theorem** (or, briefly, **Lagrange inversion**), will be a very important tool in our study.

Theorem 1.3.1 (Lagrange)**.** *Let $\phi \in R[[\lambda]]$ be such that $\phi(0) \neq 0$. Then there exists a unique formal power series $w \in R[[x]]$ such that $w = x\phi(w)$. Moreover, for any $f \in R[[\lambda]]$ and $n > 0$ we have*

$$[x^n]\, f(w) = \frac{1}{n}[\lambda^{n-1}]\, f'(\lambda)\phi^n(\lambda).$$

□

Example 1.3.2. Let $R = \mathbb{Q}[u, q_1, q_2, \ldots]$ and consider the series $\phi(\lambda) = e^{uQ(\lambda)} \in R[[\lambda]]$, where $Q(\lambda) = q_1\lambda + q_2\lambda^2 + q_3\lambda^3 + \cdots$. Then there is a unique series $w \in R[[x]]$ which satisfies the functional equation

$$w = x\phi(w) = xe^{uQ(w)}.$$

Moreover, we can apply Lagrange inversion to determine the coefficient of the generic monomial $q_\beta u^r x^n$ in the composition $Q(w)$. Here we have used the notation $q_\beta = q_{\beta_1} q_{\beta_2} \cdots q_{\beta_m}$, where $\beta = (\beta_1, \beta_2, \ldots, \beta_m)$ is a partition. Lagrange's theorem gives

$$[x^n]\, Q(w) = \frac{1}{n}\, [\lambda^{n-1}]\, Q'(\lambda) e^{nuQ(\lambda)},$$

whence it follows that

$$\begin{aligned}
[q_\beta u^r x^n]\, Q(w) &= \frac{1}{n}\, [q_\beta \lambda^{n-1}]\, Q'(\lambda) \frac{n^r\, Q(\lambda)^r}{r!} \\
&= \frac{n^{r-1}}{r!}\, [q_\beta \lambda^{n-1}] \frac{1}{r+1} \frac{d}{d\lambda} Q(\lambda)^{r+1} \\
&= \frac{n^r}{(r+1)!}\, [q_\beta \lambda^n] \left(q_1\lambda + q_2\lambda^2 + q_3\lambda^3 + \cdots\right)^{r+1} \\
&= \frac{n^r}{(r+1)!}\, [q_\beta \lambda^n] \sum_{k \geq 1} \sum_{\substack{\alpha \models k \\ \ell(\alpha) = r+1}} q_\alpha \lambda^k \\
&= \frac{n^r}{(r+1)!}\, [q_\beta] \sum_{\substack{\alpha \models n \\ \ell(\alpha) = r+1}} q_\alpha .
\end{aligned}$$

If $\beta \vdash n$ and $\ell(\beta) = r + 1$, then the term q_β appears exactly $(r + 1)!/|\operatorname{Aut}(\beta)|$ times in the final summation, giving

$$[q_\beta u^r x^n]\, Q(w) = \frac{n^{\ell(\beta)-1}}{|\operatorname{Aut}(\beta)|}$$

in this case. The coefficient is zero under any other conditions. □

For further information regarding generating series and their combinatorial applications, we direct the reader to any of the standard references on combinatorial enumeration, such as [67], [68], [74] and [26].

1.3.3 Complete Symmetric Functions and Umbral Composition

Fix $k \geq 0$ and $m \geq 1$. Then the **complete symmetric function** of total degree k in the indeterminates $x_1, \ldots, x_m$ is the series $h_k \in \mathbb{Q}[x_1, \ldots, x_m]$ defined as follows:

$$h_k(x_1, \ldots, x_m) = \sum_{\substack{i_1,\ldots,i_m \geq 0 \\ i_1+\cdots+i_m=k}} x_1^{i_1} \cdots x_m^{i_m}.$$

We also introduce the series h_k^+ consisting of all terms of h_k of positive degree. That is,

$$h_k^+(x_1, \ldots, x_m) = \sum_{\substack{i_1,\ldots,i_m \geq 1 \\ i_1+\cdots+i_m=k}} x_1^{i_1} \cdots x_m^{i_m}.$$

Finally, with $\mathbf{x} = (x_1, \ldots, x_m)$, we define the generating series

$$\Delta(t\,;\,\mathbf{x}) = \prod_{i=1}^{m} \frac{1}{1 - tx_i} = \sum_{k\geq 0} h_k(\mathbf{x})t^k \qquad \text{and} \qquad \Delta^+(t\,;\,\mathbf{x}) = \prod_{i=1}^{m} \frac{tx_i}{1 - tx_i} = \sum_{k\geq 0} h_k^+(\mathbf{x})t^k$$

for $h_k(\mathbf{x})$ and $h_k^+(\mathbf{x})$, respectively.

Let $A(t) = \sum_k a_k t^k$ be any formal power series over a commutative ring R (with unit). Then we define the series $A(t) \circ \Delta^+(t\,;\,\mathbf{x}) \in R[[\mathbf{x}]]$ as the following umbral composition of $A(t)$ with the complete symmetric functions,

$$A(t) \circ \Delta^+(t\,;\,\mathbf{x}) = \sum_{k\geq 0} \sum_{\substack{i_1,\ldots,i_m \geq 1 \\ i_1+\cdots+i_m=k}} a_k x_1^{i_1} \cdots x_m^{i_m}. \tag{1.1}$$

The indeterminate t here is obviously a dummy variable. Thus $A(t) \circ \Delta^+(t\,;\,\mathbf{x})$ is obtained from $A(t)$ by replacing t^k with the sum of all monomials $x_1^{i_1} \cdots x_m^{i_m}$ of total degree k and positive degree in each x_i.

Lemma 1.3.3. *Let $A(t) \in R[[t]]$ be any formal power series over the commutative ring R. Then, for $m \geq 2$,*

$$A(t) \circ \Delta^+(t\,;\,x_1, \ldots, x_m) = (-1)^m A(0) + \sum_{i=1}^{m} A(x_i) \prod_{\substack{1\leq j\leq m \\ j\neq i}} \frac{x_j}{x_i - x_j}.$$

□

1.3.4 Cyclic Lists

Two sequences $(a_0, a_1, \ldots, a_n)$ and $(b_0, b_1, \ldots, b_n)$ are *equivalent up to cyclic shift* if there is some integer j such that $b_i = a_{\overline{i+j}}$ for all $0 \le i \le n$, where $\overline{i+j}$ denotes the least nonnegative residue of $i + j$ modulo $n + 1$. We call the equivalence classes under this relation **cyclic lists** (or **circular sequences**), and use the notation $(a_0, \ldots, a_n)^\circ$ to indicate the class containing the sequence $(a_0, \ldots, a_n)$. Thus, for example, $(1, 3, 2, 4, 2)^\circ = (4, 2, 1, 3, 2)^\circ = (3, 2, 4, 2, 1)^\circ$.

Generally speaking, use of the notation $(a_m, a_{m+1}, \ldots, a_n)^\circ$ indicates that the symbol a_k is to be interpreted as $a_{\bar{k}}$, where $\bar{k}$ is the unique residue of k modulo $n - m + 1$ in the range $m \le \bar{k} \le n$. For instance, use of the notation $(a_0, a_1, a_2, a_3)^\circ$ implies $a_{-1} = a_3$ and $a_9 = a_1$.

A cyclic list $L = (a_0, \ldots, a_n)^\circ$ of real numbers is said to be **increasing** if one of its representative sequences is strictly increasing. A similar definition holds for **nondecreasing** cyclic lists. A pair (a_{i-1}, a_i) satisfying $a_{i-1} \ge a_i$ is called a **descent** of L. Thus L is increasing if and only if it has no descents. For example, $(3, 4, 1, 2)^\circ$ is increasing, whereas $(3, 1, 2, 4)^\circ$ is contains two descents, namely $3 \ge 1$ and $4 \ge 3$.

Finally, if S is a finite set of real numbers, then we write S° for the unique increasing circular sequence composed of the elements of S. For example, $S = \{2, 1, 5, 4, 0\}$ gives $S^\circ = (0, 1, 2, 4, 5)^\circ$.

1.3.5 The Symmetric Group

If X is a finite nonempty set then the **symmetric group** $\mathfrak{S}_X$ is the group of permutations on X. For a positive integer n, we write $\mathfrak{S}_n$ in place of $\mathfrak{S}_{[n]}$.

The symbol ι will be used to denote the identity element of $\mathfrak{S}_n$ (the parameter n being understood from context). We multiply permutations from right to left; that is, in a manner consistent with the usual composition of functions:

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 6 & 2 & 1 & 4 & 5 \end{pmatrix} \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 1 & 4 & 2 & 6 & 5 & 3 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 3 & 1 & 6 & 5 & 4 & 2 \end{pmatrix}$$

The **support** of $\pi \in \mathfrak{S}_n$ is the subset $S \subseteq [n]$ of symbols which are not fixed by π, that is, $i \in S$ if and only if $\pi(i) \ne i$. Thus, for example, the identity has empty support. We call π a ***k*-cycle** if its support can be arranged in a cyclic list $(a_1, \ldots, a_k)^\circ$ such that $\pi(a_i) = a_{i+1}$ for all i. We write $(a_1 \cdots a_k)$ for this k-cycle. We usually refer to 2-cycles as **transpositions** and n-cycles in $\mathfrak{S}_n$ as **full cycles**.

Each permutation $\pi \in \mathfrak{S}_n$ acts on $[n]$ in the obvious way, and we let $\operatorname{orb} \pi$ denote the collection of **orbits** under this action. If $\operatorname{orb}\pi = \{\mathcal{O}_1, \ldots, \mathcal{O}_m\}$, then $\mathcal{O}_i = \{\pi^j(a_i) : 0 \le j < k_i\}$ for some $a_i \in [n]$ and a minimal $k_i > 0$. Thus $\pi = \pi_1 \cdots \pi_m$, where π_i is the k_i-cycle $(a_i\ \pi(a_i)\ \cdots\ \pi^{k_i-1}(a_i))$ supported by $\mathcal{O}_i$. We call the π_i the **(disjoint) cycles** of π. The decomposition of π into disjoint cycles is unique. When it causes no confusion, we suppress cycles of length 1 (fixed points) from a permutation written in disjoint cycle form. For example:

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 4 & 1 & 8 & 6 & 2 & 9 & 5 \end{pmatrix} = (1\,3\,4)(2\,7)(5\,8\,9).$$

If $\pi \in \mathfrak{S}_n$ has m_i disjoint i-cycles, then the **cycle type** of π is the partition $[1^{m_1}\,2^{m_2}\,\cdots]$ of n. We write $\ell(\pi)$ for the number of cycles of π. Clearly $\ell(\pi) = \ell(\alpha)$ when π has cycle type α. For example, σ above has cycle type $[1\,2\,3^2]$ and $\ell(\sigma) = 4$, while $\iota \in \mathfrak{S}_n$ has cycle type $[1^n]$ and $\ell(\iota) = n$.

The **conjugacy class** $\{\sigma^{-1}\pi\sigma : \sigma \in \mathfrak{S}_n\}$ of a permutation π contains all those permutations having the same cycle type as π. If $\alpha \vdash n$ then we write $\mathscr{C}_\alpha$ for the conjugacy class in $\mathfrak{S}_n$ consisting of all permutations having cycle type $\alpha \vdash n$. Thus we have

$$|\mathscr{C}_\alpha| = \frac{n!}{|\operatorname{Aut}(\alpha)| \cdot \prod_{j=1}^{m} \alpha_j},$$

For example, k-cycles in $\mathfrak{S}_n$ have cycle type $[1^{n-k}\,k]$, and there are $|\mathscr{C}_{[1^{n-k}\,k]}| = \binom{n}{k}(k-1)!$ of them.

1.3.6 The Group Algebra of the Symmetric Group

Recall that the **group algebra** of $\mathfrak{S}_n$ over $\mathbb{C}$ is the algebra $\mathbb{C}\mathfrak{S}_n$ of all formal linear combinations of permutations on n symbols with scalars in $\mathbb{C}$. It is well known that its centre $Z(\mathbb{C}\mathfrak{S}_n)$ has a basis $\{\mathsf{K}_\alpha : \alpha \vdash n\}$ consisting of the **class sums** $\mathsf{K}_\alpha = \sum_{\sigma \in \mathscr{C}_\alpha} \sigma$. That is, any element of $Z(\mathbb{C}\mathfrak{S}_n)$ can be resolved into a linear combination of class sums. Thus, for $\alpha \vdash n$ and $z \in Z(\mathbb{C}\mathfrak{S}_n)$, we extend the usual coefficient operator notation and write $[\mathsf{K}_\alpha]\, z$ for the coefficient of K_α in the expansion of z into these basis elements. The scalars $c^{\alpha}_{\beta_1,\ldots,\beta_r} = [\mathsf{K}_\alpha]\, \mathsf{K}_{\beta_1} \cdots \mathsf{K}_{\beta_r}$ are known as the **connection coefficients** of $Z(\mathbb{C}\mathfrak{S}_n)$.

There is another important basis $\{\mathsf{F}_\theta : \theta \vdash n\}$ of $Z(\mathbb{C}\mathfrak{S}_n)$, this one consisting of **orthogonal idempotents**. That is, $\mathsf{F}_\theta \mathsf{F}_\rho = \delta_{\theta,\rho} \mathsf{F}_\theta$ for all $\theta, \rho \vdash n$, where $\delta_{\theta,\rho}$ is 1 if $\theta = \rho$ and 0 otherwise.

The idempotents F_θ are related to the class sums K_α by

$$\mathsf{K}_\alpha = |\mathscr{C}_\alpha| \sum_{\theta \vdash n} \frac{\chi_\alpha^\theta}{f^\theta} \mathsf{F}_\theta \qquad \text{and} \qquad \mathsf{F}_\theta = \frac{f^\theta}{n!} \sum_{\alpha \vdash n} \chi_\alpha^\theta \mathsf{K}_\alpha, \tag{1.2}$$

where χ_α^θ is the value of the character of the irreducible representation of $\mathfrak{S}_n$ indexed by $\theta \vdash n$ at any element of the class $\mathscr{C}_\alpha$, and $f^\theta = \chi_{[1^n]}^\theta$ is the degree of this character. Further details can be found in [63].

1.3.7 Graphs

We define **graphs** as usual, with loops and multiple edges allowed. That is, a **graph** is a tuple $\mathscr{G} = (V, E, \phi)$, where V and E are finite disjoint sets and ϕ is a function which assigns, to each $e \in E$, a multiset $\{u, v\}$ with $u, v \in V$. The elements of V and E are called the **vertices** and **edges** of $\mathscr{G}$, respectively. A **subgraph** of $\mathscr{G} = (V, E, \phi)$ is a graph $\mathscr{G}' = (V', E', \phi')$ such that $V' \subseteq V$, $E' \subseteq E$, $\phi' = \phi|_{E'}$, and $\phi'(e') \subseteq V'$ for all $e' \in E'$.

Of course, graphs have their usual representation in $\mathbb{R}^3$ as collections of points (vertices) connected by curves (edges). In particular, if $\phi(e) = \{u, v\}$, then edge e is a curve joining vertices u and v. We frequently abuse terminology and refer to $\{u, v\}$, rather than e itself, as an edge. This generally allows us to suppress mention of the incidence function ϕ entirely.

The vertex v is **incident** with the edge e if $v \in \phi(e)$. We write $\delta(v)$ for the set of all edges incident with v. A **loop** is an edge incident with only one vertex. Two distinct vertices u and v are **adjacent** if they are both incident with a common edge. The **degree** of a vertex v, written $\deg(v)$, is the number of edges incident with v, with loops counted twice.

An **isomorphism** of the graphs $\mathscr{G} = (V, E, \phi)$ and $\mathscr{G}' = (V', E', \phi')$ consists of a pair (f, g) of bijections $f : V \longrightarrow V'$ and $g : E \longrightarrow E'$ such that $v \in V$ is incident with $e \in E$ in $\mathscr{G}$ if and only if $f(v)$ is incident with $g(e)$ in $\mathscr{G}'$. Thus an isomorphism of graphs preserves edge incidence.

A **walk** of length $k + 1$ in $\mathscr{G}$ is a sequence $v_0, e_0, v_1, e_1, \ldots, v_k, e_k, v_{k+1}$ of vertices v_i and edges e_i such that e_i is incident with both v_i and v_{i+1} for all $0 \le i \le k$. A walk of length 0 is a single vertex. If no vertex or edge is duplicated in a walk then it is called a **path**. The walk $v_0, e_0, \ldots, e_k, v_{k+1}$ is said to be **closed** if $v_0 = v_{k+1}$, and in this case we identify it with the circular sequence $((v_0, e_0), \ldots, (v_k, e_k))^\circ$.

A graph is **connected** if there is a walk between any two of its vertices. A **component** of the graph $\mathscr{G}$ is a maximal connected subgraph of $\mathscr{G}$. We write $\mathscr{G} \backslash e$ for the graph obtained by deleting edge e from $\mathscr{G}$, and e is said to be a **bridge** if $\mathscr{G} \backslash e$ has more components than $\mathscr{G}$. In fact, if $\mathscr{G}$ is connected then e is a bridge of $\mathscr{G}$ precisely when $\mathscr{G}\backslash e$ has exactly two components. Furthermore, it can be shown that an edge is a bridge if and only if it does not belong to a cycle. A **tree** is a connected graph without cycles, and thus every edge of a tree is a bridge.

A **vertex-labelling** of the graph $\mathscr{G} = (V, E, \phi)$ is a function $\lambda : V \longrightarrow L$, where L is an arbitrary set. The elements of L are called **vertex labels**, and the pair $(\mathscr{G}, \lambda)$ constitutes a **vertex-labelled graph**. Note that the vertex labels *need not be distinct*. An **isomorphism** of vertex-labelled graphs must preserve labels as well as incidence. That is, the vertex-labelled graphs $(\mathscr{G}, \lambda)$ and $(\mathscr{G}', \lambda')$ are isomorphic if there is an isomorphism (f, g) of the graphs $\mathscr{G}$ and $\mathscr{G}'$ such that $\lambda'(f(v)) = \lambda(v)$ for all $v \in V$. Edges can be assigned labels in a like manner to give **edge-labelled** graphs. Moreover, various labellings may be superimposed upon each other, even if they label the same objects. For instance, (vertex)-rooted, vertex-labelled graphs are obtained by superimposing two vertex-labellings on graphs: the first uses distinct labels $1, \ldots, n$, and the second assigns 0 to all vertices except one, to which it assigns the symbol R, thereby distinguishing it as the root.

We adopt two conventions concerning labelled graphs. First, if v is a vertex of the vertex-labelled graph $(\mathscr{G}, \lambda)$, then we abuse notation and also use the symbol v to represent the label $\lambda(v)$ of v. Our particular meaning will always be clear from context, and if labels are distinct then such usage is unambiguous in any case. Second, we generally suppress all mention of particular labelling schemes, making simple reference to *the vertex-labelled graph* $\mathscr{G}$, for example. Here it is to be understood that the labelling under consideration is a bijection with $[n]$, for some $n \in \mathbb{N}$. Similarly, if the vertices of $\mathscr{G}$ are said to be *labelled with the set* L, then the labelling is supposed to be a bijection with L. Analogous conventions also apply for edge-labelled graphs. Thus, by convention, a *vertex- and edge-labelled graph on n vertices and m edges* has vertices labelled (distinctly) with the integers $1, \ldots, n$, and edges labelled (distinctly) with the integers $1, \ldots, m$.

1.3.8 Maps

A **map** is a 2-dimensional cellular complex whose polyhedron (*i.e.* geometric realization) is homeomorphic to some orientable surface. A **surface**, in this context, is a compact, connected, 2-manifold without boundary. The reader is directed to any text on combinatorial surface topology for further details on cellular decomposition and surfaces. See, for example, [69].

The 0-cells, 1-cells, and 2-cells of a map $\mathscr{M}$ are referred to as its **vertices**, **edges**, and **faces**, respectively. The **genus** of $\mathscr{M}$ is the genus of its polyhedron. If $\mathscr{M}$ has V vertices, E edges, and F faces, then its genus g is determined by the **Euler-Poincaré** formula, $V - E + F = 2 - 2g$. In what follows we make no effort to distinguish between a map and its polyhedron.

In practice, it is convenient to consider maps from a less technical perspective than is indicated by the definition given above. A map can be thought of as an embedding of a connected graph $\mathscr{G}$ in a canonical orientable surface $\mathcal{S}$ of given genus (*i.e.* a sphere with a prescribed number of handles). More precisely, the embedding $\phi : \mathscr{G} \longrightarrow \mathcal{S}$ defines a map if each of the connected components of $\mathcal{S} - \phi(\mathscr{G})$ is homeomorphic to an open disc. Of course, the map so defined has skeleton $\mathscr{G}$, and its faces correspond with the components of $\mathcal{S} - \phi(\mathscr{G})$.

Much of our work will concern maps of genus 0, also known as **planar maps**. Of course, these are maps that arise as embeddings of graphs on the sphere or, equivalently, the plane. When such a map is rendered in the plane, one of its faces is unbounded. We call this the **outer face**.

The vertices and edges of the map $\mathscr{M}$, along with their associated incidence relations, form a connected graph known as the **skeleton** of $\mathscr{M}$. Graph theoretic terminology (*e.g.* walk, loop, bridge) applied to $\mathscr{M}$ invariably refers to its skeleton. We write $\mathscr{M} \backslash e$ for the structure resulting from the deletion of edge e from $\mathscr{M}$. If e is not a bridge of $\mathscr{M}$, then $\mathscr{M} \backslash e$ is itself a map having the same genus as $\mathscr{M}$, but one fewer edges and one fewer faces. If $e = \{a, b\}$ is a bridge, then $\mathscr{M} \backslash e$ naturally separates into two maps $\mathscr{M}_a$ and $\mathscr{M}_b$ whose genera sum to the genus of $\mathscr{M}$.

By definition, orientability guarantees the existence of a consistent **clockwise** sense of rotation everywhere on a map. This, in turn, allows for the unambiguous definition of **right** and **left**. Let F be a face of the map $\mathscr{M}$. Then there is one (and only one) closed walk in $\mathscr{M}$ which traverses precisely those edges incident with F and, in doing so, keeps F to the left of the line of traversal. We call this walk the **boundary walk** of F. Its length is called the **degree** of F as is denoted by $\deg(F)$. If F has boundary walk $W = ((v_0, e_0), \ldots, (v_k, e_k))^\circ$, then any subsequence (e_{i-1}, v_i, e_i) of W consisting of two consecutive edges and their common incident vertex is called a **corner** of F. Plainly, a corner cannot belong to more than one face of a loopless map.

An **isomorphism** of the maps $\mathscr{M}$ and $\mathscr{M}'$ is an orientation-preserving homeomorphism between them which sends i-cells to i-cells and preserves incidence. An **automorphism** of $\mathscr{M}$ is an isomorphism from $\mathscr{M}$ to itself. The condition that an isomorphism be orientation-preserving is a natural one for various reasons. In essence, it asserts that turning a map "inside out" is not a valid symmetry. In direct analogy with the case of graphs, the various cells of a map can be **labelled** with arbitrary sets. In fact, all substructures of a map which are preserved by isomorphism (such as corners) can also be labelled. The notion of isomorphism is amplified for each class of labelled maps to force the preservation of all labels. We adopt the same conventions for labellings of maps as we do for labellings of graphs.

Each edge of a map can be considered to be composed of two **half-edges**, one for each "end" of the edge. Thus a half-edge is uniquely determined by a pair (v, e) consisting of a vertex v and an incident edge e. In particular, the closed walk $((v_0, e_0), \dots, (v_k, e_k))^\circ$ is fully specified by a cyclic list $(h_0, \dots, h_k)^\circ$ of half-edges, where h_i determined by the pair (v_i, e_i). Note that a clockwise tour about any vertex v, via a circle of small radius centred at v, encounters all half-edges incident with v exactly once in some cyclic order. We call the cyclic list of half-edges so produced the **circulator** of v.

Half-edges get sent to half-edges under an isomorphism of maps, and thus they can be labelled. Let $\mathscr{M}$ be a half-edge-labelled map with vertex set V and edge set E. Note that, by convention, this implies the half-edges of $\mathscr{M}$ are labelled with the set $\{1, \dots, 2m\}$, where $m = |E|$. We associate with each edge $e \in E$ the transposition $\tau_e = (h\,h') \in \mathfrak{S}_{2m}$, where h and h' are the half-edges that compose e. With each vertex $v \in V$ we associate the k-cycle $c_v = (h_1\,h_2\,\cdots\,h_k) \in \mathfrak{S}_{2m}$, where $(h_1, \dots, h_k)^\circ$ is the circulator of v. Geometrically, c_v can be interpreted as an instruction to "pivot clockwise" around vertex v from one of its incident half-edges to the next. Similarly, the action of τ_e is interpreted as that of "traversing the edge" e. If we define $\epsilon, \nu \in \mathfrak{S}_{2m}$ by $\epsilon = \prod_{e \in E} \tau_e$ and $\nu = \prod_{v \in V} c_v$, then a cycle of $\epsilon\nu$ is seen to be cyclic list of the half-edges encountered along the boundary walk of a face of $\mathscr{M}$. Hence the cycles of $\epsilon\nu$ completely determine the boundary walks of the faces of $\mathscr{M}$.

To put this more formally, define a **rotation system** on the symbols $\{1, \dots, 2m\}$ to be a pair $\mathscr{R} = (\epsilon, \nu)$ of permutations in $\mathfrak{S}_{2n}$ such that $\epsilon \in \mathscr{C}_{[2^m]}$. (That is, all cycles of ϵ are transpositions.) The rotation system (ϵ, ν) is said to be **transitive** if the permutations ϵ and ν together generate the full symmetric group $\mathfrak{S}_n$. Then we have the following theorem [15, 42, 71], which serves to completely combinatorialize half-edge-labelled maps:

Theorem 1.3.4 (Embedding Theorem)**.** *There is a bijection between transitive rotation systems on the symbols* $\{1, 2, \dots, 2m\}$ *and half-edge-labelled maps on m edges. Moreover, if* $\mathscr{R} = (\epsilon, \nu)$ *is a transitive rotation system, and if* $\mathscr{M}$ *is the half-edge-labelled map corresponding to* $\mathscr{R}$ *under this bijection, then the vertices, edges, and faces of* $\mathscr{M}$ *are in correspondence with the cycles of the permutations* ν, ϵ, *and* $\epsilon\nu$, *respectively.* □

The correspondence referred to in the theorem, between the cells of $\mathscr{M}$ and the cycles of ν, ϵ, and $\epsilon\nu$, is precisely that which is described above. That is: (1) the cycle $(i\ j)$ of ϵ corresponds with an edge whose ends are labelled i and j, (2) the cycle $(i_1, \cdots i_k)$ of ν corresponds with a vertex whose circulator is $(i_1, \dots, i_k)^\circ$, and (3) the cycle $(j_1\ \cdots\ j_m)$ of $\epsilon\nu$ corresponds with a face whose boundary walk is determined by $(j_1, \dots, j_m)^\circ$.

Chapter 2

Factorizations into Transpositions

2.1 Introduction

It is well known that the set of transpositions $\{(i\ j) : 1 \leq i < j \leq n\}$ generates all of $\mathfrak{S}_n$. That is, any permutation $\pi \in \mathfrak{S}_n$ can be expressed as a product of these transpositions. This leads to the following definitions.

Definition 2.1.1. *A* ***factorization*** *of $\pi \in \mathfrak{S}_n$ is a tuple $(\tau_r, \ldots, \tau_1)$ of transpositions $\tau_i \in \mathfrak{S}_n$ such that $\tau_r \cdots \tau_1 = \pi$. The* ***length*** *of this factorization is r and its* ***class*** *is the cycle type of π.*

For example, $((1\,4), (2\,3), (3\,5), (2\,4), (1\,3), (1\,5))$ is a factorization of $(1\,2)(3\,4)(5)$ of length 6 and of class $[1\,2^2]$, since

$$(1\,2)(3\,4)(5) = (1\,4)(2\,3)(3\,5)(2\,4)(1\,3)(1\,5). \tag{2.1}$$

We often circumvent the formality of Definition 2.1.1 and refer to an expression such as (2.1) as a factorization. Later, in Chapter 3, we shall consider factorizations whose factors are of arbitrary cycle type, but throughout this chapter the term *factorization* will always have the meaning described above.

Definition 2.1.2. *The factorization $f = (\tau_r, \ldots, \tau_1)$ in $\mathfrak{S}_n$ is* ***transitive*** *if the group $\langle \tau_1, \ldots, \tau_r \rangle$ generated by its factors acts transitively on $[n]$. That is, f is transitive if for any $a, b \in [n]$ there is a permutation $\sigma \in \langle \tau_1, \ldots, \tau_r \rangle$ such that $\sigma(a) = b$.*

For instance, the factorization (2.1) is transitive, whereas $(1\,2)(3\,4)(5) = (3\,5)(3\,4)(1\,2)(4\,5)$ is not. Transitive factorizations are a natural and very important class of factorizations to consider.

More will be said on this shortly; for now, suffice it to say that transitive factorizations play a similar rôle in the study of factorizations as do connected graphs in the study of graphs. In fact, we shall soon see that this is far from being a loose analogy.

The primary focus of this chapter is the enumeration of transitive factorizations through graphical constructions. To this end, we begin with a few comments concerning the permissible lengths of factorizations under various conditions. We shall then be in a position to discuss some of the known results concerning the enumeration of factorizations. Finally, we introduce a variety of graphical representations of factorizations and devote the balance of the chapter to their enumerative applications.

2.2 The Length of a Factorization

The aim of this section is to determine the number of transpositions required to factor a given permutation. The problem is straightforward if no conditions are placed on the factors, but if we restrict our attention to transitive factorizations then more thought is required.

2.2.1 Cuts and Joins

Complete information about the possible lengths of factorizations follows from the following lemma. It describes the effect that multiplication by a transposition has on the number of cycles in a permutation.

Lemma 2.2.1. *Let $\pi \in \mathfrak{S}_n$ and let $(a\,b) \in \mathfrak{S}_n$ be any transposition. Then*

$$\ell((a\,b)\pi) = \begin{cases} \ell(\pi) + 1 & \text{if } a \text{ and } b \text{ are on the same cycle of } \pi\,, \\ \ell(\pi) - 1 & \text{if } a \text{ and } b \text{ are on different cycles of } \pi\,. \end{cases}$$

Proof. Suppose a and b are on the same cycle of π, so that it has the form $(a \cdots a'\, b \cdots b')$. Then $\ell((a\,b)\pi) = \ell(\pi) + 1$ follows since $(a\,b)(a \cdots a'\, b \cdots b') = (a \cdots a')(b \cdots b')$. Similarly if a and b appear on distinct cycles $(a \cdots a')$ and $(b \cdots b')$ of π, then $\ell((a\,b)\pi) = \ell(\pi) - 1$ since $(a\,b)(a \cdots a')(b \cdots b') = (a \cdots a'\, b \cdots b')$. □

The proof of the lemma is to observe that multiplying a permutation on the left by a transposition either cuts one of the permutation's cycles in two, or joins two of its cycles into one. The following terminology reflects this description.

Definition 2.2.2. *The transposition* $(a\,b)$ *is called a* ***cut*** *for* π *if* $\ell((a\,b)\pi) = \ell(\pi) + 1$*, and it is a called a* ***join*** *for* π *if* $\ell((a\,b)\pi) = \ell(\pi) - 1$*.*

This definition extends to factorizations, as follows.

Definition 2.2.3. *Let* $f = (\tau_1, \dots, \tau_r)$ *be a factorization. Then the factor* τ_i *is a* ***cut*** *of* f *if* τ_i *is a cut for the initial product* $\tau_{i-1} \cdots \tau_1$*. Similarly,* τ_i *is a* ***join*** *of* f *if it is a join for* $\tau_{i-1} \cdots \tau_1$*.*

Example 2.2.4. The factor $(2\,4)$ is a cut of the factorization

$$(1\,2\,3)(4\,5) = (4\,5)(2\,4)(2\,3)(1\,3)(1\,4),$$

and each of the remaining factors is a join. □

The following fundamental lemma relates numbers of cuts and joins in a factorization to the number of cycles in its target permutation.

Lemma 2.2.5. *Let* $\pi \in \mathfrak{S}_n$ *be a factorization with* C *cuts and* J *joins. Then* $\ell(\pi) = n + C - J$*.*

Proof. We use induction on the length r of the factorization. If $r = 1$ then clearly $\ell(\pi) = n - 1$, $C = 0$ and $J = 1$, as desired. Suppose the result holds for $r = k$ and let $\pi = \tau_{k+1}\tau_k \cdots \tau_1$ be a factorization having C cuts and J joins amongst its factors. First let us assume τ_{k+1} is a cut in this factorization. Then τ_{k+1} is a cut of $\sigma = \tau_k \cdots \tau_1$, giving $\ell(\pi) = \ell(\tau_{k+1}\sigma) = \ell(\sigma) + 1$. But the factorization $\sigma = \tau_k \cdots \tau_1$ has $C - 1$ cuts and J joins, implying $\ell(\sigma) = n + (C - 1) - J$ by hypothesis. Thus $\ell(\pi) = (n + (C - 1) - J) + 1 = n + C - J$. A similar argument applies when τ_{k+1} is a join, and the result follows by induction. □

2.2.2 Minimal Factorizations

Let f be a factorization of the permutation π. It is well known that if f is of even (respectively, odd) length then *all* factorizations of π are of even (odd) length. The next result establishes this elementary fact and also provides a lower bound on the length of f when no conditions are placed on its factors.

Proposition 2.2.6. *If* $\pi \in \mathfrak{S}_n$ *admits a factorization into* r *transpositions, then* $r \geq n - \ell(\pi)$ *and* $r \equiv n - \ell(\pi) \pmod 2$*. In particular, either all factorizations of* π *are of even length, or all factorizations of* π *are of odd length.*

Proof. Suppose we have a factorization of π with r factors, C of which are cuts and J of which are joins. Then Lemma 2.2.5 gives $\ell(\pi) = n + C - J$, so that $n - \ell(\pi) = J - C$. But clearly $r = J + C$, so we have $n - \ell(\pi) \leq r$ and $n - \ell(\pi) \equiv J - C \equiv r \pmod 2$ as required. □

Note that the k-cycle $(i_1\, i_2\, \cdots\, i_k)$ admits the following factorization into $k-1$ transpositions:

$$(i_1\, i_2\, \cdots\, i_k) = (i_1\, i_2)(i_2\, i_3)\cdots(i_{k-2}\, i_{k-1})(i_{k-1}\, i_k).$$

It follows immediately that any permutation having cycle type $\alpha = (\alpha_1, \alpha_2, \ldots) \vdash n$ admits a factorization into $\sum_i (\alpha_i - 1) = n - \ell(\alpha)$ transpositions. Thus the lower bound given by Proposition 2.2.6 for the length of a factorization is attainable.

Definition 2.2.7. *A factorization of $\pi \in \mathfrak{S}_n$ into exactly $n - \ell(\pi)$ factors is said to be **minimal**. The number $n - \ell(\pi)$ is called the **rank** of π.*

2.2.3 Components

Proposition 2.2.6 identifies all possible lengths of a factorization of $\pi \in \mathfrak{S}_n$ in the case that no restrictions are placed on the factors. We now investigate a lower bound for the length of a factorization of π whose factors are restricted by a generalization of the transitivity condition.

Any subgroup S of $\mathfrak{S}_n$ acts on the set $[n]$ in a natural way. That is, if $\sigma \in S$ and $i \in [n]$ then σ acts on i to give $\sigma \cdot i = \sigma(i)$. This action partitions $[n]$ into disjoint **orbits**. We write $\mathcal{O}_S^i = \{\sigma(i) : \sigma \in S\}$ for the unique orbit containing $i \in [n]$, and orb S for the set of orbits under the action of S. Note that either $\mathcal{O}_S^i = \mathcal{O}_S^j$ or $\mathcal{O}_S^i \cap \mathcal{O}_S^j = \emptyset$ for all $i, j \in [n]$.

Let $f = (\tau_r, \ldots, \tau_1)$ be a factorization of $\pi \in \mathfrak{S}_n$. Then $S = \langle \tau_r, \ldots, \tau_1 \rangle$ acts on $[n]$ as just described. Recall that f is transitive if S acts transitively on $[n]$. More generally, let $C_1, \ldots, C_c$ be the orbits of this action and, for $i = 1, \ldots, c$, let $\pi_i = \pi|_{C_i}$ be the restriction of π to the set C_i. Then, by selecting those factors of f which act nontrivially on C_i, we naturally obtain a transitive factorization f_i of π_i. For example, for the factorization

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

we have

$$\begin{array}{lll} C_1 = \{1,2,3\} & C_2 = \{4,5,6\} & C_3 = \{7,8\} \\ \pi_1 = (1\,2\,3) & \pi_2 = (4)(5\,6) & \pi_3 = (7)(8) \\ f_1 = ((1\,2),(2\,3)) & f_2 = ((4\,5),(5\,6),(4\,6)) & f_3 = ((7\,8),(7\,8)). \end{array}$$

The transitive factorizations $f_1, \ldots, f_c$ are called the **components** of f. Clearly the transposition factors of f_i and f_j commute for $i \neq j$. Hence every c-component factorization f is a shuffling of the factors of c transitive factorizations. This identifies the transitive factorizations as the basic "connected" blocks out of which all factorizations are built. Of course, f is transitive precisely when it has exactly one component.

We would like to find a lower bound on the length of a factorization of π having c components. To do so we require the following technical lemma:

Lemma 2.2.8. *Let S be a subgroup of $\mathfrak{S}_n$. Let $(a\,b) \in \mathfrak{S}_n$ be any transposition, and let T be the subgroup generated by S and $(a\,b)$. Then*

$$|\operatorname{orb} T| = \begin{cases} |\operatorname{orb} S| & \text{if } \mathcal{O}_S^a = \mathcal{O}_S^b \\ |\operatorname{orb} S| - 1 & \text{otherwise.} \end{cases}$$

Proof. Fix any $i \notin \mathcal{O}_S^a \cup \mathcal{O}_S^b$. If $j \in \mathcal{O}_S^i$ then $j \neq a, b$, so that $(a\,b) \cdot j = j$ and hence $\pi \cdot j \in \mathcal{O}_S^i$ for all $\pi \in T$. In particular, $\pi \cdot i \in \mathcal{O}_S^i$ for all $\pi \in T$, so that $\mathcal{O}_T^i \subset \mathcal{O}_S^i$. As S is a subgroup of T we must also have $\mathcal{O}_S^i \subset \mathcal{O}_T^i$, and therefore $\mathcal{O}_T^i = \mathcal{O}_S^i$.

Now consider the case $i \in \mathcal{O}_S^a \cup \mathcal{O}_S^a$. Without loss of generality assume $i \in \mathcal{O}_S^a$. If $j \in \mathcal{O}_S^b$ then clearly $j \in \mathcal{O}_S^i \subset \mathcal{O}_T^i$. If $j \in \mathcal{O}_S^b$ then there exist $\pi, \sigma \in S$ such that $\pi \cdot i = a$ and $\sigma \cdot j = b$, implying $\sigma^{-1}(a\,b)\pi \cdot i = j$ and hence again $j \in \mathcal{O}_T^i$. Thus $\mathcal{O}_S^a \cup \mathcal{O}_S^b \subset \mathcal{O}_{S'}^i$. But if $j \in \mathcal{O}_S^a \cup \mathcal{O}_S^b$ then we also have $(a\,b) \cdot j \in \mathcal{O}_S^a \cup \mathcal{O}_S^b$, and hence we find that $\pi \cdot i \in \mathcal{O}_S^a \cup \mathcal{O}_S^b$ for all $\pi \in T$. It follows that $\mathcal{O}_T^i \subset \mathcal{O}_S^a \cup \mathcal{O}_S^b$.

Altogether we have $\mathcal{O}_T^i = \mathcal{O}_S^a \cup \mathcal{O}_S^b$ if $i \in \mathcal{O}_S^a \cup \mathcal{O}_S^b$ and $\mathcal{O}_T^i = \mathcal{O}_S^i$ otherwise. Thus in the case $\mathcal{O}_S^a = \mathcal{O}_S^b$ we have $\operatorname{orb} T = \operatorname{orb} S$, whereas otherwise we have $\operatorname{orb} T = (\operatorname{orb} S) \cup \{\mathcal{O}_S^a \cup \mathcal{O}_S^b\} - \{\mathcal{O}_S^a, \mathcal{O}_S^b\}$. The result follows immediately. □

Proposition 2.2.9. *If $\pi \in \mathfrak{S}_n$ admits a c-component factorization into r transpositions then*

$$r \geq n + \ell(\pi) - 2c.$$

Proof. Proceed by induction on r. If $r = 1$ then the factorization is $\pi = \tau_1$, where τ_1 is a transposition. Thus $\ell(\pi) = n - 1$, the factorization has $n - 1$ components, and the result holds in this case. Now suppose it holds for all factorizations with $r = k$ factors, and choose any factorization $\pi = \tau_{k+1}\tau_k \cdots \tau_1$. Set $\sigma = \tau_k \cdots \tau_1$ and $S = \langle \tau_1, \ldots, \tau_k \rangle$. Then the induction hypothesis gives

$$k \geq n + \ell(\sigma) - 2\,|\operatorname{orb} S|. \tag{2.2}$$

We consider two possibilities for the transposition τ_{k+1}.

First suppose $\tau_{k+1} = (a\,b)$ is a cut of σ. Then $\ell(\pi) = \ell(\tau_{k+1}\sigma) = \ell(\sigma) + 1$. We also have $\mathcal{O}_S(a) = \mathcal{O}_S(b)$ since a and b are on the same cycle of $\sigma = \tau_k \cdots \tau_1$. Thus Lemma 2.2.8 gives $|\operatorname{orb} S| = |\operatorname{orb} T|$, where T is the group generated by S and $(a\,b)$. From (2.2) we therefore have

$$k + 1 \geq n + \ell(\pi) - 2\,|\operatorname{orb} T|. \tag{2.3}$$

Next suppose $\tau_{k+1} = (a\,b)$ is a join of σ. Then $\ell(\pi) = \ell(\tau_{k+1}\sigma) = \ell(\sigma) - 1$ and Lemma 2.2.8 gives $|\operatorname{orb} T| \geq |\operatorname{orb} S| - 1$. With (2.2) this again yields (2.3). The result follows by induction since $T = \langle \tau_1, \ldots, \tau_{k+1} \rangle$. □

Corollary 2.2.10. *Let π be a permutation having c disjoint cycles $\pi_1, \ldots, \pi_k$. If f is a minimal factorization of π, then its components are $f_1, \ldots, f_k$, where f_i is a minimal factorization of π_i.*

Proof. Let $f = (\tau_r, \ldots, \tau_1)$ be a minimal factorization of $\pi \in \mathfrak{S}_n$, and suppose $f_1, \ldots, f_c$ are the components of f. Then f_i is a transitive factorization of length r_i of some permutation σ_i acting on a subset $S_i \subset [n]$. Clearly $S_1, \ldots, S_c$ are disjoint sets, and $\pi = \sigma_1 \cdots \sigma_c$. Also note that $r_1 + \cdots + r_c = r$ and $|S_1| + \cdots + |S_k| = n$. By Proposition 2.2.9 we have

$$r_i \geq |S_i| + \ell(\sigma_i) - 2 \tag{2.4}$$

for each i. Summing over i gives $r \geq n + \sum_i \ell(\sigma_i) - 2c$, from which the minimality condition $r = n - \ell(\pi)$ gives

$$2c \geq \ell(\pi) + \sum_{i=1}^{c} \ell(\sigma_i). \tag{2.5}$$

This implies $c \geq \ell(\pi)$, as $\ell(\sigma_i) \geq 1$ for all i. But $\pi = \tau_r \cdots \tau_1$ forces $c = |\operatorname{orb}\langle \tau_1, \ldots, \tau_r \rangle| \leq \ell(\pi)$. Thus $c = \ell(\pi) = k$, and (2.5) now yields $\ell(\sigma_i) = 1$ for $i = 1, \ldots, c$. That is, σ_i is a full cycle of $\mathfrak{S}_{n_i}$. Since $\pi = \sigma_1 \cdots \sigma_k$, it follows that the permutations σ_i coincide with the disjoint cycles of

π. Finally, $\ell(\sigma_i) = 1$ forces (2.5) to be tight, which in turn forces (2.4) to be tight for all i. Thus $r_i = |S_i| - 1$, so that f_i is a minimal factorization of σ_i. □

2.2.4 Transitive Factorizations

As mentioned previously, a factorization is transitive if and only if it has one component. It is readily demonstrated that the bound of Proposition 2.2.9 is always attainable in the transitive case. (This is true for any number of components. The demonstration is similar but lengthier.) For any $\alpha = (\alpha_1, \ldots, \alpha_k) \vdash n$ we can express the generic permutation

$$\pi = (1^1 \cdots \alpha_1^1)(1^2 \cdots \alpha_2^2) \cdots (1^k \cdots \alpha_k^k)$$

of cycle type α as the product

$$\pi = (1^{k-1} 1^k) \cdots (1^2 1^3)(1^1 1^2)(1^1 \cdots \alpha_1^1 1^2 \cdots \alpha_2^2 1^3 \cdots \alpha_{k-1}^{k-1} 1^k \cdots \alpha_k^k).$$

The n-cycle on the far right of this product can be further factored (both minimally and transitively) into $n-1$ transpositions, giving a factorization of π into $(n-1)+(k-1) = n + \ell(\pi) - 2$ transpositions. As these transpositions clearly act transitively on $[n]$, the bound of Proposition 2.2.9 has been attained. Accordingly, we make the following definitions.

Definition 2.2.11. *A transitive factorization of $\pi \in \mathfrak{S}_n$ having exactly $n + \ell(\pi) - 2$ factors is said to be* ***minimal transitive****. The number $n + \ell(\pi) - 2$ itself is known as the* ***transitive rank*** *of π.*

Intuitively we expect that, of the $n + \ell(\pi) - 2$ factors in a minimal transitive factorization of $\pi \in \mathfrak{S}_n$, there must be exactly $n-1$ joins (for transitivity) and $\ell(\pi) - 1$ cuts (to obtain $\ell(\pi)$ cycles in the product). This intuition is proved correct by the following corollary of Proposition 2.2.9.

Corollary 2.2.12. *A minimal transitive factorization of $\pi \in \mathfrak{S}_n$ has $n-1$ joins and $\ell(\pi) - 1$ cuts.*

Proof. Suppose such a factorization has C cuts and J joins. Then, since it must have exactly $n + \ell(\pi) - 2$ factors, we have $C + J = n + \ell(\pi) - 2$. But Lemma 2.2.5 gives $\ell(\pi) = n + C - J$. Solving this system gives $C = \ell(\pi) - 1$ and $J = n - 1$, as desired. □

Let $f = (\tau_r, \ldots, \tau_1)$ be any transitive factorization of $\pi \in \mathfrak{S}_n$, not necessarily minimal. Then certainly $r \geq n + \ell(\pi) - 2$. Moreover, the parity restriction of Proposition 2.2.6 guarantees that r exceeds $n + \ell(\pi) - 2$ by an even integer. This leads to the following definition.

Definition 2.2.13. *Let f be a transitive factorization of $\pi \in \mathfrak{S}_n$ of length r. The* ***genus*** *of f is the nonnegative integer g defined by $r = (n + \ell(\pi) - 2) + 2g$. We write $r_g(\alpha)$ for the number $n + \ell(\alpha) + 2g - 2$ of factors in any genus g factorization of class $\alpha \vdash n$.*

The reason for this peculiar choice of terminology will be made apparent later, in §2.4. We emphasize that, by definition, a genus g factorization is transitive. Thus the phrases *"genus 0 factorization"* and *"**minimal transitive** factorization"* are synonymous.

2.2.5 Additional Notes

With the possible exception of Proposition 2.2.9, the material of this section is folklore. In [29], it is shown that a transitive factorization f of $\pi \in \mathfrak{S}_n$ has at least $n + \ell(\pi) - 2$ factors by considering spanning trees of the graph of f. This is Proposition 2.2.9 in the case $c = 1$. The approach followed here is suggested in [70].

2.3 Enumeration of Factorizations and Hurwitz's Problem

The study of factorizations has quite a long history, dating back at least to the late 19th century and the work of Hurwitz, so a good deal is known about their structure and how to count them. In this section we review some techniques which have been successfully applied to analyzing these objects.

We begin by looking at a very general algebraic technique for counting factorizations, based on computations in the group algebra of the symmetric group. In principle, this method is applicable to the enumeration of factorizations of any prescribed length, but, in practice, it can be applied only in the simplest circumstances. Next we turn our attention to transitive factorizations, our principal objects of study. We present an elegant formula of Hurwitz for the number of minimal transitive factorizations of a permutation of arbitrary cycle type, and summarize one of its proofs. The method of proof we discuss is based on a simple combinatorial decomposition, but is heavily supported by a purely algebraic argument — one which seems to belie the simplicity of the formula that it verifies. This having been said, the same method has recently been extended to factorizations of higher genus with considerable success, and no alternative path to these new results is currently known. We conclude the section by briefly commenting on a link between transitive factorizations and geometry.

2.3.1 Factorizations of a Prescribed Length

Let $F_r(\alpha)$ denote the number of factorizations (not necessarily transitive) of any permutation $\pi \in \mathscr{C}_\alpha$ into exactly r transpositions. We shall now quickly derive the generating series

$$\Upsilon(z, \mathbf{p}, u) = \sum_{n,r\geq 1} \sum_{\alpha\vdash n} |\mathscr{C}_\alpha| F_r(\alpha) \frac{z^n}{n!} \frac{u^r}{r!} p_\alpha$$

of these numbers, where $\mathbf{p} = (p_1, p_2, \ldots)$ is a vector of indeterminates and $p_\alpha = p_{\alpha_1} p_{\alpha_2} \cdots$ for $\alpha = (\alpha_1, \alpha_2, \ldots)$.

If $\pi \in \mathscr{C}_\alpha$ then observe that the connection coefficient $[\mathsf{K}_\alpha]\,\mathsf{K}_{\beta_1} \cdots \mathsf{K}_{\beta_r}$ of $\mathbb{C}\mathfrak{S}_n$ is equal to the number of r-tuples $(\sigma_1, \ldots, \sigma_r)$ of permutations with $\sigma_i \in \mathscr{C}_{\beta_i}$ that satisfy $\sigma_1 \cdots \sigma_r = \pi$. (See §1.3.6.) In particular, we have $F_r(\alpha) = [\mathsf{K}_\alpha]\,(\mathsf{K}_{[1^{n-2}\,2]})^r$. One can exploit the relations (1.2) between the orthogonal idempotents of $\mathbb{C}\mathfrak{S}_n$ and the class sums to express an arbitrary connection coefficient as a character sum:

$$[\mathsf{K}_\alpha]\,\mathsf{K}_{\beta_1} \cdots \mathsf{K}_{\beta_r} = \frac{1}{n!} |\mathscr{C}_{\beta_1}| \cdots |\mathscr{C}_{\beta_r}| \sum_{\theta\vdash n} \left(\frac{1}{f^\theta}\right)^{r-1} \chi^\theta_{\beta_1} \cdots \chi^\theta_{\beta_r} \chi^\theta_\alpha. \tag{2.6}$$

Setting $\beta_i = [1^{n-2}\,2]$ for all i in (2.6) then yields

$$F_r(\alpha) = \frac{1}{n!} \sum_{\theta\vdash n} f^\theta (\xi_\theta)^r \chi^\theta_\alpha, \qquad \text{where} \quad \xi_\theta = \binom{n}{2} \frac{1}{f^\theta} \chi^\theta_{[1^{n-2}\,2]}. \tag{2.7}$$

Recall that the **Schur symmetric functions** $\{s_\theta : \theta \vdash n\}$ and the **power sum symmetric functions** $\{p_\alpha : \alpha \vdash n\}$ are related [51], [61] through

$$p_\alpha = \sum_{\theta\vdash n} \chi^\theta_\alpha s_\theta \qquad \text{and} \qquad s_\theta = \frac{1}{n!} \sum_{\alpha\vdash n} |\mathscr{C}_\alpha|\, \chi^\theta_\alpha\, p_\alpha. \tag{2.8}$$

We can therefore express χ^θ_α as a scaling of the coefficient of p_α in the resolution of s_θ into power sums. By doing so, (2.7) leads to the expression

$$\Upsilon(z, \mathbf{p}, u) = \sum_{n\geq 1} \frac{z^n}{n!} \sum_{\theta\vdash n} f^\theta s_\theta e^{u\xi_\theta}. \tag{2.9}$$

Equations (2.7) and (2.9) can, in principle, be used to determine $F_r(\alpha)$. However, while ξ_θ is easy to evaluate (see [51], p.118), the evaluation of the arbitrary character appearing in (2.7) or, equivalently, the extraction of the required coefficient from (2.9), is generally intractable. One particularly nice exception is the special case when $\alpha = (n)$, corresponding to factorizations of a full cycle. In this case the Murnaghan-Nakayama rule implies that χ^{θ}_{α} vanishes unless θ is a *hook* of the form $[1^{n-k}\, k]$. Straightforward computation using (2.7) then leads to the following result, which first appeared in [47].

Theorem 2.3.1. *There are*

$$\frac{1}{n!}\binom{n}{2}^r \sum_{k=0}^{n-1} (-1)^k \binom{n-1}{k} (n-2k-1)^r$$

factorizations of any full cycle of $\mathfrak{S}_n$ into r transpositions. □

2.3.2 Transitive Factorizations and Hurwitz Numbers

It happens that the number of topologically inequivalent, almost simple, n-fold coverings of the sphere by a Riemann surface of genus g is directly related to the number of genus g factorizations in $\mathfrak{S}_n$. (A brief description of this connection can be found in §2.3.6.) It was in this context that the study of minimal transitive factorizations began, in the late 19th century, with Hurwitz's investigation of branched coverings of the sphere by the sphere.

Definition 2.3.2. *We write $H_g(\alpha)$ for the number of genus g factorizations of any fixed permutation $\pi \in \mathscr{C}_\alpha$. The numbers $H_g(\alpha)$ are known as **Hurwitz numbers**. The generating series*

$$\Phi^{(g)}(z, \mathbf{p}, u) = \sum_{n\geq 1} \sum_{\alpha\vdash n} |\mathscr{C}_\alpha| H_g(\alpha) \frac{z^n}{n!} \frac{u^{r_g(\alpha)}}{r_g(\alpha)!} p_\alpha, \tag{2.10}$$

*where $\mathbf{p} = (p_1, p_2, \ldots)$ and $p_\alpha = p_{\alpha_1} p_{\alpha_2} \cdots$, will be called the **Hurwitz series**.*

In the literature, various scalings of the numbers $H_g(\alpha)$ are also referred to as Hurwitz numbers. The determination of these numbers is commonly referred to as the *Hurwitz Enumeration Problem*.

Note that the scaling factor $|\mathscr{C}_\alpha|$ appearing in the series (2.10) is a natural one. Since there are $H_g(\alpha)$ genus g factorizations of each of the $|\mathscr{C}_\alpha|$ permutations in the conjugacy class $\mathscr{C}_\alpha$, there are $|\mathscr{C}_\alpha| H_g(\alpha)$ genus g factorizations of class α in total.

Recall that (2.7) counts factorizations only by their length and class, without regard to the number of components. However, since a multi-component factorization is a shuffling of the factors of a collection of transitive ones, a standard exponential generating series argument for connected structures yields the following relationship between the classes of all factorizations and their transitive atoms:

$$1 + \Upsilon(z, \mathbf{p}, u) = \exp\left(\sum_{g\geq 0} \Phi^{(g)}(z, \mathbf{p}, u)\right). \tag{2.11}$$

With (2.7) or (2.9) this connection can, in principle, be used to determine the Hurwitz series, but one is still confronted with all the computational pitfalls of the non-transitive case, and the logarithm now involved compounds these troubles further. Furthermore, this expression offers no combinatorial insight into the nature of transitivity, nor is it amenable to simplification even in cases where simple formulas are known to exist, such as in genus 0.

Hurwitz [44] discovered the following remarkably simple formula for the number $H_0(\alpha)$ of minimal transitive (*i.e.* genus 0) factorizations of any permutation of cycle type α.

Theorem 2.3.3 (The Hurwitz Formula)**.** *For* $\alpha = (\alpha_1, \ldots, \alpha_m) \vdash n$ *we have*

$$H_0(\alpha) = n^{m-3}(n + m - 2)! \prod_{i=1}^{m} \frac{\alpha_i^{\alpha_i}}{(\alpha_i - 1)!}. \tag{2.12}$$

□

Hurwitz did not actually provide a complete proof of Theorem 2.3.3. This did not come until a century later, when Goulden and Jackson [29] rediscovered and fully proved the formula. At least three other proofs are now known, of analytic, geometric and combinatorial flavours. See the Additional Notes at the end of this section further information.

Although obtained independently of Hurwitz's work, the proof of Theorem 2.3.3 offered by Goulden and Jackson begins with essentially the same combinatorial argument that Hurwitz had followed. Through a *cut and join* analysis (details will follow in §2.3.3) they develop a recurrence relation for the numbers $H_0(\alpha)$ in the form of a differential equation satisfied by $\Phi^{(g)}(z, \mathbf{p}, u)$. By applying a change of variables and following an algebraic argument centred around Lagrange inversion, they then demonstrate that the numbers generated by (2.12) satisfy this same recurrence. Finally, Theorem 2.3.3 is established by the uniqueness of solutions with given initial conditions.

Much more is known in the way of explicit formulae for Hurwitz numbers, but none of the higher genus analogues of Theorem 2.3.3 shares its simple multiplicative form. Some further details can be found in §2.3.7. Here we mention only the following evaluation of the special Hurwitz number $H_g((n))$, which counts genus g factorizations of a fixed full cycle of $\mathfrak{S}_n$. The result was first published in this form in [64].

Theorem 2.3.4. *For $n \geq 1$ and any $g \geq 0$ we have*

$$H_g((n)) = \frac{(n-1+2g)!}{2^{2g} n!} n^{n-1+2g} \, [x^{2g}] \left(\frac{\sinh x}{x} \right)^{n-1}.$$

Proof. This comes routinely from Theorem 2.3.1 upon setting $r = r_g((n)) = n - 1 + 2g$. □

2.3.3 Lagrangian Structure in the Hurwitz Series

A recurrence relation for $H_0(\alpha)$ is obtained by noting that the final factor τ_r in a minimal transitive factorization $f = (\tau_r, \ldots, \tau_1)$ must either be a cut of f, and therefore cut an $(i + j)$-cycle of $\tau_{r-1} \cdots \tau_1$ into an i-cycle and a j-cycle, or be a join of f, and do the reverse. In the former case, $(\tau_{r-1}, \ldots, \tau_r)$ is a transitive factorization, while in the latter it has exactly two components. Thus deleting the final factor τ_r, and considering these two cases separately, leads to the following differential equation for the Hurwitz series $\Phi = \Phi^{(0)}(z, \mathbf{p}, u)$:

$$\frac{\partial \Phi}{\partial u} = \frac{1}{2} \sum_{i,j \geq 1} \Big(\overbrace{ij p_{i+j} \frac{\partial \Phi}{\partial p_i} \frac{\partial \Phi}{\partial p_j}}^{\text{join}} + \overbrace{(i+j) p_i p_j \frac{\partial \Phi}{\partial p_{i+j}}}^{\text{cut}} \Big). \tag{2.13}$$

A good amount of technical work is required to show that if $H_0(\alpha)$ is given by (2.12) then $\Phi^{(0)}$ does indeed satisfy this differential equation. The verification in [29] involves complicated summations and essential use of Lagrange's implicit function theorem. One particular implicitly defined series is of central importance. For its definition, first set

$$\phi_0(z, \mathbf{p}) = \sum_{n \geq 1} n^n p_n \frac{z^n}{n!}.$$

Now let $s = s(z, \mathbf{p})$ be the unique formal power series solution of the functional equation

$$s = z \exp \phi_0(s, \mathbf{p}). \tag{2.14}$$

This series arises in the verification of (2.13) as a result of the many connections between it and $\Phi^{(0)}$. We mention only one such relationship presently, namely

$$\left(z \frac{\partial}{\partial z} \right)^2 \Phi^{(0)}(z, \mathbf{p}, 1) = \phi_0(s, \mathbf{p}), \tag{2.15}$$

which can be verified using Lagrange inversion and (2.12). Note that the indeterminate u has been suppressed. This does not result in a loss of information, as u can be recovered through $\Phi^{(g)}(z, \mathbf{p}, u) = u^{2g-2}\Phi^{(g)}(uz, u\mathbf{p}, 1)$. We shall comment on possible combinatorial interpretations of (2.15) later, in §2.4.2.

Recently, these results have been generalized to factorizations of arbitrary genus. The following theorem of Goulden, Jackson and Vakil [34] demonstrates that the strong link between factorizations and the series $s(z, \mathbf{p})$ persists for all genera.

Theorem 2.3.5. *Let* $\Phi^{(g)}(z, \mathbf{p}) = \Phi^{(g)}(z, \mathbf{p}, 1)$. *For each* $i \geq 0$ *set*

$$\phi_i(z, \mathbf{p}) = \sum_{n \geq 1} n^{n+i} p_n \frac{z^n}{n!},$$

and let $s = s(z, \mathbf{p})$ *be defined as in* (2.14). *Then*

$$\left(z\frac{\partial}{\partial z}\right)^2 \Phi^{(0)}(z, \mathbf{p}) = \phi_0(s, \mathbf{p}),$$
$$\Phi^{(1)}(z, \mathbf{p}) = \frac{1}{24}\left(\log(1 - \phi_1(s, \mathbf{p}))^{-1} - \phi_0(s, \mathbf{p})\right),$$

while, for arbitrary genus $g \geq 2$,

$$\Phi^{(g)}(z, \mathbf{p}) = \sum_{e=2g-1}^{5g-5} \frac{1}{(1 - \phi_1(s, \mathbf{p}))^e} \cdot \sum_{n=e-1}^{e+g-1} \sum_{\theta} K_\theta^g \, \phi_{\theta_1}(s, \mathbf{p})\phi_{\theta_2}(s, \mathbf{p}) \cdots ,$$

where the innermost summation is over all partitions $\theta = (\theta_1, \theta_2, \ldots) \vdash n$ *of length* $e - 2(g-1)$ *having no part equal to 1, and where the coefficients* K_θ^g *are known rational constants.* □

The rational constants K_θ^g in the theorem are, up to sign, important numbers known as *Hodge integrals* which properly belong to the realm of algebraic geometry. Their appearance here reflects the deep connections between the combinatorics of the symmetric group and geometry. See the Additional Notes for further references.

2.3.4 Labelled Trees

A standard combinatorial argument shows that the generating series $w = w(x, u)$ counting rooted vertex-labelled trees with respect to number of vertices (marked by x) and edges (marked by u)

satisfies the functional equation

$$w = xe^{uw}. \tag{2.16}$$

Applying Lagrange inversion gives

$$w = \sum_{n\geq 1} n^{n-1}u^{n-1}\frac{x^n}{n!}, \tag{2.17}$$

and thus there are n^{n-1} rooted trees on n labelled vertices. Reversing the rooting process by dividing by n immediately gives the following theorem, typically credited to Cayley [10].

Theorem 2.3.6. *There are n^{n-2} trees on n labelled vertices.* □

Of course, there are n^n doubly rooted trees on n vertices and, in general, n^{n+i} labelled trees with $i+2$ independently marked vertices. Numbers of the form n^{n+i} might therefore be called *tree numbers*. The appearance of such numbers in Hurwitz's formula (2.12) makes it unsurprising that the series w of (2.17) will play a fundamental rôle in our analysis of transitive factorizations.

Definition 2.3.7. *Throughout this chapter the symbol w will be used exclusively as defined in* (2.17). *We refer to $w = w(x, u)$ as the* ***tree series****.*

Implicit differentiation of (2.16) yields the following useful formula for the generating series for doubly rooted labelled trees

$$x\frac{dw}{dx} = \frac{w}{1-uw}. \tag{2.18}$$

Combinatorially, this is reflected by the observation that the unique directed path between the roots of a doubly rooted tree decomposes the tree into an ordered sequence of rooted trees.

2.3.5 The Symmetrized Hurwitz Series

The following symmetrization of the Hurwitz series $\Phi^{(g)}$ appears in [33]. For $m \geq 1$, let Π_m be the operator

$$\Pi_m(p_\alpha) = \begin{cases} \sum_{\pi\in\mathfrak{S}_m} x_1^{\alpha_{\pi(1)}}\cdots x_m^{\alpha_{\pi(m)}} & \text{if } \ell(\alpha) = m, \\ 0 & \text{otherwise,} \end{cases} \tag{2.19}$$

extended linearly to act on all series in the p_i's. Then, with $\mathbf{x} = (x_1, \ldots, x_m)$, let $\Psi_m^{(g)}(\mathbf{x}, u)$ be the image of $\Phi^{(g)}(1, \mathbf{p}, u)$ under Π_m. It follows that

$$\Psi_m^{(g)}(\mathbf{x}, u) = \sum_{n \geq 1} \sum_{\substack{\alpha \models n \\ \ell(\alpha)=m}} H_g(\alpha) \frac{x_1^{\alpha_1}}{\alpha_1} \cdots \frac{x_m^{\alpha_m}}{\alpha_m} \frac{u^{r_g(\alpha)}}{r_g(\alpha)!}. \tag{2.20}$$

In genus 0, Theorem 2.3.3 can be used to obtain an expression for the symmetrized Hurwitz series in terms of the tree series. Before stating this result, we first introduce some further notation.

Definition 2.3.8. *The symbol w_i will be used throughout this chapter to denote the series $w(x_i, u)$. That is, $w_i = w_i(x_i, u)$ is the unique series solution to the functional equation $w_i = x_i e^{u w_i}$.*

Theorem 2.3.9. *Let $m \geq 1$ and, for $1 \leq i \leq m$, let $w_i = w(x_i, u)$. Then*

$$\Psi_m^{(0)}(\mathbf{x}, u) = u^{2m-2} \left(\sum_{i=1}^{m} x_i \frac{\partial}{\partial x_i} \right)^{m-3} \prod_{i=1}^{m} x_i \frac{d w_i}{d x_i}.$$

Proof. Use Theorem 2.3.3 together with (2.17) and (2.20). □

From (2.18) we have

$$x_i \frac{\partial}{\partial x_i} = \frac{w_i}{1 - u w_i} \frac{\partial}{\partial w_i}, \tag{2.21}$$

so that Theorem 2.3.9 can be rewritten as

$$\Psi_m^{(0)}(\mathbf{x}, u) = u^{2m-2} \left(\sum_{i=1}^{m} \frac{w_i}{1 - u w_i} \frac{\partial}{\partial w_i} \right)^{m-3} \prod_{i=1}^{m} \frac{w_i}{1 - u w_i}.$$

For $m \geq 3$, this expresses $\Psi_m^{(0)}(\mathbf{x}, u)$ as a rational function in the tree series $w_1, \ldots, w_m$. The situation is similar for all genera. In light of Theorem 2.3.5, it is known that

$$\Psi_m^{(g)}(\mathbf{x}, u) = u^{2m+2g-2} P_m^{(g)} \left(x_1 \frac{\partial}{\partial x_1}, \ldots, x_m \frac{\partial}{\partial x_m} \right) \prod_{i=1}^{m} \frac{w_i}{1 - u w_i}, \tag{2.22}$$

for all $m \geq 1$ and $g \geq 1$, where $P_m^{(g)}(a_1, \ldots, a_m)$ is a unique symmetric polynomial of total degree $m + 3g - 3$. In fact, for the partition $\alpha = (\alpha_1, \ldots, \alpha_m)$ we have

$$H_g(\alpha) = P_m^{(g)}(\alpha_1, \ldots, \alpha_m) \cdot |\mathscr{C}_\alpha| r_g(\alpha)! \prod_{i=1}^{m} \frac{\alpha_i^{\alpha_i}}{(\alpha_i - 1)!}.$$

Further details can be found in [33] and [34]. A primary goal of this thesis is to explain the combinatorial significance of the dependence of $\Psi_m^{(g)}(\mathbf{x}, u)$ on the tree series w.

2.3.6 Geometry and Hurwitz Numbers

Below we provide a sketch of the connection between the combinatorics of transitive factorizations and the geometry of branched coverings. Our description is not intended to be technically complete. For more extensive coverage, see [69], [22], or [19]. These references are listed in increasing order of the level of detail they provide.

Let $\mathbb{S}^2$ be the Riemann sphere, or, equivalently, the extended complex plane $\mathbb{C} \cup \{\infty\}$ with its usual topology. A **branched *n*-fold covering** of the sphere by a surface S of given genus is a non-constant meromorphic function $f : S \longrightarrow \mathbb{S}^2$ such that $|f^{-1}(p)| = n$ for all but a finite number of points $p \in \mathbb{S}^2$. The points p for which $|f^{-1}(p)| < n$ are called **branch points** of the covering, and all others are **regular** points. The number n is also called the **degree** of the covering. For example, the map $z \mapsto z^2$ is a degree 2 branched covering of the sphere by the sphere with branch points 0 and ∞. The covering $f' : S \longrightarrow \mathbb{S}^2$ is **equivalent** to f if there is a homeomorphism $\phi : S \longrightarrow S$ such that $f = f'\phi$.

For each $p \in \mathbb{S}^2$ there is a partition $\alpha = [1^{a_1} 2^{a_2} \cdots] \vdash n$, called the **branching type** of p, such that f behaves like $z \mapsto z^i$ locally around exactly a_i of the points in $f^{-1}(p)$. Regular points are those with branching type $[1^n]$, and a branch point with branching type $[1^{n-2}\,2]$ is said to be **simple**. An **almost simple** covering is one in which all branch points, except possibly one, are simple.

Roughly speaking, one can view a branched n-fold covering of the sphere as n labelled sheets (*i.e.* copies of the extended complex plane) wrapped about the sphere in such a way that they interact over only a finite number of points. These points are the branch points of the cover. The manner in which the sheets interact over a given branch point p is dictated by a permutation $\pi \in \mathfrak{S}_n$ of the their labels. In particular, starting on sheet i, a counterclockwise tour on the covering surface over p will terminate on sheet $\pi(i)$. The cycle type of π is the branching type of p. For example, π is the identity precisely when p has branching type $[1^n]$, in which case p is a regular point and the sheets over p are mutually disjoint. If π is a transposition, then the branching type of p is $[1^{n-2}\,2]$, so that p is a simple branch point. In this case, only two sheets interact over p.

If $P_1, \ldots, P_m$ are the branch points of a degree n covering f, and $\pi_1, \ldots, \pi_m$ are their associated permutations, then the consistency relation $\iota = \pi_1 \cdots \pi_m$ can be deduced geometrically. In the case where $P_1, \ldots, P_{m-1}$ are simple branch points, so that π_i is a transposition τ_i for $1 \le i \le m-1$, it follows that $\pi_m = \tau_{m-1} \cdots \tau_1$. Thus we obtain from f a factorization of π_m into transpositions.

Moreover, this factorization is transitive precisely when f is a *connected* covering, and its genus is the genus of the covering surface.

Up to a known scaling that accounts for the (artificial) labelling of sheets, the Hurwitz number $H_g(\alpha)$ is therefore seen to count inequivalent, connected, almost simple, degree n coverings of the sphere by a surface of genus g, for which ∞ has branching type α. The formula $r_g(a) = n + \ell(\alpha) + 2g - 2$ for the length of the corresponding factorizations is, in the geometrical context, a consequence of the *Riemann-Hurwitz formula*.

2.3.7 Additional Notes

The application of representation theory to the problem of enumerating permutation factorizations was initiated by Hurwitz [44], who showed that the answers to such problems could be expressed in terms of the irreducible characters of $\mathfrak{S}_n$. Equation (2.9) appears in [29]. In fact, Goulden and Jackson used this expression and (2.11) to generate the data which led them to conjecture the Hurwitz formula. (Hurwitz's work was not known to them at the time.) Mednykh [52, 53] also gives a complete solution of the general Hurwitz enumeration problem in terms of complicated expressions involving character sums.

Computations in $\mathbb{C}\mathfrak{S}_n$ like those of §2.3.1 were first used by Stanley [66] to count factorizations of permutations into full cycles. Jackson [47] applied these same methods to obtain more general results, including Theorem 2.3.1. A combinatorial proof of Theorem 2.3.1 can be found in [25].

Hurwitz first stated the formula bearing his name in [45]. His work was largely forgotten for nearly a century, during which time various authors rediscovered the formula, in whole or in part. Dénes [13] showed combinatorially that $H_0((n)) = n^{n-2}$, and the physicists Crescimanno and Taylor [12] found the expression $H_0([1^n]) = n^{n-3}(2n-2)!$. Arnol'd [2] was able to obtain a formula for $H_0((p,q))$. A detailed exposition of Goulden and Jackson's proof of Theorem 2.3.3 is contained in [57]. Hurwitz himself had obtained (2.13), in the form of a recurrence, but did not fully prove that it is satisfied by the numbers that bear his name. He did, however, provide insight on how such a proof might proceed. Strehl offers a possible reconstruction of Hurwitz's ideas in [70].

Bousquet-Mélou and Schaeffer [8] have recently derived the Hurwitz formula, in full generality, as a consequence of a bijection between a class of factorizations more general than those considered here and certain rooted trees. Their proof of the formula is not truly bijective, however, as the final stage of their argument requires inclusion-exclusion to restrict to factorizations into transpositions. We defer further discussion on their approach until §3.2.7. At the time of writing, no bijective proof of Theorem 2.3.3 has been found.

Branched coverings of the sphere have been counted by analytic methods that altogether avoid Hurwitz's encoding of the problem in terms of permutation factorizations. In particular, singularity theory and the analysis of the *Lyashko-Looijenga map* have led to substantial results. Briefly, the Lyashko-Looijenga map assigns to a meromorphic function f the polynomial whose roots are the critical values of f. When its domain is restricted to almost simple n-fold coverings of the sphere by a surface of genus g for which ∞ has branched type $\alpha \vdash n$, the Lyashko-Looijenga map is a finite covering of the space of monic polynomials of degree $r_g(\alpha) = n + \ell(\alpha) + 2g - 2$. The degree of this covering bears a simple relation to the Hurwitz number $H_g(\alpha)$, and can be computed through other methods. The formula for $H_0(n)$ follows from Looijenga's inaugural work [50]. This was extended by Arnol'd [1] to evaluate $H_0((p,q))$, and then by Goryunov and Lando [23] to arrive at the general formula for $H_0(\alpha)$, with arbitrary α. Later, in the seminal paper [17], Ekedahl, Lando, Shapiro, and Vainshtein pushed these ideas much further to prove that the Hurwitz numbers (of all genera) are related to particular *Hodge integrals*, which are intersection numbers for the Chern classes of certain line bundles on the moduli space of complex curves.

Vakil [72] has also given a proof of the Hurwitz formula in the context of enumerative geometry. Using the theory of stable maps, he derives recursions satisfied by $H_g(\alpha)$, for arbitrary α, in genera $g = 0$ and $g = 1$. He then observes that these recursions are also satisfied by the solutions of certain straightforward graph enumeration problems. For $g = 0$, counting the relevant graphs leads to Theorem 2.3.3. When $g = 1$, the graph counting problem also admits a closed form solution, resulting in the formula

$$H_1(\alpha) = \frac{1}{24}(n+m)!\left(n^n - n^{n-1} - \sum_{i=2}^{m}(i-2)!e_i n^{m-i}\right)\prod_{i=1}^{m}\frac{\alpha_i^{\alpha_i}}{(\alpha_i - 1)!}.$$

Here $\alpha = (\alpha_1, \ldots, \alpha_m)$ is a partition of n and e_i is the i-th elementary symmetric function evaluated at $(\alpha_1, \ldots, \alpha_m)$. Although the classes of labelled graphs that are enumerated to obtain these results are very simple to describe, no bijection between them and factorizations has been found. Combinatorializing the geometric argument that leads to Vakil's recursions appears to be difficult. The formula above was originally conjectured by Goulden and Jackson in [33], and also proved by them in [30] using methods completely different from Vakil's. See also Appendix B.

Theorem 2.3.4 is first stated explicitly in [64], though the result upon which it is based (namely, Theorem 2.3.1) appears earlier in [47] and [25]. The proof given in [64] is identical to that of [47], but the the connection with geometry was not observed in the latter paper. Formulas for $H_1((p,q))$ for special values of p and q are also given in [64].

Minimal transitive factorizations are also known to be related to parking functions. In particular, [6] contains a bijection between parking functions and minimal factorizations of full cycles, thus giving another proof that $H_0((n)) = n^{n-2}$. More recently, a bijection between *prime* parking functions and transitive factorizations of class $(1, n-1)$ has been found by Kim and Seo [48], thereby proving that $H_0((n-1, 1)) = (n-1)^n$. Their methods have been extended by Rattan [60].

2.4 Graphical Representation of Factorizations

None of the approaches to the enumeration of factorizations described in the previous section is fully satisfying from a purely combinatorial perspective, as each relies on algebraic arguments for which no combinatorial interpretation is known. In fact, at present, there is no known bijective proof of Theorem 2.3.3, despite its strikingly simple form. We wish to better understand Hurwitz's formula from a combinatorial standpoint.

In this section we introduce a general, graphical representation of factorizations that we shall exploit throughout the remainder of the chapter (indeed, throughout this entire thesis). We begin with Dénes' well known encoding of minimal factorizations of full cycles as vertex- and edge-labelled trees. The balance of the section is devoted to extending of this encoding to give bijections between arbitrary factorizations and certain classes of labelled maps.

2.4.1 Counting Minimal Factorizations

Let π be a full cycle in $\mathfrak{S}_n$. Since π acts transitively on $[n]$, every factorization of π is transitive. In particular, minimal transitive factorizations of π are identified with minimal factorizations of π, which are precisely the factorizations of length $n-1$.

Dénes [13] discovered the following formula for the number $H_0((n))$ of minimal transitive factorizations of the full cycle $(1\,2\,\cdots\,n)$. The formula is interesting in its own right, but the method of proof is truly intriguing. We shall actually reprove this result in a more general setting in §2.4.7, so some details are suppressed in the proof given here.

Theorem 2.4.1 (Dénes)**.** *There are n^{n-2} minimal transitive factorizations of any full cycle in $\mathfrak{S}_n$.*

Sketch proof: Let $f = (\tau_{n-1}, \ldots, \tau_1)$ be a minimal (transitive) factorization of a full cycle in $\mathfrak{S}_n$, where $\tau_i = (a_i\, b_i)$ for $1 \le i \le n-1$. Construct the graph T_f having labelled vertices $\{1, 2, \ldots, n\}$ and edges $\{\{a_i\, b_i\} : 1 \le i \le n-1\}$. Assign label i to edge $\{a_i, b_i\}$, for $1 \le i \le n-1$. Note that the transitivity of f implies T_f is connected. Thus T_f is a vertex- and edge-labelled tree. This

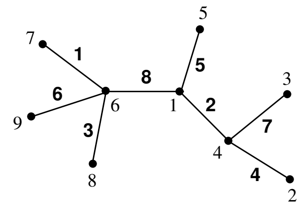

Figure 2.1: The tree corresponding to the factorization (2.23).

construction is reversible, and $f \mapsto T_f$ is seen to be a bijection between minimal factorizations of full cycles in $\mathfrak{S}_n$ and vertex- and edge-labelled trees on n vertices. As there are n^{n-2} trees on n labelled vertices, and $(n-1)!$ edge labellings of each, there are $(n-1)!\, n^{n-2}$ minimal factorizations of full cycles in $\mathfrak{S}_n$. The result follows by symmetry, since there are $(n-1)!$ full cycles on n symbols. □

Example 2.4.2. For example, the tree corresponding to the factorization

$$(1\,2\,3\,4\,5\,6\,7\,8\,9) = (1\,6)(3\,4)(6\,9)(1\,5)(2\,4)(6\,8)(1\,4)(7\,6) \tag{2.23}$$

under Dénes' correspondence is drawn in Figure 2.1. □

Consider now a minimal factorization f of the permutation $\pi \in \mathfrak{S}_n$. From Corollary 2.2.10, a minimal factorization of π is a shuffling of minimal factorizations of its disjoint cycles. If π has cycle type $(\alpha_1, \ldots, \alpha_m)$, then by Theorem 2.4.1 its cycles can be minimally factored in $\alpha_1^{\alpha_1-2} \cdots \alpha_m^{\alpha_m-2}$ ways, and the $\sum_i (\alpha_i - 1) = n - m$ resulting factors can be shuffled in $\binom{n-m}{\alpha_1-1,\ldots,\alpha_m-1}$ ways. This proves the following result, which can also be found in [13].

Corollary 2.4.3. *There are*

$$(n-m)! \prod_{i=1}^{m} \frac{\alpha_i^{\alpha_i-1}}{\alpha_i!}$$

minimal factorizations of any permutation π having cycle type $(\alpha_1, \ldots, \alpha_m) \vdash n$. □

This result can also be put in a graphical context. By mimicking Dénes' proof, we find that minimal factorizations of class $\alpha = (\alpha_1, \ldots, \alpha_m) \vdash n$ are in correspondence with vertex- and edge-labelled forests consisting of trees $T_1, \ldots, T_m$ having $\alpha_1, \ldots, \alpha_m$ vertices, respectively. It is easy to count such forests and, upon doing so and dividing by a symmetry factor, we obtain Corollary 2.4.3.

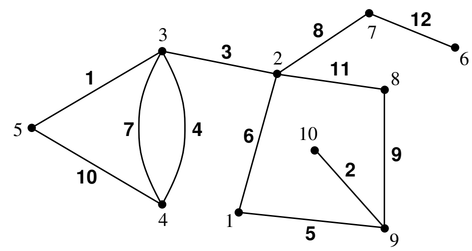

Figure 2.2: The graph of the factorization (2.24).

2.4.2 The Graph of a Factorization

Let $f = (\tau_r, \dots, \tau_1)$ be a factorization in $\mathfrak{S}_n$ with factors $\tau_i = (a_i\, b_i)$ for $1 \le i \le r$. Following the proof of Theorem 2.4.1, we construct a graph $\mathscr{G}_f$ on the vertices $\{1, 2, \dots, n\}$ by interpreting the transposition τ_i as the edge $\{a_i, b_i\}$, for $1 \le i \le r$, and assigning this edge the label i. Factors occurring more than once in f correspond to multiple edges of $\mathscr{G}_f$.

We refer to $\mathscr{G}_f$ as the **graph of** $\boldsymbol{f}$. For example, Figure 2.2 shows the graph of the factorization

$$(1\,2\,3\,4\,5)(6\,7\,8\,9\,10) = (6\,7)(2\,8)(4\,5)(8\,9)(2\,7)(3\,4)(1\,2)(1\,9)(3\,4)(2\,5)(9\,10)(3\,5). \tag{2.24}$$

Clearly, $\mathscr{G}_f$ completely encodes the factorization f. Thus $f \mapsto \mathscr{G}_f$ is a one-one correspondence between factorizations in $\mathfrak{S}_n$ of length r and loopless graphs on n labelled vertices and r labelled edges.

Dénes' combinatorial derivation of the number of minimal transitive factorizations of a full cycle (Theorem 2.4.1) naturally compels us to seek a similar proof of the more general Hurwitz formula. In analogy with the graphical derivation of Corollary 2.4.3, it is reasonable to conjecture that the factors $\alpha_i^{\alpha_i}$ in Hurwitz's formula correspond to tree-like structures in the graph of a minimal transitive factorization. The factor $(n + \ell(\alpha) - 2)!$ probably again corresponds to an edge-labelling of this graph, but the factor $n^{\ell(\alpha)-3}$ seems difficult to explain combinatorially; in particular, the appearance of n^{-3} may well correspond to an elusive symmetry.

The simple relationship (2.15) between the generating series for transitive factorizations and the implicitly defined series s of (2.14) also calls for a combinatorial explanation along these lines. For instance, the differential operator on the left-hand side of (2.15) corresponds to the marking of two vertices in the graph of a factorization. The series s, on the other hand, is reminiscent of the functional equation $T = ze^T$ for the generating series $T = T(z)$ for labelled rooted trees. In fact, the number n^n of doubly-rooted labelled trees on n vertices appears in $\phi_0(s, \mathbf{p})$, and so the indeterminate p_n may serve to record the number of trees of a given size that should be pasted together to form the graph of a factorization. No such combinatorial understanding of the rôle of s is currently known, though the interest in finding one is underscored by the ubiquity of s in Theorem 2.3.5. The connections with geometry mentioned there, *viz.* Hodge integrals, suggest that such an understanding could bring with it fresh combinatorial insight into the geometry of the moduli space of curves.

Despite all this tantalizing combinatorial structure, the only transitive factorizations currently understood from a natural combinatorial standpoint are genus 0 factorizations of full cycles, for which Theorem 2.4.1 provides a simple characterization. (Schaeffer and others [41, 58] have given combinatorial interpretations of certain computations in $\mathbb{C}\mathfrak{S}_n$ that enable them to count factorizations of full cycles of arbitrary genus, but their approach is not particularly satisfying, as the combinatorics seems far from natural.) Our investigation of the graphs of factorizations is motivated by a desire to extend this understanding and, in particular, to explain the significance of tree-like structure in factorizations.

2.4.3 Carriers and Orbits

Observe that the set of edges incident with a vertex v in the graph of $f = (\tau_r, \ldots, \tau_1)$ corresponds with the set of factors of f which move the symbol v. That is, $\delta(v) = \{e \in [r] : \tau_e(v) \neq v\}$. We shall use this basic connection to translate properties of a factorization into properties of its graph.

For completeness, we begin with a formal proof of the fact that connectivity of $\mathscr{G}_f$ characterizes transitivity of f. More generally, it can be shown that the connected components of $\mathscr{G}_f$ are the graphs of the components of f.

Proposition 2.4.4. *A factorization is transitive if and only if its graph is connected.*

Proof. Let $f = (\tau_r, \ldots, \tau_1)$ be a factorization of $\pi \in \mathfrak{S}_n$ and let $S = \langle \tau_1, \ldots, \tau_r \rangle$. Recall that f is transitive if and only if S acts transitively on $[n]$.

Suppose first that $\mathscr{G}_f$ is connected. Then for any $a, b \in [n]$ with $a \neq b$ there must be a walk $v_0, e_0, \ldots, e_k, v_{k+1}$ in $\mathscr{G}_f$ from $v_0 = a$ to $v_{k+1} = b$. Thus we have $\tau_{e_j} = (v_j\, v_{j+1})$ for all $0 \leq j < k$, and so the product $\sigma = \tau_{e_k} \cdots \tau_{e_1} \in S$ satisfies $\sigma(a) = b$. Hence S acts transitively on $[n]$.

Assume now that f is transitive and choose $a, b \in [n]$ with $a \neq b$. Then there is some product $\sigma = \tau_{e_m} \cdots \tau_{e_0} \in S$ satisfying $\sigma(a) = b$, and this product determines a walk from a to b in $\mathscr{G}_f$ as follows: Let $\tau_{e_{i_0}}$ be the first (rightmost) factor that moves $v_0 = a$, and set $v_1 = \tau_{e_{i_0}}(v_0)$. Now let $\tau_{e_{i_1}}$ be the first factor after $\tau_{e_{i_0}}$ which moves v_1, and set $v_2 = \tau_{e_{i_1}}(v_1)$. Proceed in this manner until

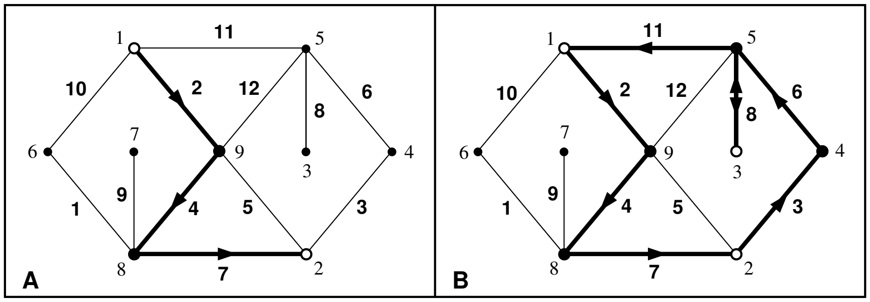

Figure 2.3: (A) The carrier and (B) the orbit of vertex 1 in the graph of a factorization.

a vertex v_k is obtained such that no factor after $\tau_{e_{i_k}}$ moves v_{k+1}. By construction we have

$$v_{k+1} = (\tau_{e_{i_k}} \cdots \tau_{e_{i_0}})(v_0) = (\tau_{e_m} \cdots \tau_{e_0})(v_0) = \sigma(a) = b,$$

and hence $v_0, e_{i_0}, v_1, e_{i_1}, \cdots, v_k, e_{i_k}, v_{k+1}$ is a walk from a to b. Therefore $\mathscr{G}_f$ is connected. $\square$

The construction in the second half of the proof of Proposition 2.4.4 can be applied to the product $\pi = \tau_r \cdots \tau_1$ and any $v \in [n]$ to define a walk $v_0, e_0, \cdots, e_k, v_{k+1}$ in $\mathscr{G}_f$ from $v_0 = v$ to $v_{k+1} = \pi(v)$. More precisely, this walk is uniquely determined by the conditions $v_0 = v$, $v_{k+1} = \pi(v)$, and

$$e_i = \begin{cases} \min \delta(v_0) & \text{if } i = 0, \\ \min\{e \in \delta(v_i) : e > e_{i-1}\} & \text{if } 0 < i < k, \\ \max \delta(v_{k+1}) & \text{if } i = k. \end{cases} \tag{2.25}$$

We call this walk the **carrier** of v. Figure 2.3A illustrates the carrier of $v = 1$ in the graph of a factorization of $\pi = (1\,2\,3)(4\,5)(6\,7\,8\,9)$. Notice that it starts at v and ends at $\pi(v) = 2$, with edge labels increasing along the walk.

If v lies on the cycle $(v\,\pi(v)\,\cdots\pi^m(v))$, then concatenating the carriers of $v, \pi(v), \ldots, \pi^m(v)$ results in a closed walk which we call the **orbit** of v. Of course, the orbit of v is equivalent to the orbit of $\pi^s(v)$ for any s. Figure 2.3B illustrates the orbit of $v = 1$ in the graph of a factorization of $\pi = (1\,2\,3)(4\,5)(6\,7\,8\,9)$. It is obtained by concatenating the carriers of $v = 1, \pi(v) = 2$, and $\pi^2(v) = 3$. Hollow vertices mark the endpoints of these carriers. Note that the orbit traverses edge 8 twice, once in each direction.

The following lemma serves to identify a closed walk in $\mathscr{G}_f$ as an orbit.

Lemma 2.4.5. *Let f be a factorization of π and let $W = ((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be a closed walk in $\mathscr{G}_f$. If $(e_{i-1}, v_i, e_i) \neq (e_{j-1}, v_j, e_j)$ for $i \neq j$ and the conditions*

$$e_i = \begin{cases} \min \delta(v_i) & \text{if } e_{i-1} = \max \delta(v_i), \\ \min\{e \in \delta(v_i) \,:\, e > e_{i-1}\} & \text{otherwise}, \end{cases} \tag{2.26}$$

are satisfied, then W is the orbit of some vertex v. In particular, if $D = \{i \,:\, e_{i-1} \geq e_i\}$ and $D^\circ = (i_0, \ldots, i_m)^\circ$, then $(v_{i_0}\, v_{i_1}\, \cdots\, v_{i_m})$ is a cycle of π and W is the orbit of v_{i_0}.

Proof. First observe that $k > 0$ since $\mathscr{G}_f$ is loopless. It follows that $D \neq \emptyset$, as otherwise we would have $e_0 < \ldots < e_k < e_0$ with $k > 0$. Choose any s and set $p = i_s$ and $q = i_{s+1}$. Then the inequalities $e_{p-1} \geq e_p$ and $e_{q-1} \geq e_q$, together with the conditions (2.26), imply that $e_p = \min \delta(v_p)$, $e_{q-1} = \max \delta(v_q)$, and $e_i = \min\{e \in \delta(v_i) \,:\, e > e_{i-1}\}$ for $p < i < q-1$. Thus the walk $v_p, e_p, \ldots, v_{q-1}, e_{q-1}, v_q$ satisfies conditions (2.25). It is therefore the carrier of v_p. It follows that $v_{i_{s+1}} = \pi(v_{i_s})$ for all s and that W is the concatenation of the carriers of $v_{i_0}, v_{i_1}, \ldots, v_{i_m}$. Finally, the condition $(e_{i-1}, v_i, e_i) \neq (e_{j-1}, v_j, e_j)$ for $i \neq j$ ensures that none of these carriers coincide. Therefore $(v_{i_0}\, v_{i_1}\, \ldots\, v_{i_m})$ is a cycle of π and W is the orbit of v_{i_0}. □

2.4.4 The Map of a Factorization

Let $\mathscr{G}$ be a connected, loopless, vertex- and edge-labelled graph, with n vertices and m edges. Through the correspondence described below, $\mathscr{G}$ is associated with a unique loopless, vertex- and edge-labelled map.

Let L be the set of all $2m$ symbols of the form e^v in which edge e is incident with vertex v of $\mathscr{G}$. For each edge $e = \{u, v\}$ of $\mathscr{G}$, let τ_e be the transposition $(e^u\, e^v) \in \mathfrak{S}_L$. For each vertex v, let c_v be the k-cycle $(e_1^v \,\cdots\, e_k^v) \in \mathfrak{S}_L$, where $(e_1, \ldots, e_k)^\circ = \delta(v)^\circ$. Let $\nu = \prod_{v \in V} c_v$ and $\epsilon = \prod_{e \in E} \tau_e$. Then the pair (ϵ, ν) is a rotation system on the symbols L. Moreover, since $\mathscr{G}$ is connected, this rotation system is transitive. (An argument similar to the proof of Proposition 2.4.4 formally proves this claim.) By Theorem 1.3.4, (ϵ, ν) corresponds to a unique loopless map with half-edges labelled by L. These half-edge labels induce vertex labels and edge labels in the obvious way, and so we obtain a loopless, vertex- and edge-labelled map whose skeleton is $\mathscr{G}$. We write $\mathscr{M}(\mathscr{G})$ for this map.

Example 2.4.6. Consider the graph $\mathscr{G}$ presented in Figure 2.4A. The corresponding transitive rotation system is (ϵ, ν), where

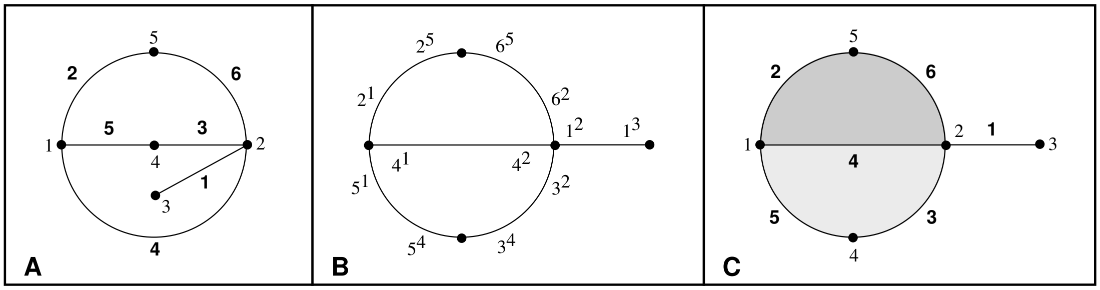

Figure 2.4: Constructing the map of a factorization from its graph.

$$\epsilon = (1^2\,1^3)(2^1\,2^5)(3^2\,3^4)(4^1\,4^2)(5^1\,5^4)(6^2\,6^5)$$
$$\nu = (2^1\,4^1\,5^1)(1^2\,3^2\,4^2\,6^2)(1^3)(3^4\,5^4)(2^5\,6^5).$$

For example, edge 3 of $\mathscr{G}$ contributes the transposition $\tau_3 = (2\,4)$ to ϵ and vertex 2 contributes the 4-cycle $c_2 = (1^2\,3^2\,4^2\,6^2)$ to ν. Figure 2.4B illustrates the (planar) half-edge-labelled map corresponding to this rotation system through Theorem 1.3.4. Finally, Figure 2.4C shows the loopless, vertex- and edge-labelled map $\mathscr{M}(\mathscr{G})$ associated with $\mathscr{G}$. The two internal faces of this map have been shaded to underscore the distinction between it and the original graph. Notice that the edge labels encountered along a clockwise tour around any vertex appear in cyclic increasing order. □

We now define the *rotator* of a vertex in an edge-labelled map, a fundamental construct that is analogous to a circulator in a half-edge-labelled map.

Definition 2.4.7. *Let $\mathscr{M}$ be an edge-labelled map. The* ***rotator*** *of a vertex v of $\mathscr{M}$ is the cyclic list of edge labels encountered along a clockwise tour of small radius about v.*

For example, in the map of Figure 2.4C, the rotators of vertices 1 and 2 are, respectively, $(2, 4, 5)^\circ$ and $(1, 3, 4, 6)^\circ$. In general, observe that $\mathscr{M}(\mathscr{G})$ is constructed so that $\delta(v)^\circ$ is the rotator of vertex v. Thus $\mathscr{G} \mapsto \mathscr{M}(\mathscr{G})$ is a bijection between connected, loopless, vertex- and edge-labelled graphs and loopless, vertex- and edge-labelled maps whose rotators are increasing.

By virtue of Proposition 2.4.4, it follows that $f \mapsto \mathscr{M}(\mathscr{G}_f)$ is a bijection between transitive factorizations and loopless, vertex- and edge-labelled maps with increasing rotators.

Definition 2.4.8. *The map $\mathscr{M}(\mathscr{G}_f)$ corresponding to the transitive factorization f is called the* ***map of f****, and will be denoted simply by $\mathscr{M}_f$.We write* MAP *for the bijection $f \mapsto \mathscr{M}_f$.*

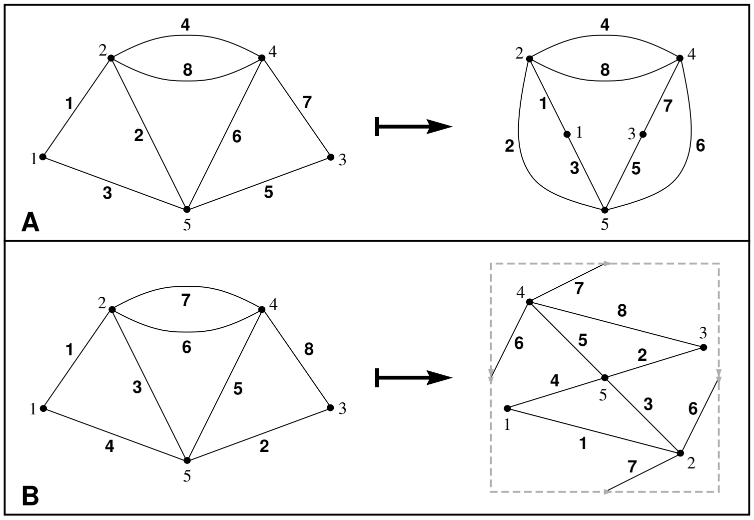

Figure 2.5: Graphs and maps of the factorizations (2.27) and (2.28).

Example 2.4.9. Figure 2.5A displays the graph (on the left) and map (on the right) of the factorization

$$(1)(2)(3)(4)(5) = (2\,4)(3\,4)(4\,5)(3\,5)(2\,4)(1\,5)(2\,5)(1\,2), \tag{2.27}$$

and Figure 2.5B does the same for the factorization

$$(1)(2\,3)(4\,5) = (3\,4)(2\,4)(2\,4)(4\,5)(1\,5)(2\,5)(3\,5)(1\,2). \tag{2.28}$$

Notice that the map of the latter factorization is of genus 1, despite the fact that the graphs of both factorizations are planar. In fact, both factorizations have genus equal to that of their map. This is not a coincidence. As we shall see in §2.4.6, the bijection MAP generally preserves genus. □

2.4.5 Edge-Labelled Maps and Descent Structure

In what follows, we shall be concerned only with maps that arise from factorizations. Since every such map is loopless and edge-labelled with increasing rotators, it will avoid a great deal of redundancy to absorb these two properties into the definition of a map. Unless otherwise stated, the following conditions are assumed throughout the remainder of Chapter 2.

- All maps are loopless.
- All maps are edge-labelled in such a way that rotators are increasing.

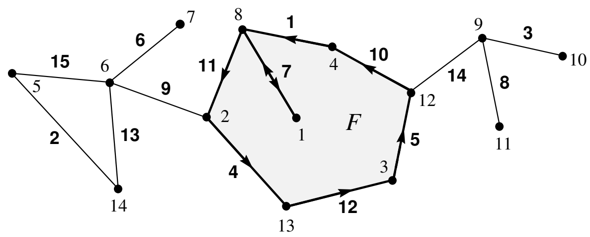

Figure 2.6: Descent structure of the map of the factorization (2.29).

For instance, under these conventions, the correspondence MAP is a bijection between transitive factorizations and vertex-labelled maps. We now introduce some fundamental definitions that apply to any (loopless, edge-labelled) map.

Let F be a face of the map $\mathscr{M}$, and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the boundary walk of F. If $e_{i-1} \geq e_i$, then we call the pair (e_{i-1}, e_i) a **descent** of face F, and we say that vertex v_i is **at a descent** of F. The corner (e_{i-1}, v_i, e_i) of F identified by the descent (e_{i-1}, e_i) is called a **descent corner**. The set $\{v_i \,:\, e_{i-1} \geq e_i\}$ of all vertices at descents of F is said to be the **descent set** of F. There is a natural cyclic ordering of this set, obtained by listing the vertices in the order in which they appear along the boundary walk of F. The resulting cyclic sequence is called the **descent cycle** of F. Finally, since the rotator of every vertex is increasing, each vertex is at a descent of exactly one face. Thus the descent sets of the faces of $\mathscr{M}$ are disjoint and partition the vertex set. If $\mathscr{M}$ has m_i faces with i descents, then $[1^{m_1}2^{m_2}\cdots]$ is called the **descent partition** of $\mathscr{M}$.

Example 2.4.10. Figure 2.6 illustrates these definitions with the map $\mathscr{M}_f$ of the factorization

$$(5\ 6)(9\ 12)(6\ 14)(3\ 13)(2\ 8)(4\ 12)(2\ 6)(9\ 11)(1\ 8)(6\ 7)(3\ 12)(2\ 13)(9\ 10)(5\ 14)(4\ 8) \qquad (2.29)$$

of $\pi = (1\,2\,3\,4)(5)(6\,7\,8\,9\,10\,11\,12\,13\,14)$. The boundary walk of the shaded face F of this map is highlighted; the direction of traversal keeps F on the left and pivots clockwise at each vertex. Formally, the boundary walk is

$$((v_0, e_0), \ldots, (v_7, e_7))^\circ = ((1, 7), (8, 11), (2, 4), (13, 12), (3, 5), (12, 10), (4, 1), (8, 7))^\circ,$$

which has four descents, namely $(11, 4)$, $(12, 5)$, $(10, 1)$, and $(7, 7)$. Vertex labels have been placed at descent corners throughout the figure. In particular, vertices 1, 2, 3 and 4 are at descents of F, so that F has descent set $\{1, 2, 3, 4\}$. The descent cycle of F is seen to be $(1, 2, 3, 4)^\circ$. Since the other faces of $\mathscr{M}_f$ have 1 and 9 descents, the descent partition of $\mathscr{M}_f$ is $(9, 4, 1) \vdash 14$.

Notice that the boundary walk of F is simply the orbit of vertex 1. Moreover, this orbit is the concatenation of the carriers of the vertices 1, 2, 3, and 4 that are at descents of F. With Lemma 2.4.5, this explains the coincidence of the descent cycle of F and the cycle $(1\,2\,3\,4)$ of the target permutation π. These observations will be formalized in §2.4.6, below. □

2.4.6 A Bijection Between Factorizations and Maps

The following theorem is central to our discussion. It describes how the genus and class of a factorization are encoded in its map.

Theorem 2.4.11. *The correspondence* MAP : $f \mapsto \mathscr{M}_f$ *restricts to a bijection between genus* g *factorizations of class* α *and genus* g *vertex-labelled maps with descent partition* α. *Moreover, if* f *is a factorization of* $\pi \in \mathfrak{S}_n$, *then the descent cycles of* $\mathscr{M}_f$ *coincide with the cycles of* π.

Proof. We have already seen that MAP is a bijection between transitive factorizations and vertex-labelled maps. Thus we need only show that the map of a genus g factorization of class α is indeed of genus g with descent partition α.

Let f be a genus g factorization of $\pi \in \mathscr{C}_\alpha$. Let F be a face of $\mathscr{M}_f$ and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be its boundary walk. Then this walk passes each corner of F exactly once, so that $(e_{i-1}, v_i, e_i) \neq (e_{j-1}, v_j, e_j)$ for $i \neq j$. Furthermore, since $\mathscr{M}_f$ is loopless we can assert unambiguously that e_i immediately follows e_{i-1} in the rotator of vertex v_i. But the rotator of v_i is $\delta(v_i)^\circ$, so we have

$$e_i = \begin{cases} \min \delta(v_i) & \text{if } e_{i-1} = \max \delta(v_i), \\ \min\{e \in \delta(v_i) \,:\, e > e_{i-1}\} & \text{otherwise.} \end{cases}$$

Let $D = \{i \,:\, e_{i-1} \geq e_i\}$ index the descents of F, and let $D^\circ = (i_0, \ldots, i_m)^\circ$. Then $(v_{i_0}, \ldots, v_{i_m})^\circ$ is the descent cycle of F, and Lemma 2.4.5 implies that this coincides with a cycle of π. Since every vertex is at a descent of exactly one face, this correspondence between faces of $\mathscr{M}_f$ and cycles of π is one-one. In particular, the descent partition of $\mathscr{M}_f$ coincides with the cycle type of π.

Finally, suppose $\alpha \vdash n$. Then $\mathscr{M}_f$ has n vertices and $\ell(\alpha)$ faces, as its descent partition is α. Since f must be of length $r_g(\alpha) = n + \ell(\alpha) + 2g - 2$ we find that $\mathscr{M}_f$ also has $n + \ell(\alpha) + 2g - 2$ edges. The Euler-Poincaré formula identifies g as the genus of $\mathscr{M}_f$. □

Observe that Theorem 2.4.11 identifies the Hurwitz series $\Phi^{(g)}(z, \mathbf{p}, u)$, defined in (2.10), as the generating series for genus g vertex-labelled maps with respect to labelled vertices (marked by z), labelled edges (marked by u), and descent partition (marked by $\mathbf{p}$).

Let f be a transitive factorization. The edge-labelling of $\mathscr{M}_f$ determines the descent structure of $\mathscr{M}_f$, and hence, by Theorem 2.4.11, the class of f. Relabelling the vertices of $\mathscr{M}_f$ results in a new map whose associated factorization is of the same class as f. The next proposition shows that relabelling almost always results in a map distinct from $\mathscr{M}_f$, which in turn corresponds to a factorization distinct from f.

Proposition 2.4.12. *A map on $n \neq 2$ vertices has no nontrivial automorphisms.*

Proof. Suppose ϕ is a nontrivial automorphism of the map $\mathscr{M}$. Since ϕ preserves edge-labels, it cannot fix all vertices. Let u and v be distinct vertices with $\phi(u) = v$. Then u and v must have the same rotator, since isomorphisms preserve rotators. This implies u and v are adjacent to each other, and nothing else. Thus $\mathscr{M}$ has exactly two vertices. □

Corollary 2.4.13. *Let $\alpha \vdash n$, where $n \neq 2$. Then there are $|\mathscr{C}_\alpha| H_g(\alpha)/n!$ genus g maps with descent partition α.*

Proof. Let $M_g(\alpha)$ be the number of genus g maps with descent partition α. If $n \neq 2$, the proposition implies there are $n! M_g(\alpha)$ vertex-labelled genus g maps with descent partition α. By Theorem 2.4.11, there are the same number of genus g factorizations of class α. That is, $n! M_g(\alpha) = |\mathscr{C}_\alpha| H_g(\alpha)$. □

The corollary may leave some doubt as to the nature of maps on only two vertices. The next proposition provides a full description of these maps, and will be used later.

Proposition 2.4.14. *For $g \geq 0$, there are exactly two maps with only two vertices. One of these maps has one face and $2g + 1$ edges, and the other has two faces and $2g + 2$ edges.*

Proof. The two possible vertex labellings of map on two vertices are obviously equivalent. Thus the number of genus g maps on two vertices with descent partition $\alpha \vdash 2$ is equal to the number of genus g factorizations of class α. If $\alpha = (2)$, then $(1\,2) = (1\,2)^{2g+1}$ is clearly the only such factorization, where the notation $(1\,2)^k$ means k copies of the factor $(1\,2)$. If $\alpha = (1, 1)$, then $(1)(2) = (1\,2)^{2g+2}$ is the only factorization. □

The previous proposition establishes a (trivial) topological result through the link between factorizations and maps forged by Theorem 2.4.11. Generally speaking, our overall purpose is to exploit this link in the opposite direction and investigate factorizations through their associated maps. We now consider a brief example that may illustrate the usefulness of this enterprise.

Let $\sigma \in \mathfrak{S}_n$. Then permuting the vertex labels of a map $\mathscr{M}$ by replacing i with $\sigma(i)$ is equivalent to conjugating the factors of the factorization f corresponding to $\mathscr{M}$. That is, if $f = (\tau_r, \ldots, \tau_1)$, then the relabelled map corresponds to the factorization $(\sigma\tau_r\sigma^{-1}, \ldots, \sigma\tau_1\sigma^{-1})$. Proposition 2.4.12, which essentially states that all vertex labellings of a map are inequivalent, is therefore equivalent to the following result. The proof given here is based entirely in $\mathfrak{S}_n$, and should be compared with the simple topological proof of Proposition 2.4.12.

Proposition 2.4.15. *Let $f = (\tau_r, \ldots, \tau_1)$ be a transitive factorization of $\pi \in \mathfrak{S}_n$. For $\sigma \in \mathfrak{S}_n$, let $f_\sigma = (\sigma\tau_r\sigma^{-1}, \ldots, \sigma\tau_1\sigma^{-1})$. Then $f_{\sigma_1} = f_{\sigma_2}$ implies $\sigma_1 = \sigma_2$, except in the case $n = 2$.*

Proof. When $n = 2$, each τ_i is the transposition $(1\,2)$ and so $f_{\sigma_1} = f_{\sigma_2}$ for all $\sigma_1, \sigma_2 \in \mathfrak{S}_2$. Now let $n > 2$, and suppose that $f_{\sigma_1} = f_{\sigma_2}$ for $\sigma_1 \neq \sigma_2$. Then $\tau_i\rho = \rho\tau_i$ for each $i = 1, \ldots, r$, where $\rho = \sigma_2\sigma_1^{-1}$. Since $\rho \neq \iota$ there is some transposition $\tau_j = (a\,b)$ such that ρ does not fix both a and b. But $(a\,b)\rho = \rho(a\,b)$, so it follows that $\rho(a) = b$ and $\rho(b) = a$. Since f is transitive and $n > 2$ there must be a factor τ_k equal to either $(a\,c)$ or $(b\,c)$, where $c \neq a, b$. Suppose, without loss of generality, that $\tau_k = (a\,c)$. Then $(a\,c)\rho = \rho(a\,c)$. But $\rho(a\,c)$ sends a to $\rho(c)$, while $(a\,c)\rho$ sends a to b, since $\rho(a) = b \neq c$. Thus $\rho(a) = \rho(c)$, which gives the contradiction $a = c$. □

2.4.7 Genus 0 Factorizations of Full Cycles

When applied in the genus 0 case with $\alpha = (n)$, Theorem 2.4.11 shows that the number of minimal transitive factorizations of full cycles in $\mathfrak{S}_n$ is equal to the number of vertex-labelled planar maps with one face. Such maps correspond with vertex- and edge-labelled trees, so we have $|\mathscr{C}_{(n)}|\, H_0((n)) = (n-1)!\, n^{n-2}$, or $H_0((n)) = n^{n-2}$. Of course, this is just a reiteration of Dénes' proof of Theorem 2.4.1. Note, however, that this derivation of $H_0((n))$ is not fully bijective, because the factor $(n-1)!$ introduced by edge-labelling must be eliminated by division. We now show how the argument can be modified to make it truly bijective.

Notice that Theorem 2.4.11 actually provides a bijection between minimal transitive factorizations of the full cycle $(1\,2\,\cdots\,n)$ and planar vertex- and edge-labelled trees whose lone descent cycle is $(1, 2, \ldots, n)^\circ$. All vertex labels save one can be stripped from such a tree without any loss of information, since the restriction on the descent cycle allows only one vertex-labelling once any

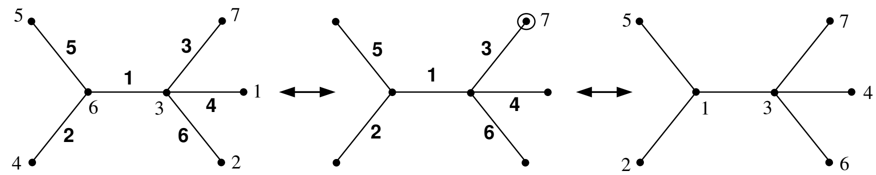

Figure 2.7: A bijection between factorizations of full cycles and labelled trees.

particular label has been assigned. Thus minimal factorizations of $(1\,2\,\cdots\,n)$ are in bijection with vertex-rooted, edge-labelled trees on n vertices, where the root vertex carries label n. Now observe that the edge labels of such a tree can be "pushed" away from the root and onto the vertices, in the sense that the label of an edge gets shifted to whichever of its endpoints is furthest from the root. This process results in a tree on n labelled vertices, and is clearly reversible.

This sequence of transformations gives a bijection between minimal factorizations of $(1\,2\,\cdots\,n)$ and trees on n labelled vertices. The correspondence is illustrated in Figure 2.7, starting with the factorization

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

The leftmost tree is the map of the factorization, and the other trees are obtained by first stripping vertex labels and then pushing edge-labels. The circled vertex in the central tree is its root. This bijection is equivalent to that given by Moszkowski in [54].

2.4.8 Genus 1 Factorizations of Full Cycles

We shall now use Theorem 2.4.11 to enumerate genus 1 factorizations of full cycles in $\mathfrak{S}_n$. This special case is substantially more complicated than that of minimal transitive factorizations treated in the previous section. The approach we take here, *viz.* pruning trees, will be substantially modified and generalized to all classes of factorizations in the next section. Our current description of the method is intended only as a preliminary to the more general case, and is accordingly abbreviated. To be succinct, we refer to genus 1 maps with one face as **one-maps** throughout our discussion. Also, all maps, graphs, and trees that we encounter are both vertex- and edge-labelled, unless otherwise specified.

Theorem 2.4.11 implies that we can count genus 1 factorizations of full cycles in $\mathfrak{S}_n$ by determining the number of one-maps on n vertices. Such maps have $r_1((n)) = n + 1$ edges, so their skeleton graphs are trees with two additional edges. These graphs can be *pruned* by first iteratively

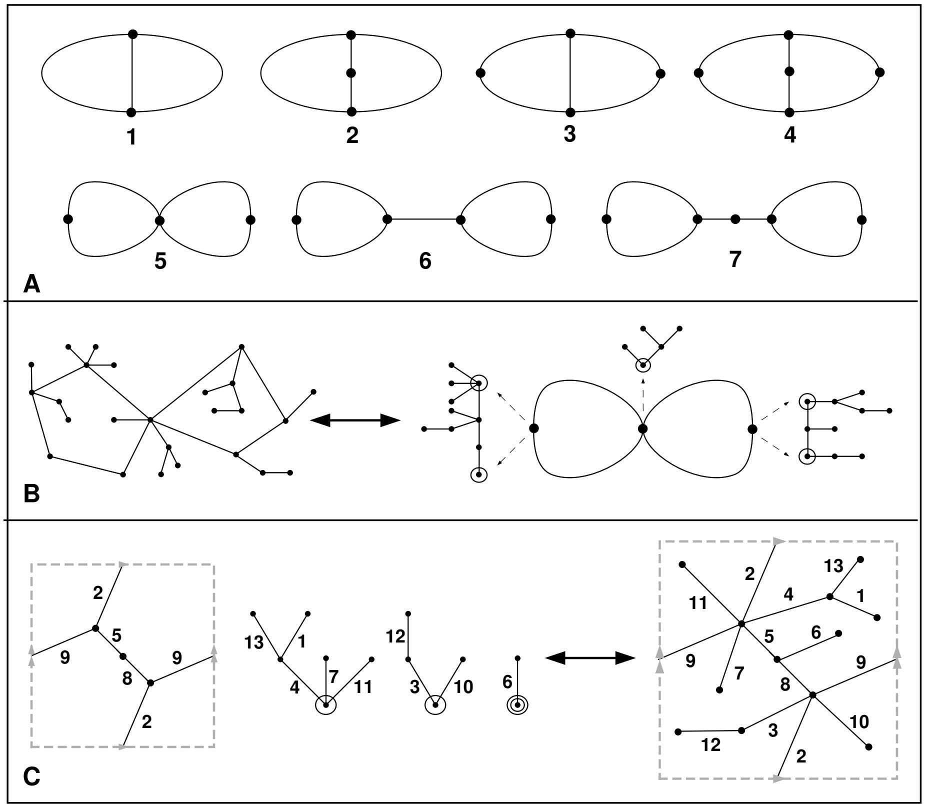

Figure 2.8: Pruning trees from graphs and maps.

removing vertices of degree one, and then contracting all edges joining vertices of degree two, provided such contractions do not result in loops. Through this process, the skeletons of one-maps can be categorized into the seven types depicted in Figure 2.8A.

Pruning is reversed by "replacing" each bivalent vertex with a doubly rooted tree and all other vertices with singly rooted trees. Replacing a vertex with a rooted tree is done in the obvious way, by identifying the vertex with the root of the tree. Replacement of a bivalent vertex v with a doubly rooted tree T is only slightly more involved: if v is incident with edges labelled i and j, where $i < j$, then v is first deleted, then the first root of T is attached to edge i, and finally the second root is attached to edge j. A schematic for pruning process and its reversal is given in Figure 2.8B. Labels have been suppressed in these diagrams for clarity.

Of course, one-maps can also be pruned. We call a one-map *irreducible* if it has no univalent vertices and if no edge connecting two bivalent vertices can be contracted without forming a loop. Thus pruning a one-map results in an irreducible one-map and a collection of rooted and doubly-

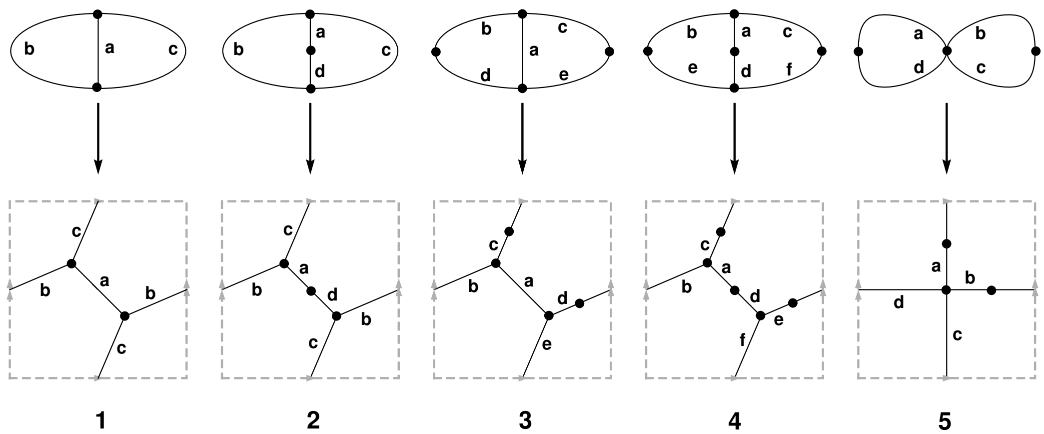

Figure 2.9: Classes of irreducible one-maps.

rooted trees. However, to reconstruct the original map from this data, the cyclic orderings of tree edges around the vertices of the irreducible map must be known. Fortunately, these orderings are completely specified by the increasing rotator condition. Thus the pruning process is reversible, as is demonstrated in Figure 2.8C. Clearly two one-maps are isomorphic if and only if the irreducible maps and trees obtained from each by pruning are isomorphic, with corresponding locations for attachment of the trees. Thus one-maps can be viewed as the composition of irreducible one-maps with trees.

Figure 2.9 illustrates the five distinct classes of irreducible one-maps, where we have used the standard polygonal representation of the torus. The skeletons of these maps are also shown for comparison with Figure 2.8A. (Note that graphs 6 and 7 of Figure 2.8A cannot be embedded on the torus to produce maps with one face.) For $i = 1, \ldots, 5$, let c_i be the number of (vertex- and edge-labelled) maps in class i. Then the generating series for one-maps with respect to labelled vertices, marked by x, and edges, marked by u, is

$$M(x,u) = c_1 \cdot \frac{w^2}{2!}\frac{u^3}{3!} + c_2 \cdot \frac{w^2 v}{3!}\frac{u^4}{4!} + c_3 \cdot \frac{w^2 v^2}{4!}\frac{u^5}{5!} + c_4 \cdot \frac{w^2 v^3}{5!}\frac{u^6}{6!} + c_5 \cdot \frac{w v^2}{3!}\frac{u^4}{4!}, \qquad (2.30)$$

where $w = w(x,u)$ is the tree series and $v = v(x,u)$ is the generating series for doubly-rooted trees. All series are exponential in x and u.

We now determine $c_1, \ldots, c_5$. To do so, we first hand-count all possible assignments of edge labels to the maps in Figure 2.9 such that rotators are increasing. We then divide by the appropriate number of automorphisms to obtain the true number of edge-labellings of each map. Finally, we use

Proposition 2.4.12 to deduce the number of distinct vertex-labellings of the resulting edge-labelled structures. These are the numbers $c_1, \ldots, c_5$. The symbols a, b, c, d, e and f used in our analysis are defined as in Figure 2.9.

Class 1: If a is the minimal label, then the order $a < b < c$ is fixed. A similar situation holds if either b or c is minimal. Thus there are 3 admissible labellings of the map. But there are also 3 automorphisms, corresponding to rotation of edges around the vertices, and therefore only $3/3 = 1$ edge-labelled map in this class. Since this map has 2 vertices, its two possible vertex labellings are equivalent. Thus we also have $c_1 = 1$.

Class 2: If a is minimal, then $a < b < c$ and $(b, c, d)^\circ$ must be increasing. Thus $a < d < b < c$ or $a < b < c < d$. A similar analysis holds if d is minimal. If b is minimal, then $b < c < a$ and $b < c < d$, giving only two possibilities, $b < c < a < d$ and $b < c < d < a$. The same holds if c is minimal. Thus there are 8 admissible labellings. There is only one nontrivial symmetry $(a \leftrightarrow d, c \leftrightarrow b)$, and hence $8/2 = 4$ inequivalent maps. There are 3 vertices, so Proposition 2.4.12 guarantees all 3! vertex-labellings are distinct. Thus $c_2 = 3! \cdot 4$.

Class 3: If a is minimal, then $a < b < c$ and $a < d < e$. There are $\binom{4}{2} = 6$ ways this can occur. If b is minimal, then $b < c < a$ and $(a, d, e)^\circ$ is increasing. A quick check shows 6 possibilities in this case, and the same is true if c, d or e is minimal. Thus there are $5 \cdot 6 = 30$ admissible labellings. There is only one nontrivial symmetry $(b \leftrightarrow d, c \leftrightarrow e)$, so there are $30/2 = 15$ inequivalent maps in this class. Hence $c_3 = 4! \cdot 15$.

Class 4: If a is minimal, then $a < b < c$ and $(d, e, f)^\circ$ is increasing. There are $3\binom{5}{2} = 30$ ways this can occur, and the same is true if b, c, d, e or f is minimal. Thus there are $6 \cdot 30 = 180$ admissible labellings. There are $3 \cdot 2 = 6$ automorphisms, obtained through all compositions of rotation around one vertex and the exchange $(a \leftrightarrow d,\ b \leftrightarrow e,\ c \leftrightarrow f)$. Hence there are $180/6 = 30$ inequivalent maps in this class, and $c_4 = 5! \cdot 30$.

Class 5: If a is minimal, then the order $a < b < c < d$ is fixed. The same holds if b, c, or d is minimal. Thus there are 4 admissible labellings of the map. There are 4 automorphisms (rotations around the central vertex), so there is only $\frac{4}{4} = 1$ edge-labelled map in this class. Thus $c_5 = 3! \cdot 1$.

Using (2.18) to write $v = w/(1 - uw)$, we can now simplify (2.30) to obtain

$$M(x, u) = \frac{u^3 w^2 (2 - uw)}{24(1 - uw)^3}. \tag{2.31}$$

By definition, $n!(n+1)!\,[x^n u^{n+1}]\,M(x,u)$ is the number of one-maps on n vertices. Equivalently, this is the number of genus 1 factorizations of full cycles in $\mathfrak{S}_n$, which is $(n-1)!H_1((n))$ since $\mathscr{C}_{(n)} = (n-1)!$. That is,

$$M(x,u) = \frac{1}{24}\sum_{n\geq 1}\frac{H_1((n))}{n(n+1)!}x^n u^{n+1}.$$

One can now apply Lagrange inversion to expand $M(x,1)$ as a series in x, thereby evaluating $H_1((n))$. Alternatively, implicit differentiation of (2.18) shows that

$$x^2\frac{d^2w}{dx^2} = \frac{uw^2(2-uw)}{(1-uw)^3}.$$

Together with (2.17) and (2.31), this gives

$$M(x,u) = \frac{1}{24}u^2x^2\frac{d^2w}{dx^2} = \frac{1}{24}\sum_{n\geq 1}\frac{n^{n-1}}{(n-2)!}u^{n+1}x^n.$$

Equating coefficients now completes the proof of the following result, which is seen to be in agreement with Theorem 2.3.4.

Theorem 2.4.16. *For any $n \geq 1$, we have $H_1((n)) = \frac{1}{24}n^{n+1}(n^2-1)$.* □

We conclude by drawing attention to the fact that

$$M(x,u) = \frac{1}{4!}u^2x^2\frac{d^2w}{dx^2}.$$

The series on the right counts doubly vertex-rooted and singly edge-rooted trees with an additional two edges adjoined (where both vertices and edges are labelled) up to a factor of 4!, which possibly accounts for some symmetry. This is probably coincidental, as such pleasant forms are not apparent for higher genus one-face maps, but perhaps there is a combinatorial construction based on these observations that bypasses the case-analytic path we have followed.

2.4.9 Face-Labelled Maps

The examples of the previous two sections demonstrate how the connection between maps and factorizations can be gainfully applied to the study of factorizations. We now develop a variation of the bijection MAP that links factorizations to face-labelled maps.

We first require some new terminology. The face labels of a face-labelled map naturally induce an ordering on the parts of its descent partition. It is natural, then, to let the resulting composition

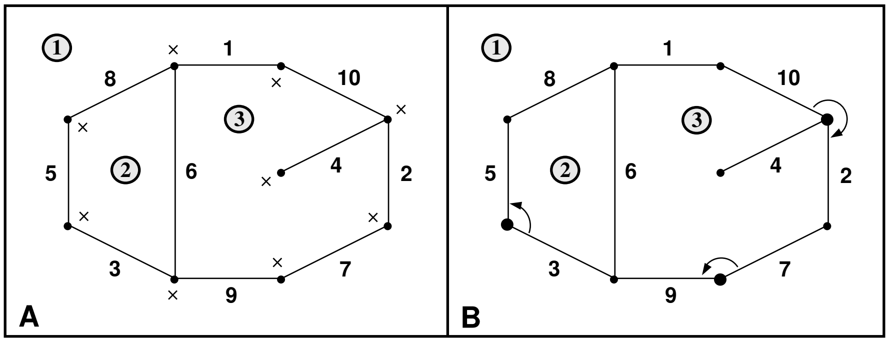

Figure 2.10: A face-labelled map of descent class $(3, 2, 4) \models 9$.

represent the descent structure of such a map.

Definition 2.4.17. *Let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be an m-part composition. A face-labelled map with m faces is said to be of **descent class** α if face s has exactly α_s descents, for $1 \leq s \leq m$.*

For example, Figure 2.10A shows a face-labelled map of descent class $(3, 2, 4) \models 9$, with crosses placed at descent corners.

The next theorem is a corollary of Theorem 2.4.11 and the fact that vertex-labelled maps have no nontrivial automorphisms. It gives an alternative interpretation of the Hurwitz numbers $H_g(\alpha)$ in terms of face-labelled maps with distinguished descents.

Theorem 2.4.18. *Let α be a composition and fix $\pi \in \mathscr{C}_\alpha$. Then there is a bijection between genus g factorizations of π and face-labelled genus g maps of descent class α in which one descent of each face has been distinguished.*

Proof. Suppose α has m parts. Let $\pi_1, \ldots, \pi_m$ be the cycles of π, and let p_i be the minimal symbol of π_i. Without loss of generality, assume that the cycles π_i have been indexed so that $p_1 \leq \cdots \leq p_m$. Let f be a genus g factorization of π. By Theorem 2.4.11, $\mathscr{M}_f$ is a genus g vertex-labelled map with m faces with descent cycles $\pi_1, \ldots, \pi_m$. For $1 \leq s \leq m$, assign label s to the face of $\mathscr{M}_f$ having descent cycle π_s. This yields a face-labelled map of descent class α. Distinguish the vertices with labels $p_1, \ldots, p_m$ in some way, and then strip all vertex labels from this map. This transformation is reversible, since the locations of all labels are uniquely determined by the descent cycles from the locations of $p_1, \ldots, p_m$. We therefore obtain a face-labelled genus g map of descent class α, with one descent of each face distinguished. Clearly any such map can be constructed in this way and, since vertex-labelled maps have no nontrivial automorphisms, two different factorizations never lead to the same maps. $\square$

For example, the bijection of Theorem 2.4.18 associates the map drawn in Figure 2.10B with the factorization

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

In fact, we have already made use of a special case of Theorem 2.4.18. In genus 0 with $\pi = (1\,2\,\cdots\,n)$, it was the basis of the bijective proof of Theorem 2.4.1 given in §2.4.7. It will be used again in §2.8 as a basis for further bijections of a similar nature.

2.4.10 Properly Labelled Maps

In Theorem 2.4.11 we established the bijection MAP between factorizations and vertex-labelled maps, and in Theorem 2.4.18 we described a close relative of this bijection that connects factorizations with certain face-labelled maps. In this section, we consider another modification of MAP, this one associating factorizations with vertex- and face-labelled maps. Although these correspondences are extremely similar, each provides a slightly different representation of factorizations which is particularly well suited for certain applications. Theorem 2.4.21, below, will be convenient when we extend the method of pruning trees introduced in §2.4.8 in to arbitrary factorizations.

With any composition $\alpha = (\alpha_1, \ldots, \alpha_m) \models n$ we associate a sequence $\mathbb{D}_1(\alpha), \ldots, \mathbb{D}_m(\alpha)$ of subsets $[n]$, defined as follows:

$$\mathbb{D}_s(\alpha) = \{\alpha_1 + \cdots + \alpha_{s-1} + 1, \quad \ldots, \quad \alpha_1 + \cdots + \alpha_s\}, \qquad \text{for } 1 \le s \le m.$$

We refer to these sets as the **canonical descent sets** associated with α. For example, if $\alpha = (3, 2, 4)$ then its associated canonical descent sets are $\mathbb{D}_1(\alpha) = \{1, 2, 3\}$, $\mathbb{D}_2(\alpha) = \{4, 5\}$ and $\mathbb{D}_3(\alpha) = \{6, 7, 8, 9\}$. We write $\mathfrak{S}(\alpha)$ for the set of all permutations π such that $\operatorname{orb} \pi = \{\mathbb{D}_1(\alpha), \ldots, \mathbb{D}_m(\alpha)\}$. That is, $\mathfrak{S}(\alpha)$ contains the $\prod_i (\alpha_i - 1)!$ permutations whose cycles are supported by the canonical descent sets associated with α.

Definition 2.4.19. *A vertex- and face-labelled map is said to be **properly labelled** if it is of descent class $\alpha = (\alpha_1, \ldots, \alpha_m)$ and if face s has descent set $\mathbb{D}_s(\alpha)$, for $1 \le s \le m$.*

Example 2.4.20. Figure 2.11 shows a properly labelled map of descent class $\alpha = (10, 3, 2) \models 15$. Its descent sets are the canonical descent sets of α, namely

$$\mathbb{D}_1(\alpha) = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}, \quad \mathbb{D}_2(\alpha) = \{11, 12, 13\}, \quad \text{and} \quad \mathbb{D}_3(\alpha) = \{14, 15\}.$$

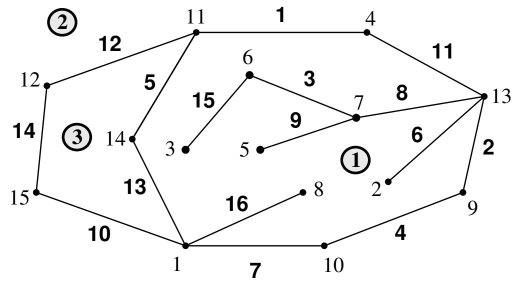

Figure 2.11: A properly-labelled map of descent class $(10, 3, 2) \models 15$.

Stripped of its face labels, this is actually the map of a factorization of

$$\pi = (1\ 10\ 9\ 2\ 6\ 3\ 7\ 5\ 4\ 8)(13\ 12\ 11)(14\ 15),$$

which is one of the $9!\,2!\,1!$ members of $\mathfrak{S}(\alpha)$. □

The relationship between factorizations and properly labelled maps is formalized in the next theorem. Its proof is nearly identical to that of Theorem 2.4.18, and is not given in full detail.

Theorem 2.4.21. *Let α be a composition. There is a bijection between genus g factorizations of permutations in $\mathfrak{S}(\alpha)$ and properly labelled genus g maps of descent class α.*

Proof. Let α have m parts. If f is a factorization of $\pi \in \mathfrak{S}(\alpha)$, then the descent cycles of $\mathscr{M}_f$ are supported by the canonical descent sets $\mathbb{D}_1(\alpha), \ldots, \mathbb{D}_m(\alpha)$. Assigning label s to the face of $\mathscr{M}_f$ having descent set $\mathbb{D}_s(\alpha)$ yields the properly labelled map corresponding to f. □

Corollary 2.4.22. *There are*

$$H_g(\alpha) \prod_i (\alpha_i - 1)!$$

properly labelled genus g maps of descent class $\alpha = (\alpha_1, \ldots, \alpha_m)$.

Proof. This follows immediately from the theorem since $|\mathfrak{S}(\alpha)| = \prod_i (\alpha_i - 1)!$. □

Thus the number of properly labelled genus g maps of descent class α is a simple scaling of the Hurwitz number $H_g(\alpha)$. Much of the remainder of this chapter is devoted to the study of the generating series for such maps, which is defined as follows.

Definition 2.4.23. *We write $M_g(\alpha)$ for the number of properly labelled genus g maps of descent class α. For fixed $m \geq 1$, set $\mathbf{x} = (x_1, \ldots, x_m)$ and let*

$$\Psi_m^{(g)}(\mathbf{x}, u) = \sum_{n \geq 1} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} M_g(\alpha) \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \frac{u^{r_g(\alpha)}}{r_g(\alpha)!}, \tag{2.32}$$

be the generating series for the numbers $\{M_g(\alpha) : \ell(\alpha) = m\}$. When considering the genus 0 series, we often write Ψ_m in place of $\Psi_m^{(0)}$.

Notice that $x_1, \ldots, x_m$ and u are naturally exponential indeterminates in (2.32), with x_i marking vertices at descents of face i of a properly labelled map (these are labelled with the i-th canonical descent set), and u marking labelled edges. Throughout this chapter, the symbol $\mathbf{x}$ will always represent the vector $(x_1, \ldots, x_m)$, where m is understood from context.

The apparent clash of notation between the definition of $\Psi_m^{(g)}$ given here and the one presented in §2.3.5 is resolved by Corollary 2.4.22, since the identity $M_g(\alpha) = H_g(\alpha) \prod_i (\alpha_i - 1)!$ shows the series (2.32) for properly labelled maps to be equal to the symmetrized Hurwitz series (2.20). The combinatorial effect of the operator (2.19), which transformed the Hurwitz series $\Phi^{(g)}$ into the symmetrized series $\Psi_m^{(g)}$ (see §2.3.5), is seen to be that of applying face labels to the maps of factorizations.

We remark that we have actually already evaluated $\Psi_m^{(g)}(\mathbf{x}, u)$ in two special cases. In particular, Theorem 2.4.1 implies

$$x \frac{d}{dx} \Psi_1^{(0)}(x, u) = w, \tag{2.33}$$

while (2.31) asserts that

$$\Psi_1^{(1)}(x, u) = \frac{u^3 w^2 (2 - uw)}{24(1 - uw)^3}.$$

Note the dependence of both expressions on the tree series w.

2.4.11 Comments on Labelling

Proposition 2.4.12 implies that all vertex-labellings of a map are distinct (*i.e.* result in nonisomorphic vertex-labelled maps) unless the map has exactly two vertices. Thus the number of properly labelled maps should be easily obtained from the number of face-labelled maps. The following technical results make this notion precise.

Proposition 2.4.24. *A face-labelled map with more than two vertices or more than one face has no nontrivial isomorphisms.*

Proof. By Proposition 2.4.12, we need only consider face-labelled maps with exactly two vertices and at least two faces. Let $\mathscr{M}$ be such a map, say of genus g with vertices u and v. Proposition 2.4.14 implies $\mathscr{M}$ has two faces and $2g+2$ edges. Moreover, the increasing rotator condition forces these faces to have boundary walks

$$((u,1),(v,2),\ldots,(u,2g+1),(v,2g+2))^{\circ} \quad \text{and} \quad ((v,1),(u,2),\ldots,(v,2g+1),(u,2g+2))^{\circ}.$$

Switching u and v therefore interchanges the labelled faces of $\mathscr{M}$, resulting in a distinct map. □

Corollary 2.4.25. *Let* $\alpha = (\alpha_1,\ldots,\alpha_m) \neq (2)$. *Then there are* $M_g(\alpha)/(\alpha_1!\cdots\alpha_m!)$ *genus* g *face-labelled maps of descent class* α.

Proof. Let $\mathscr{M}$ be a genus g face-labelled map of descent class α. Propositions 2.4.12 and 2.4.24 imply that the vertices of face i can be labelled with $\mathbb{D}_i(\alpha)$ in $\alpha_i!$ distinct ways. Doing so for each face results in a properly labelled map, and the result follows. □

The corollary indicates that $\Psi_m^{(g)}(\mathbf{x},u)$ can widely be regarded as the generating series for genus g maps with m labelled faces, where x_i is an *ordinary* marker for descents in face i, and u is an exponential marker for labelled edges. The sole exception occurs for maps with only one face and two vertices. These maps correspond to the x^2 term of $\Psi_1^{(g)}(x,u)$, where we have

$$\left[x^2\frac{u^{2g+1}}{(2g+1)!}\right]\Psi_1^{(g)}(x,u) = \frac{1}{2}.$$

In what follows, we shall often ignore this anomaly and interpret $\Psi_m^{(g)}$ as the series for face-labelled maps rather than properly labelled maps. This usually has the effect of simplifying our combinatorial manipulations, since we need not worry about preserving vertex labels. We adopt this alternative interpretation of $\Psi_m^{(g)}$ only when $m \neq 1$, or when a differential operator such as $x\partial/\partial x$ is being applied to $\Psi_1^{(g)}(x,u)$. In the latter case, note that the series $(x\partial/\partial x)\Psi_1^{(g)}(x,u)$ does faithfully count one-face genus g maps in which one vertex has been distinguished.

2.4.12 Additional Notes

A number of authors have given bijective proofs of Theorem 2.4.1. Moszkowski [54] was the first among these, but see also [38], [39], and [57]. Both [39] and [57] contain an alternative description of Moszkowski's bijection like the one presented in §2.4.7. A different, but related, correspondence between trees and factorizations of full cycles also appears in [39]. We shall encounter a generalization of this bijection in §3.4.6. The graph of a general factorization is considered in [4].

Arnol'd [2] is often credited with being the first to assign "map-like" properties to the graph of a factorization, though this seems somewhat generous. Through analytic methods he determines $H_0((p,q))$, and then, as a corollary, he makes the corresponding graph-theoretic claim regarding the number of vertex- and edge-labelled graphs with the property that the product of the transpositions induced by the edges is equal to a permutation of cycle type (p,q). Arnol'd refers to the graph of a factorization as a *monodromy graph*.

The link established by Theorem 2.4.11 between factorizations and maps with certain descent structure also appears, independently, in [56] and [57]. Poulalhon's description [57] is essentially identical to Theorem 2.4.11, whereas in [56] the correspondence arises from geometrical considerations and is presented in different form.

2.5 Differential Equations for Labelled Maps

In this section we investigate a differential decomposition for face-labelled maps, and show how it provides an algebraic rationale for the dependence of the symmetrized Hurwitz series on the tree series. Throughout, we avoid vertex labellings altogether and regard $\Psi_m^{(g)}$ as the generating series for genus g maps with m labelled faces. (See §2.4.11 for comments on labelling.)

2.5.1 Decomposition of Planar Maps

Theorem 2.5.1, below, gives a differential equation satisfied by the generating series $\Psi_m(\mathbf{x},u)$ for planar, properly labelled maps. It first appeared in [33], in a rough form, and then again in [36] in a form identical to that given here. The proof offered in both cases is algebraic, consisting of an analysis of the action of the symmetrization operator (2.19) on the cut-join equation (2.13). The proof we give here relies on a decomposition for planar face-labelled maps.

Theorem 2.5.1. *Fix $m \geq 2$. For any subset $\lambda = \{\lambda_1,\dots,\lambda_k\} \subseteq [m]$, where $\lambda_1 < \cdots < \lambda_k$, let $\mathbf{x}_\lambda = (x_{\lambda_1},\dots,x_{\lambda_k})$. For $1 \leq i \leq m$, let $\overline{\mathbf{x}}_i = \mathbf{x}_{[m]\setminus\{i\}}$. Also, for each i, let ∂_i denote the operator $x_i\partial/\partial x_i$, and let $\mathcal{P}_i$ be the collection of all pairs $\{\gamma,\lambda\}$ of subsets of $[m]$ such that $\gamma\cap\lambda=\{i\}$ and $\gamma\cup\lambda=[m]$. Then*

$$\frac{\partial}{\partial u}\Psi_m(\mathbf{x}, u) = \sum_{i=1}^{m} \sum_{\{\gamma,\lambda\}\in\mathcal{P}_i} \partial_i \Psi_{|\gamma|}(\mathbf{x}_\gamma, u) \cdot \partial_i \Psi_{|\lambda|}(\mathbf{x}_\lambda, u) + \sum_{1\le i<j\le m} \frac{x_j \partial_i \Psi_{m-1}(\overline{\mathbf{x}}_j, u) - x_i \partial_j \Psi_{m-1}(\overline{\mathbf{x}}_i, u)}{x_i - x_j}. \quad (2.34)$$

Proof. The series on the left-hand side of (2.34) counts all possible structures $\mathscr{M}\backslash e$, where $\mathscr{M}$ is a face-labelled planar map with m faces and e is its maximal edge. The enumeration is with respect to labelled edges of $\mathscr{M}\backslash e$ and the descent class of $\mathscr{M}$. We show that the series on the right counts these same objects. To this end, let $\mathscr{M}$ be a face-labelled planar map with m faces, and let $e = \{a, b\}$ be its maximal edge. Consider the effect of deleting e from $\mathscr{M}$.

Suppose first that e is incident with only one face F of $\mathscr{M}$, labelled i. As shown below, deletion of e separates $\mathscr{M}$ into two planar maps, $\mathscr{M}_a$ and $\mathscr{M}_b$, containing vertices a and b, respectively. These maps inherit labels from $\mathscr{M}$ in the obvious way.

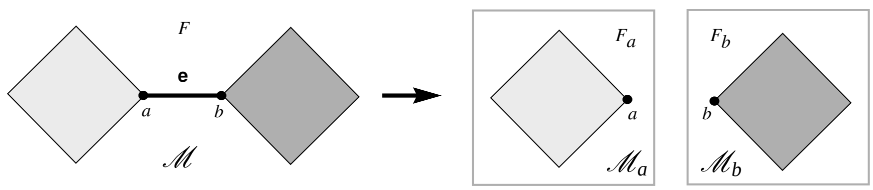

The faces of $\mathscr{M}_a$ and $\mathscr{M}_b$ are labelled with $\gamma \subseteq [m]$ and $\lambda \subseteq [m]$, respectively, where $\gamma \cap \lambda = \{i\}$ and $\gamma \cup \lambda = [m]$. Let F_a and F_b be the faces of $\mathscr{M}_a$ and $\mathscr{M}_b$ with label i. Since e is maximal, both a and b are at descents of F. Thus a is at a descent of F_a, and b is at a descent of F_b. The series counting all pairs $(\mathscr{M}_a, \mathscr{M}_b)$ is therefore

$$\partial_i \Psi_{|\gamma|}(\mathbf{x}_\gamma, u) \cdot \partial_i \Psi_{|\lambda|}(\mathbf{x}_\lambda, u)$$

where the operator ∂_i has the effect of distinguishing vertices a and b at descents of F_a and F_b. Summing over i and over permissible pairs $\{\lambda, \gamma\}$ gives the first summation on the right-hand side of (2.34).

Now suppose e is incident with two distinct faces F_i and F_j of $\mathscr{M}$, labelled i and j, respectively. Since e is maximal, a is at a descent of one of these faces, and b is at a descent of the other. Without loss of generality, assume a is at a descent of F_i. Deletion of e creates a new map $\mathscr{M}_0$ by fusing F_i and F_j into a single face F_0, which we label 0. All other faces and edges of $\mathscr{M}_0$ inherit labels from $\mathscr{M}$. The deletion of e is illustrated below.

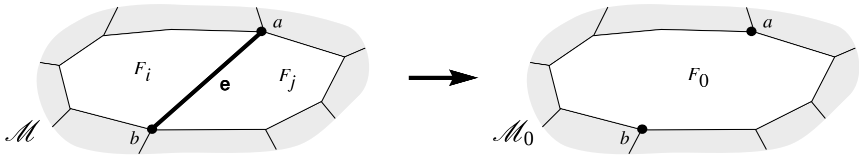

Observe that a vertex is at a descent of F_0 if and only if it is at a descent of F_i or F_j. If F_0 has n descents, faces F_i and F_j therefore have d and $n-d$ descents, respectively, for some d with $1 \leq d \leq n-1$. Moreover, b is uniquely determined by the location of a and this value of d.

Let $\overline{\mathbf{x}}_{ij} = \mathbf{x}_{[m]\setminus\{i,j\}}$, and define

$$G(x, \overline{\mathbf{x}}_{ij}, u) = x \frac{\partial}{\partial x} \Psi_{m-1}(x, \overline{\mathbf{x}}_{ij}, u).$$

Regarding x as a marker for descents of face 0, this series counts maps $\mathscr{M}_0$ with $m-1$ faces labelled $\{0, \ldots, m\}\setminus\{i, j\}$ in which a vertex a at a descent of face 0 has been distinguished. From the considerations above, the series counting all possible structures $\mathscr{M}\setminus e$ is therefore obtained from G by replacing x^n with $\sum_{d=1}^{n-1} x_i^d x_j^{n-d}$. By (1.1) and Lemma 1.3.3, this yields

$$G(x, \overline{\mathbf{x}}_{ij}, u) \circ \Delta^{+}(x\,;\, x_i, x_j) = \frac{x_j G(x_i, \overline{\mathbf{x}}_{ij}, u) - x_i G(x_j, \overline{\mathbf{x}}_{ij}, u)}{x_i - x_j}.$$

But, since Ψ_{m-1} is symmetric, $G(x_i, \overline{\mathbf{x}}_{ij}, u) = \partial_i \Psi_{m-1}(\overline{\mathbf{x}}_j, u)$ and $G(x_j, \overline{\mathbf{x}}_{ij}, u) = \partial_j \Psi_{m-1}(\overline{\mathbf{x}}_i, u)$. Summing over all pairs $\{i, j\} \subseteq [m]$ gives the second summation on the right-hand side of (2.34). □

Corollary 2.5.2. *With the same notation as in Theorem 2.5.1, we have*

$$\left(\frac{\partial}{\partial u} - \sum_{i=1}^{m} w_i \partial_i\right) \Psi_m(\mathbf{x}, u) = \sum_{i=1}^{m} \sum_{\substack{\{\gamma,\lambda\}\in\mathcal{P}_i \\ |\gamma|,|\lambda|\geq 2}} \partial_i \Psi_{|\gamma|}(\mathbf{x}_\gamma, u) \cdot \partial_i \Psi_{|\lambda|}(\mathbf{x}_\lambda, u) + \sum_{\substack{1\leq i,j\leq m \\ i\neq j}} \frac{x_j \partial_i \Psi_{m-1}(\overline{\mathbf{x}}_j, u)}{x_i - x_j}.$$

Proof. If $\{\gamma, \lambda\} \in \mathcal{P}_i$ and $|\gamma| = 1$, then $\gamma = \{i\}$ for some $i \in [m]$ and $\lambda = [m]$. Hence, by (2.33), $\partial_i \Psi_{|\gamma|}(\mathbf{x}_\gamma, u) = \partial_i \Psi_1(x_i, u) = w_i$, and $\partial_i \Psi_{|\lambda|}(\mathbf{x}_\lambda, u) = \partial_i \Psi_m(\mathbf{x}, u)$. The result follows immediately upon rearranging (2.34). □

Both [33] and [36] also give a differential equation satisfied by the series $\Psi_m^{(g)}(\mathbf{x}, u)$ with positive genus g. The restriction here to planar maps is intended only to simplify our presentation. It is straightforward, though not particularly enlightening, to modify the proof of Theorem 2.5.1 to obtain decompositions for maps of any genus.

2.5.2 A Change of Variables

The significance of the seemingly obscure differential operator on the left-hand side of Corollary 2.5.2 will now be explained. For this purpose, we momentarily regard $w_1, \ldots, w_m$ as algebraically independent indeterminates, forgetting the usual definition of these symbols as tree series. Then, following [36], we change variables by substituting

$$x_i = w_i e^{-uw_i} \in \mathbb{Q}[u][[w_i]] \tag{2.35}$$

for each occurrence of x_i in $\Psi_m(\mathbf{x}, u)$. That is, we introduce the series

$$\Gamma_m(\mathbf{w}, u) = \Psi_m(w_1 e^{-uw_1}, \ldots, w_m e^{-uw_m}, u) \in \mathbb{Q}[u][[\mathbf{w}]], \tag{2.36}$$

where $\mathbf{w} = (w_1, \ldots, w_m)$. Of course, the substitution (2.35) can be inverted to identify w_i as a series in x_i and u. In particular, we have $w_i = x_i e^{uw_i}$. Comparing with (2.16), we see that w_i is indeed the tree series $w(x_i, u) \in \mathbb{Q}[u][[x_i]]$, which explains our choice of notation.

Having described the change of variables (2.35) and its inverse, we may now pass freely between the rings $\mathbb{Q}[u][[\mathbf{x}]]$ and $\mathbb{Q}[u][[\mathbf{w}]]$. For instance, we can rewrite (2.36) as

$$\Gamma_m(\mathbf{w}, u) = \Psi_m(\mathbf{x}, u), \tag{2.37}$$

where both sides are to be interpreted either as series in the independent variables $\mathbf{w}$ and u, or as series in the independent variables $\mathbf{x}$ and u. Under the former interpretation, differentiating (2.37) with the chain rule gives

$$\begin{aligned}
\frac{\partial}{\partial u}\Gamma_m(\mathbf{w}, u) &= (D_{m+1}\Psi_m)(\mathbf{x}, u) \cdot \frac{\partial u}{\partial u} + \sum_{i=1}^{m} (D_i \Psi_m)(\mathbf{x}, u) \cdot \frac{\partial x_i}{\partial u} \\
&= (D_{m+1}\Psi_m)(\mathbf{x}, u) - \sum_{i=1}^{m} (D_i \Psi_m)(\mathbf{x}, u) \cdot w_i x_i \\
&= \left(D_{m+1} - \sum_{i=1}^{m} w_i x_i D_i \right) \Psi_m(\mathbf{x}, u),
\end{aligned}$$

where D_i represents differentiation with respect to the i-th argument. Note that we have used (2.35) to evaluate $\partial x_i / \partial u = -w_i^2 e^{uw_i} = -w_i x_i$. The expression above can be rewritten as

$$\frac{\partial}{\partial u}\Gamma_m(\mathbf{w}, u) = \left(\frac{\partial}{\partial u} - \sum_{i=1}^{m} w_i x_i \frac{\partial}{\partial x_i} \right) \Psi_m(\mathbf{x}, u), \tag{2.38}$$

though we caution that the operator $\partial/\partial u$ has different meanings on the left- and right-hand sides of this equation. In particular, both the w_i on the left and the x_i on the right are to be regarded as constants (independent of u) for the purposes of this operator.

Equation (2.38) shows that the differential operator of Corollary 2.5.2 has a pleasant form in terms of the tree series w_i. In fact, we can invert the operator to obtain the following recursive expression for $\Gamma_m(\mathbf{w}, u)$.

Theorem 2.5.3. *Fix $m \geq 2$. For any subset $\lambda = \{\lambda_1, \ldots, \lambda_k\} \subseteq [m]$, where $\lambda_1 < \cdots < \lambda_k$, define $\mathbf{w}_\lambda = (w_{\lambda_1}, \ldots, w_{\lambda_k})$. For each $i \in [m]$, set $\overline{\mathbf{w}}_i = \mathbf{w}_{[m]\setminus\{i\}}$. Also, for each i, let*

$$\partial_i = \frac{w_i}{1 - uw_i} \frac{\partial}{\partial w_i}$$

and let $\mathcal{P}_i$ be the set of all pairs $\{\gamma, \lambda\}$ with $\gamma, \lambda \subset [m]$ such that $\gamma \cap \lambda = \{i\}$ and $\gamma \cup \lambda = [m]$. Then

$$\Gamma_m(\mathbf{w}, u) = \sum_{i=1}^{m} \sum_{\substack{\{\gamma,\lambda\}\in\mathcal{P}_i \\ |\gamma|,|\lambda|\geq 2}} \int \partial_i \Gamma_{|\gamma|}(\mathbf{w}_\gamma, u) \cdot \partial_i \Gamma_{|\lambda|}(\mathbf{w}_\lambda, u)\, du + \sum_{\substack{1\leq i,j\leq m \\ i\neq j}} \int \frac{w_j e^{-uw_j} \partial_i \Gamma_{m-1}(\overline{\mathbf{w}}_j, u)}{w_i e^{-uw_i} - w_j e^{-uw_j}}\, du,$$

where $w_1, \ldots, w_m$ are considered to be constants independent of u in the integrations.

Proof. This follows immediately from Corollary 2.5.2 and equations (2.38), (2.35), and (2.21). □

Corollary 2.5.4.

$$\left(x_1 \frac{\partial}{\partial x_1} + x_2 \frac{\partial}{\partial x_2} \right) \Psi_2(x_1, x_2, u) = \frac{u^2 w_1 w_2}{(1 - uw_1)(1 - uw_2)}.$$

Proof. Directly applying Theorem 2.5.3 in the case $m = 2$ produces

$$\Gamma_2(w_1, w_2, u) = \int \frac{w_2 e^{-uw_2} \partial_1 \Gamma_1(w_1, u) - w_1 e^{-uw_1} \partial_2 \Gamma_1(w_2, u)}{w_1 e^{-uw_1} - w_2 e^{-uw_2}}\, du.$$

From (2.21) and (2.33) we have $\partial_i \Gamma_1(w_i, u) = w_i$, and thus

$$\begin{aligned}
\Gamma_2(w_1, w_2, u) &= \int \frac{w_1 w_2 e^{-uw_2} - w_1 w_2 e^{-uw_1}}{w_1 e^{-uw_1} - w_2 e^{-uw_2}}\, du \\
&= \int \left(\frac{w_1^2 e^{-uw_1} - w_2^2 e^{-uw_2}}{w_1 e^{-uw_1} - w_2 e^{-uw_2}} - (w_1 + w_2) \right) du \\
&= \log\left(\frac{w_1 - w_2}{w_1 e^{-uw_1} - w_2 e^{-uw_2}} \right) - u(w_1 + w_2). \qquad (2.39)
\end{aligned}$$

Therefore (2.35) and (2.37) yield

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

Using (2.21), this gives

$$x_1 \frac{\partial}{\partial x_1}\Psi_2(x_1, x_2, u) = \frac{w_2}{(w_1 - w_2)(1 - uw_1)} - \frac{x_2}{x_1 - x_2} \tag{2.40}$$
$$x_2 \frac{\partial}{\partial x_2}\Psi_2(x_1, x_2, u) = \frac{w_1}{(w_2 - w_1)(1 - uw_2)} - \frac{x_1}{x_2 - x_1},$$

from which the result follows. □

Corollary 2.5.5.

$$\Psi_3(x_1, x_2, x_3, u) = \frac{u^4 w_1 w_2 w_3}{(1 - uw_1)(1 - uw_2)(1 - uw_3)}.$$

Proof. This follows from Theorem 2.5.3 and (2.40). Details can be found in [36]. □

Of course, the previous two corollaries are seen to be in agreement with Theorem 2.3.9 and, in general, the recursive formula of Theorem 2.5.3 can be applied (as above) to compute closed form expressions for $\Psi_m(\mathbf{x}, u)$ for any $m \geq 2$. However, it is not known how to obtain Theorem 2.3.9 through this method. In fact, it is not even clear from the recurrence that the series $\Psi_m(\mathbf{x}, u)$ is rational in $w_1, \ldots, w_m$. This last point, at least, is cleared up by the following simplification of Theorem 2.5.3 in the case $m \geq 4$.

Theorem 2.5.6. *Fix $m \geq 4$. With the same notation as in Theorem 2.5.3, we have*

$$\Gamma_m(\mathbf{w}, u) = \sum_{i=1}^{m} \sum_{\substack{\{\gamma,\lambda\}\in\mathcal{P}_i \\ |\gamma|,|\lambda|\geq 3}} \int \partial_i \Gamma_{|\gamma|}(\mathbf{w}_\gamma, u)\cdot\partial_i \Gamma_{|\lambda|}(\mathbf{w}_\lambda, u)\, du + \sum_{\substack{1\leq i,j\leq m \\ i\neq j}} \int \frac{w_j \partial_i \Gamma_{m-1}(\overline{\mathbf{w}}_j, u)}{(1 - uw_i)(w_i - w_j)}\, du.$$

Proof. Suppose $\{\gamma, \lambda\} \in \mathcal{P}_i$ with $|\gamma| = 2$. Then $\gamma = \{i, j\}$ and $\lambda = [m]\backslash\{j\}$ for some $i \neq j$. We therefore have

$$x_i \frac{\partial}{\partial x_i}\Psi_{|\gamma|}(\mathbf{x}_\gamma, u) = \frac{w_j}{(w_i - w_j)(1 - uw_i)} - \frac{x_j}{x_i - x_j},$$
$$x_i \frac{\partial}{\partial x_i}\Psi_{|\lambda|}(\mathbf{x}_\lambda, u) = x_i \frac{\partial}{\partial x_i}\Psi_{m-1}(\overline{\mathbf{x}}_j, u),$$

where the first equation comes from (2.40). Upon substituting these expressions in the differential

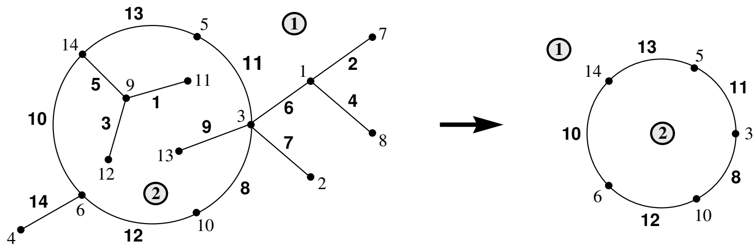

Figure 2.12: A properly labelled map and its core.

equation of Corollary 2.5.2, all terms with denominator $x_i - x_j$ cancel. Integrating the resulting equation with respect to u completes the proof. □

Together with Corollary 2.5.5, this theorem demonstrates that $\Psi_m(\mathbf{x}, u)$ is a rational function of $w_1, \ldots, w_m$. However, the combinatorial rationale for this dependence on the tree series is unclear, since the combinatorics of Theorem 2.5.1 is lost in the algebraic contortions used to deduce Theorem 2.5.6. We now abandon this algebraic approach and return to the combinatorics of properly labelled maps.

2.6 Smooth Maps and Pruning Trees

In §2.4.8, we found that *pruning trees* was a key step toward the enumeration of one-face maps on the torus. In this section we consider the extension of this method to arbitrary factorizations.

2.6.1 Cores and Branches

A **leaf** of a map is a vertex of degree one, and a map is **smooth** if it has no leaves. Iteratively removing leaves (and their incident edges) from a map clearly results in a smooth map of the same genus. Moreover, if the original map is not a plane tree (*i.e.* one-face planar map) then the smooth map obtained in this way is unique. We call the map resulting from the reduction of $\mathscr{M}$ the **core** of $\mathscr{M}$, and denote it by $\mathscr{M}^c$. It inherits labels from $\mathscr{M}$ in the obvious way. See Figure 2.12.

There is a natural correspondence between the faces of a map and those of its core, since faces are not destroyed by the removal of leaves. If F is a face of $\mathscr{M}$ and F^c is the corresponding face of $\mathscr{M}^c$, then the boundary walk of F^c is obtained from that of F by removing all occurrences of vertices and edges not in the core. The rotator of a vertex v in $\mathscr{M}^c$ is therefore obtained from its

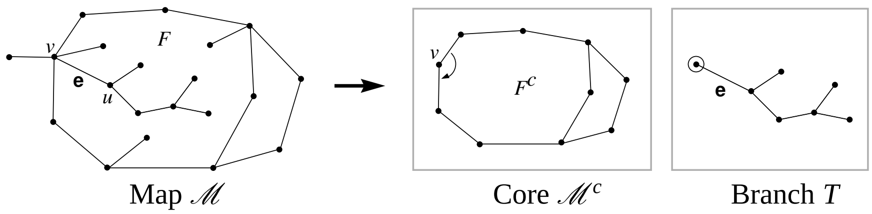

Figure 2.13: A map, its core, and one of its branches.

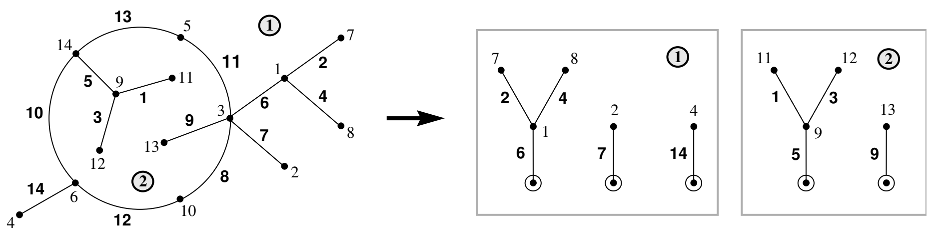

Figure 2.14: A properly labelled map and its branches.

rotator in $\mathscr{M}$ by deleting edges not in $\mathscr{M}^c$. It follows that v is at a descent of F^c if and only if it is at a descent of F, as can be verified in Figure 2.12.

Let $\mathscr{M}$ be any map that is not a plane tree, and let $e = \{u, v\}$ be an edge of $\mathscr{M}$ such that v lies in $\mathscr{M}^c$ but u does not. Then detaching e from v results in two maps; one of these contains v, and the other is a vertex-rooted plane tree T whose root is a leaf, incident only with edge e. The root of T may be regarded as "missing" so that this decomposition preserves vertices. We call T a **branch** of face F, and edge e the **stem** of this branch. Vertex v is known as the **base vertex** of T, and its **base corner** is the corner of F^c at which e was attached. See Figure 2.13 for an illustration. The base corner of T in F^c is indicated with an arrow in the diagram.

If the vertices of $\mathscr{M}$ are labelled, then the non-root vertices of its branches are also naturally labelled, while their roots are not. For example, Figure 2.14 displays the branches of properly labelled map, grouped by the face to which they belong.

The next two results are clear from the definitions above and the increasing rotator condition. The first of these lemmas makes the pruning of trees a plausible method for decomposing generic maps, and the second allows descent structure to be preserved in the pruning process.

Lemma 2.6.1. *Two maps are isomorphic if and only if their cores are isomorphic and the branches based at corresponding vertices coincide.* $\square$

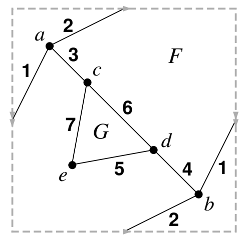

Figure 2.15: A two-face map on the torus.

Lemma 2.6.2. *Let F be a face of the map $\mathscr{M}$, and let T be a branch of F. Then every vertex of T is at a descent of F.* □

2.6.2 Normally Indexed Boundary Walks

Let $\mathscr{M}$ be a vertex-labelled map, and let F be a face of $\mathscr{M}$ of degree $k+1$. Then there are $k+1$ distinct vertex-edge pairs (v, e) occurring along the boundary walk W of F. If we fix one such pair, (v_0, e_0), then the symbols v_i and e_i, for $1 \leq i \leq k$, are well-defined by the assertion that $W = ((v_0, e_0), \ldots, (v_k, e_k))^\circ$. We therefore say that W can be *indexed* in $k+1$ distinct ways by the symbols v_i and e_i. We would like to distinguish one of these $k+1$ possibilities as a canonical indexing scheme for W. Phrased differently, we wish to determine a canonical "starting point" for boundary walks.

Definition 2.6.3. *Let $W = ((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be a boundary walk in a vertex-labelled map. For $0 \leq i \leq k$, define the list $L_i = (e_i, e_{i+1}, \ldots, e_{i+k}, v_i, v_{i+1}, \ldots, v_{i+k}) \in \mathbb{Z}^{2k+2}$. We say W is **normally indexed** by the symbols $\{v_0, \ldots, v_k\}$ and $\{e_0, \ldots, e_k\}$ if L_0 is minimal, under standard lexicographic order, amongst the lists $\{L_0, \ldots, L_k\}$.*

The fact that the vertex-edge pairs (v, e) of a boundary walk are distinct implies that any such walk of length $k+1$ admits a unique normal indexing by $\{v_0, \ldots, v_k\}$ and $\{e_0, \ldots, e_k\}$. Therefore asserting that $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ is a normally indexed boundary walk unambiguously defines the symbols v_i and e_i.

Example 2.6.4. The map in Figure 2.15 has two faces, F and G, of degrees 11 and 3, respectively. The boundary walk $((v_0, e_0), \ldots, (v_{10}, e_{10}))^\circ$ of F is normally indexed when

$$\begin{aligned}&((v_0, e_0), \ldots, (v_{10}, e_{10}))\\ &= ((a, 1), (b, 2), (a, 3), (c, 6), (d, 4), (b, 1), (a, 2), (b, 4), (d, 5), (e, 7), (c, 3)).\end{aligned}$$

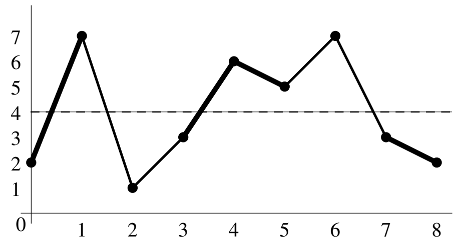

Figure 2.16: The proof of Lemma 2.6.5, with $L = (2, 7, 1, 3, 6, 5, 7, 3)^\circ$ and $e = 4$.

Similarly, the boundary walk $((u_0,\ f_0), (u_1,\ f_1), (u_2,\ f_2))^\circ$ of G is normally indexed precisely when $((u_0,\ f_0), (u_1,\ f_1), (u_2,\ f_2)) = ((d, 5), (c, 6), (e, 7))$. □

The definition we have given for normal indexing may seem somewhat unnatural. In particular, it would be far simpler to say that $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ is normally indexed when (v_0, e_0) is minimal amongst all vertex-edge pairs (v_i, e_i). Indeed, this alternative definition would serve our immediate purposes very well. The rationale supporting Definition 2.6.3 will be unveiled later, in §2.8.1, where we prove that it usually makes vertex labels irrelevant in the determination of normal indexing. Thus Definition 2.6.3 extends naturally to all maps, with or without vertex labels.

2.6.3 The Index of a Branch

We begin this section with a lemma concerning cyclic sequences. Its purpose may not be clear initially, but we shall see shortly that it plays a central rôle in everything to follow.

Lemma 2.6.5. *Let $L = (e_0, \ldots, e_k)^\circ$ be a cyclic list of real numbers with d descents. If $e \in \mathbb{R}$ is not in the list L, then there are exactly d values of i with $0 \le i \le k$ such that $(e_{i-1}, e, e_i)^\circ$ is nondecreasing.*

Proof. Let P be the polygonal path in the plane connecting the points $(0, e_0), \ldots, (k, e_k), (k+1, e_0)$, in that order. Let s_i be the i-th step of P. We call s_i an *up step* if $e_i > e_{i-1}$ and a *down step* otherwise. Thus down steps of P correspond with descents of L. For example, Figure 2.16 shows the path P corresponding to the list $L = (2, 7, 1, 3, 6, 5, 7, 3)^\circ$.

Note that $(e_{i-1}, e, e_i)^\circ$ is nondecreasing if and only if either $e_{i-1} < e < e_i$, or $e > e_{i-1} \ge e_i$, or $e_{i-1} \ge e_i > e$. Plainly, one of these conditions holds if and only if either (A) s_i is an up step which the line $y = e$ crosses, or (B) s_i is a down step which this line misses. Since the origin and terminus

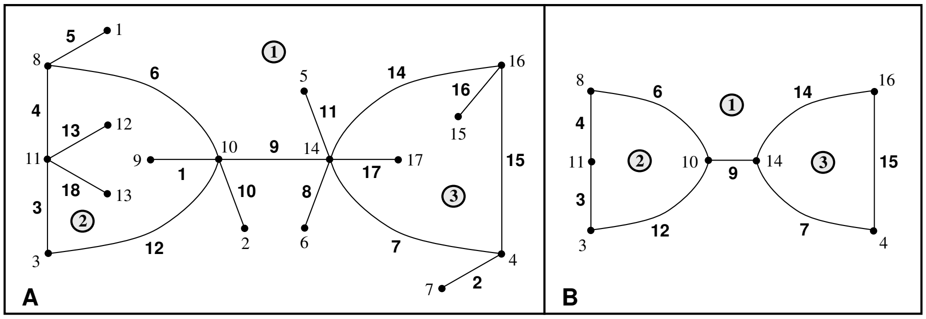

Figure 2.17: (A) The map $\mathscr{N}$ and (B) its core $\mathscr{N}^c$.

of P have the same y-coordinate, the numbers of up steps and down steps crossed by $y = e$ must be equal. Thus the number of indices i for which (A) or (B) is satisfied is equal to d, the total number of down steps of P, and this completes the proof. See Figure 2.16 for an illustration. The dashed line is $e = 4$, and steps for which $(e_{i-1}, e, e_i)^\circ$ is nondecreasing have been thickened. □

Let $\mathscr{M}$ be a properly labelled map. Let F and F^c be corresponding faces of $\mathscr{M}$ and $\mathscr{M}^c$, and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the normally indexed boundary walk of F^c. If T is a branch of F with stem e and base vertex v, then the base corner of T is (e_{b-1}, v, e_b) for a unique b with $0 \le b \le k$. Hence $(e_{b-1}, e, e_b)^\circ$ is increasing, as it is a subsequence of the rotator of v in $\mathscr{M}$. However, Lemma 2.6.5 implies that $(e_{j-1}, e, e_j)^\circ$ is increasing for exactly d values of j in the range $0 \le j \le k$, where d is the number of descents of F^c. Let these values of j be $j_1 < \cdots < j_d$. Then the **index** of branch T is the unique value of $i \in \{1, \ldots, d\}$ such that $j_i = b$.

Example 2.6.6. Consider the properly labelled map $\mathscr{N}$ and its core $\mathscr{N}^c$ drawn in Figure 2.17. For $s = 1, 2, 3$, let F_s and F_s^c, respectively, be the faces of $\mathscr{N}$ and $\mathscr{N}^c$ with label s. It will be convenient here to identify the branches of $\mathscr{N}$ by their stems; we write B_e for the branch with stem $e = 5$.

To compute the indices of the various branches $\{B_5, B_{11}, B_2, B_8, B_{10}\}$ of F_1, first note that the normally indexed boundary walk of F_1^c is

$$((v_0, e_0), \ldots, (v_8, e_8))^\circ = ((3, 3), (11, 4), (8, 6), (10, 9), (14, 14), (16, 15), (4, 7), (14, 9), (10, 12))^\circ.$$

Thus F_1^c has $d = 2$ descents, namely $e_5 \ge e_6$ and $e_8 \ge e_0$.

Consider the branch $T = B_5$ of F_1. The base corner of T is $(4, 8, 6) = (e_1, v_2, e_2)$, hence $b = 2$. The $d = 2$ values of j with $0 \le j \le 8$ such that $(e_{j-1}, 5, e_j)^\circ$ is increasing are $j_1 = 2$ and $j_2 = 6$. Since $b = 2 = j_1$, we have $i = 1$. Thus branch B_5 has index 1.

Face	**Branches**	**Indices**
1	$B_5,\ B_{11},\ B_2,\ B_8,\ B_{10}$	1, 1, 2, 2, 2
2	$B_{13},\ B_{18},\ B_1$	1, 1, 2
3	$B_{17},\ B_{16}$	1, 2

Table 2.1: The indices of the branches of $\mathscr{N}$.

Now consider branch $T = B_8$ of F_1. The base corner of this branch is $(7, 14, 9) = (e_6, v_7, e_7)$, so that $b = 7$. The $d = 2$ values of j with $0 \leq j \leq 8$ such that $(e_{j-1}, 8, e_j)^\circ$ is increasing are now $j_1 = 3$ and $j_2 = 7$. Since $b = 7 = j_2$, the index of B_8 is $i = 2$.

The branches of F_2 are $\{B_{13}, B_{18}, B_1\}$, and the normally indexed boundary walk of F_2^c is

$$((v_0, e_0), \ldots, (v_3, e_3))^\circ = ((11, 3), (3, 12), (10, 6), (8, 4))^\circ.$$

Thus F_2^c has $d = 3$ descents, namely $e_1 \geq e_2$, $e_2 \geq e_3$, and $e_3 \geq e_0$. The base corner of branch $T = B_1$ is $(12, 10, 6) = (e_1, v_2, e_2)$, so $b = 2$. The $d = 3$ values of j with $0 \leq j \leq 3$ for which $(e_{j-1}, 1, e_j)^\circ$ is increasing are $j_1 = 0$, $j_2 = 2$, and $j_3 = 3$. Since $b = 2 = j_2$, B_1 has index 2.

Computing the indices of the remaining branches of $\mathscr{N}$ in a like manner leads to the data listed in Table 2.1. $\square$

2.6.4 Pruning Trees

Consider again the map $\mathscr{N}$ drawn in Figure 2.17 and analyzed in Example 2.6.6. Let $(\theta_1, \theta_2, \theta_3) = (2, 3, 2)$ be the descent class of $\mathscr{N}^c$. For $s = 1, 2, 3$, and for each i with $1 \leq i \leq \theta_s$, let $\mathcal{B}_i^s$ be the set of all branches of face s of $\mathscr{N}$ that are of index i. From each of these sets $\mathcal{B}_i^s$, construct a new rooted tree T_i^s by identifying the roots of the various branches it contains. Finally, group these trees into the ordered forests

$$\mathcal{F}_1 = (T_1^1, T_2^1), \qquad \mathcal{F}_2 = (T_1^2, T_2^2, T_3^2), \qquad \text{and} \qquad \mathcal{F}_3 = (T_1^3, T_2^3).$$

In this way, $\mathscr{N}$ decomposes into the smooth map $\mathscr{N}^c$ and the forests $\mathcal{F}_1, \mathcal{F}_2$ and $\mathcal{F}_3$ depicted in Figure 2.18. Notice that these forests provide a complete encoding of the information in Table 2.1. We could therefore reverse this construction and fully recover $\mathscr{N}$ from the data $(\mathscr{N}^c, \mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3)$.

The process outlined above effectively prunes trees from the properly labelled map $\mathscr{N}$, and does so in a reversible manner. The next theorem formally specifies this process for arbitrary maps. We shall henceforth refer to the bijection described by this theorem as the **tree pruning bijection**.

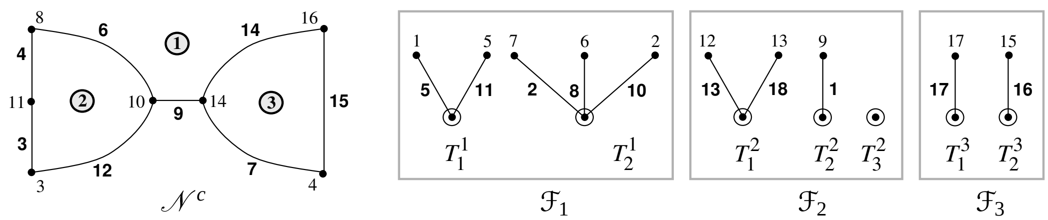

Figure 2.18: The smooth map and forests obtained from the map $\mathscr{N}$ of Figure 2.17A.

Theorem 2.6.7 (Tree Pruning Bijection). *Let $g \geq 0$ and $m \geq 1$, with $(g, m) \neq (0, 1)$, and let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be any m-part composition. Then there is a bijection between properly labelled genus g maps of descent class α and tuples $(\theta, \mathscr{S}, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ with the following properties:*

(a) $\theta = (\theta_1, \ldots, \theta_m)$ is a composition with $\theta_i \leq \alpha_i$.

(b) $\mathscr{S}$ is a smooth, vertex- and face-labelled genus g map of descent class θ.

(c) $\mathcal{F}_s$ is an ordered forest of θ_s rooted trees with labelled non-root vertices and edges.

(d) The descent set of face s of $\mathscr{S}$ together with the vertex labels of $\mathcal{F}_s$ partition $\mathbb{D}_s(\alpha)$.

(e) The edge labels of $\mathscr{S}$ together with those of $\mathcal{F}_1, \ldots, \mathcal{F}_m$ partition $\{1, 2, \ldots, r_g(\alpha)\}$.

Proof. Let $\mathscr{M}$ be a properly labelled genus g map of descent class $\alpha = (\alpha_1, \ldots, \alpha_m)$, and suppose its core $\mathscr{M}^c$ has descent class $\theta = (\theta_1, \ldots, \theta_m)$. Then, for $1 \leq s \leq m$ and $1 \leq i \leq \theta_s$, assemble all branches of face s of $\mathscr{M}$ that are of index i into a rooted tree T_i^s by identifying their roots as a common new root vertex. Let $\mathcal{F}_s = (T_1^s, \ldots, T_{\theta_s}^s)$ be the ordered forest consisting of the trees obtained from face s. We claim the tuple $(\theta, \mathscr{M}^c, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ satisfies properties (a) through (e). In fact, all conditions but (d) are immediate from the construction, and (d) is a direct result of Lemma 2.6.2.

Lemma 2.6.1 implies that the correspondence $\mathscr{M} \mapsto (\theta, \mathscr{M}^c, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ described above is one-one. We now prove it is also surjective, onto the set of all tuples $(\theta, \mathscr{S}, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ satisfying (a) through (e). To this end, let $(\theta, \mathscr{S}, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ be such a tuple, where $\mathcal{F}_s = (T_1^s, \ldots, T_{\theta_s}^s)$. Fix $s \in \{1, \ldots, m\}$ and $i \in \{1, \ldots, \theta_i\}$. Let e be an edge of T_i^s incident with the root. Detaching e from the root leaves another rooted tree B_e whose root is incident only with e. Now let F be the face of $\mathscr{S}$ labelled s, so that F has θ_s descents, and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be its normally indexed boundary walk. Then Lemma 2.6.5 implies that $(e_{j-1}, e, e_j)^\circ$ is increasing for exactly θ_s values of j with $0 \leq j \leq k$, say $j_1 < \cdots < j_{\theta_s}$. Attach the tree B_e to $\mathscr{S}$ at v_{j_i}, doing so in the unique manner

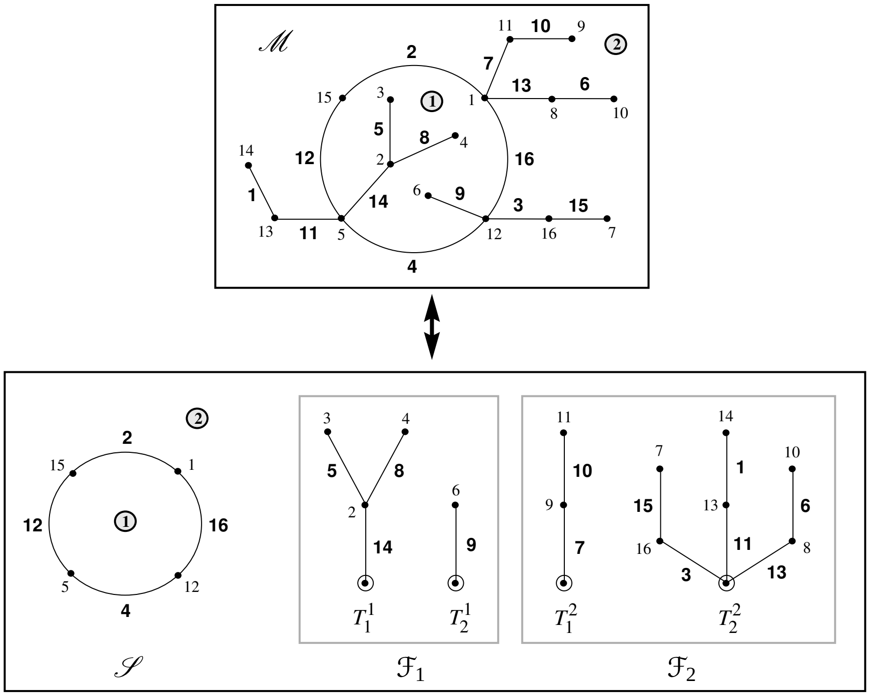

Figure 2.19: Tree pruning bijection.

that leaves the rotator of v_{j_i} increasing. Repeat this process for all s, i and e to obtain a vertex- and face-labelled map $\mathscr{M}$. Clearly $\mathscr{M}^c = \mathscr{S}$, and the fact that $\mathscr{M}$ is of descent class α follows from conditions (a) through (e). □

Example 2.6.8. Figure 2.19 illustrates the tree pruning bijection. The two-face planar map $\mathscr{M}$ in the upper panel corresponds with the tuple $((2,2), \mathscr{S}, \mathfrak{F}_1, \mathfrak{F}_2)$, whose components are shown in the lower panel. □

The tree pruning bijection suggests that understanding the nature of transitive factorizations is tantamount to understanding the structure of smooth properly labelled maps. In light of this revelation we make the following definitions.

Definition 2.6.9. *Let $S_g(\theta)$ be the number of smooth, properly labelled, genus g maps of descent class θ. For $m \geq 1$ we define the generating series for the numbers $\{S_g(\theta) : \ell(\theta) = m\}$ by*

$$\Gamma_m^{(g)}(\mathbf{z}, u) = \sum_{k \geq 1} \sum_{\substack{\theta \models k \\ \ell(\theta) = m}} S_g(\theta) \frac{\mathbf{z}^{\boldsymbol{\theta}}}{\boldsymbol{\theta}!} \frac{u^{r_g(\theta)}}{r_g(\theta)!},$$

where $\mathbf{z} = (z_1, \ldots, z_m)$. When considering the genus 0 series, we often write Γ_m in place of $\Gamma_m^{(0)}$.

The apparent discrepancy between this definition of Γ_m and the notation used in §2.5 is resolved by comparing (2.36) with the next theorem, which makes precise the connection between the series $\Gamma_m^{(g)}$ and $\Psi_m^{(g)}$. This result finally identifies the combinatorial significance of the dependence of $\Psi_m^{(g)}$ on the tree series. We remind the reader that the symbol w_i in the statement of theorem represents the tree series $w_i = w(x_i, u)$. (See Definition 2.3.8).

Theorem 2.6.10. *For any $g \geq 0$ and $m \geq 1$ with $(g, m) \neq (0, 1)$ we have*

$$\Psi_m^{(g)}(\mathbf{x}, u) = \Gamma_m^{(g)}(\mathbf{w}, u),$$

where $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{w} = (w_1, \ldots, w_m)$.

Proof. Let α be any m-part composition. Then the number of tuples $(\theta, \mathscr{S}, \mathcal{F}_1, \ldots, \mathcal{F}_m)$ satisfying properties (a) through (e) of the Theorem 2.6.7 is equal (by the theorem) to $M_g(\alpha)$. We now count these objects directly.

The genus g map $\mathscr{S}$ is of descent class θ. There are, by definition, exactly $S_g(\theta)$ properly labelled maps of this type. The forest $\mathcal{F}_i$ consists of θ_i rooted trees, with labelled non-root vertices and edges. The series counting such trees is $\exp(uw_i)$, where u marks edges and x_i marks labelled vertices. Thus the series counting forests $\mathcal{F}_i$ is $\exp(uw_i)^{\theta_i}$. Distributing labels properly between $\mathscr{S}$ and the forests $\mathcal{F}_1, \ldots, \mathcal{F}_m$ (*i.e.* according to (d) and (e) of Theorem 2.6.7) amounts to multiplying the generating series of these structures. Doing so, and summing over the parameter θ, gives

$$\sum_{n\geq 1} \sum_{\substack{\alpha \models n \\ \ell(\alpha)=m}} M_g(\alpha) \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \frac{u^{r_g(\alpha)}}{r_g(\alpha)!} = \sum_{k\geq 1} \sum_{\substack{\theta \models k \\ \ell(\theta)=m}} S_g(\theta) \frac{x_1^{\theta_1}}{\theta_1!} \cdots \frac{x_m^{\theta_m}}{\theta_m!} \frac{u^{r_g(\theta)}}{r_g(\theta)!} \cdot \left(e^{uw_1}\right)^{\theta_1} \cdots \left(e^{uw_m}\right)^{\theta_m}.$$

The result follows at once upon simplifying this expression with the aid of (2.16). □

Perhaps surprisingly, Theorem 2.6.10 and the identity (2.22) combine to show that $\Gamma_m^{(g)}(\mathbf{z}, u)$ is rational in $\mathbf{z}$ and u for all $m \geq 3$. For example, in genus 0 we immediately deduce the following theorem.

Theorem 2.6.11. *For any $m \geq 2$, we have*

$$\Gamma_m(\mathbf{z}, u) = u^{2m-2} \left(\sum_{i=1}^{m} \frac{z_i}{1 - uz_i} \frac{\partial}{\partial z_i} \right)^{m-3} \prod_{i=1}^{m} \frac{z_i}{1 - uz_i}.$$

□

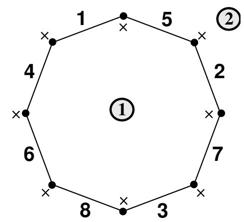

Figure 2.20: The 2-face planar map associated with $(1, 5, 2, 7, 3, 8, 6, 4)^\circ$.

Through the tree pruning bijection, a combinatorial proof of this result would provide a combinatorial proof of Theorem 2.3.9, which is equivalent to Hurwitz's formula. In the next section we describe some progress that has been made along these lines.

2.7 Combinatorial Constructions for Smooth Maps

We have seen that the combinatorics of transitive factorizations is essentially equivalent to that of smooth properly labelled maps. In this section we investigate these maps in detail, focusing on the planar case. In particular, we present bijections that illuminate the structure of smooth planar maps with two and three faces, and thereby offer combinatorial proofs of Theorem 2.6.11 for $m = 2, 3$. We also describe a general differential decomposition for maps with at least four faces. This yields another proof of Theorem 2.5.6, and explains the serendipitous algebraic simplifications exploited in the earlier, algebraic derivation of that result.

The comments made in §2.4.11 concerning interpretations of $\Psi_m^{(g)}(\mathbf{x}, u)$ also apply to $\Gamma_m^{(g)}(\mathbf{z}, u)$. That is, the latter series can largely be regarded as counting smooth face-labelled maps with respect to labelled edges and descent class. We take this perspective throughout this section.

2.7.1 Two-Face Smooth Planar Maps

Recall that a **circular permutation** of $[n]$ is a cyclic ordering of the elements of $[n]$. That is, $\sigma = (a_0, \ldots, a_{n-1})^\circ$ is a circular permutation of $[n]$ if $\{a_0, a_1, \ldots, a_{n-1}\} = \{1, 2, \ldots, n\}$. The pair (a_{i-1}, a_i) of consecutive elements of σ is called a **rise** if $a_{i-1} < a_i$, and a **fall** if $a_{i-1} > a_i$.

A smooth planar map with two faces is simply a cycle, so smooth face-labelled planar maps of descent class $\alpha = (a, b)$ correspond with circular permutations having a rises and b falls. For instance, the map corresponding to the circular permutation $\sigma = (1, 5, 2, 7, 3, 8, 6, 4)^\circ$ is drawn in Figure 2.20. Note that σ has 3 rises and 5 falls, and the map is of descent class $(3, 5)$.

We require the following well-known result. The proof given here is based on an inclusion-exclusion argument.

Lemma 2.7.1. *Let* $a, b \geq 1$ *and set* $n = a + b$*. Then there are*

$$n!\,[r^a f^b]\log\left(\frac{r-f}{re^f - fe^r}\right)$$

circular permutations of $\{1, \ldots, n\}$ *having exactly* a *rises and* b *falls.*

Proof. Let $\sigma = (a_0, \ldots, a_{n-1})^\circ$ be a circular permutation of $[n]$. Let $I_\sigma = \{i : a_i < a_{i+1}\}$ index the rises of σ, and let $\mathcal{S}_\sigma$ be the set of circular sequences that can be obtained by choosing a subset $I \subset I_\sigma$ and, for all $i \in I$, replacing the pair a_i, a_{i+1} with the pattern $a_i, *, a_{i+1}$. For example, if $\sigma = (1, 3, 2, 5, 4)^\circ$ then

$$\mathcal{S}_\sigma = \{(1,3,2,5,4)^\circ, (1,*,3,2,5,4)^\circ, (1,*,3,2,*,5,4)^\circ, (1,3,2,*,5,4)^\circ\}.$$

Note that each element of $\mathcal{S}_\sigma$ corresponds with a unique cyclic list of maximal contiguous patterns of the form $a_i * \cdots * a_{i+j}$, where $a_i < \cdots < a_{i+j}$. For example, if $\sigma = (1, 3, 4, 2, 5, 6)^\circ$ then $(1, *, 3, *, 4, 2, *, 5, 6)^\circ \in \mathcal{S}_\sigma$ corresponds to the list $(1*3*4, 2*5, 6)^\circ$. Let $F(x, u)$ be the generating series for such lists, with x marking symbols of $[n]$ (exponentially) and u marking occurrences of $*$ (ordinarily). The generating series, with respect to these markers, for patterns $a_i * \cdots * a_{i+j}$ satisfying $a_i < \cdots < a_{i+j}$ is $\sum_{k\geq 1} u^{k-1}x^k/k! = (e^{ux} - 1)/u$. Since $\log(1-z)^{-1}$ is the exponential generating series for cycles, it follows that

$$F(x,u) = \log\left(1 - \frac{e^{ux}-1}{u}\right)^{-1}.$$

Let $G(x, r)$ be the generating series for circular permutations on $n \geq 2$ symbols, where x marks these symbols (exponentially) and r marks rises between them (ordinarily). Then the above replacement argument gives $G(x, r + 1) = F(x, r) - x$, where the subtraction eliminates the single permutation on one symbol. Thus

$$G(x,r) = F(x, r-1) - x = \log\left(\frac{r-1}{r - e^{(r-1)x}}\right) - x. \tag{2.41}$$

A circular permutation on n symbols with a rises has $b = n - a$ falls. The number of such permutations is therefore $(a + b)!\,[r^a f^b]\,G(f, rf^{-1})$. The result follows from (2.41). □

Proposition 2.7.2.
$$\Gamma_2^{(0)}(z_1, z_2, u) = \log\left(\frac{z_1 - z_2}{z_1 e^{uz_2} - z_2 e^{uz_1}}\right).$$

Proof. Lemma 2.7.1 gives

$$\Gamma_2^{(0)}(z_1, z_2, u) = \sum_{a,b\geq 1}\left((a+b)!\,[r^a f^b]\,\log\left(\frac{r-f}{re^f - fe^r}\right)\right) z_1^a z_2^b \frac{u^{a+b}}{(a+b)!},$$

and the result follows immediately upon simplification. □

Corollary 2.7.3.
$$\Psi_2^{(0)}(x_1, x_2, u) = \log\left(\frac{w_1 - w_2}{w_1 e^{uw_2} - w_2 e^{uw_1}}\right).$$

Proof. This follows immediately from Proposition 2.7.2, Theorem 2.6.10 and (2.16). □

With some routine algebra, the $m = 2$ case of Theorem 2.6.11 can be deduced from these results. In fact, Corollary 2.7.3 gives precisely the identity (2.39) that was manipulated to prove Corollary 2.5.4, which is in turn equivalent to the Theorem 2.6.11 when $m = 2$.

This derivation of the series $\Psi_2^{(0)}(x_1, x_2, u)$ is certainly more explanative than that given in §2.5.2, but it still is not bijective. Moreover, it is not at all clear how this method can be extended to count maps with more than two faces. We shall therefore now make a fresh attempt at proving Theorem 2.6.11, first in the case $m = 2$, and then more generally.

2.7.2 Attaching Edges to a Map

The following lemma, which is a slight modification of Lemma 2.6.5, lies at the heart of all the remaining results of this chapter. Its specific rôle will be made clear below.

Lemma 2.7.4. *Let $L = (e_0, \ldots, e_k)^\circ$ be a cyclic list of real numbers with d descents. Let e be any number not in the list L, and let $i_1 < \cdots < i_d$ be the values of i with $0 \leq i \leq k$ such that $(e_{i-1}, e, e_i)^\circ$ is nondecreasing. Then, for any $1 \leq t < s \leq d$, the cyclic list $(e_{i_t}, e_{i_t+1}, \ldots, e_{i_s-1}, e)^\circ$ has exactly $s - t$ descents.*

Proof. Let $f = e + \delta$, where $0 < \delta < \min_i |e - e_i|$. Then clearly $(e_{j-1}, f, e_j)^\circ$ is nondecreasing if and only if $(e_{j-1}, e, e_j)^\circ$ is nondecreasing. In particular, there are exactly $s - t - 1$ indices j with $i_t < j < i_s$ such that $(e_{j-1}, f, e_j)^\circ$ is nondecreasing. Also note that $(e, f, e_{i_t})^\circ$ is nondecreasing, while $(e_{i_s-1}, f, e)^\circ$ is not. Let $(f_0, \ldots, f_{i_t+i_s})^\circ = (e_{i_t}, e_{i_t+1}, \ldots, e_{i_s-1}, e)^\circ$. Then $(f_{i-1}, f, f_i)^\circ$ is nondecreasing for exactly $(s-t-1)+1 = s-t$ values of i with $0 \leq i \leq m$. Lemma 2.6.5 therefore implies that $(f_0, \ldots, f_{i_t+i_s})^\circ$ has $s - t$ descents, as required. □

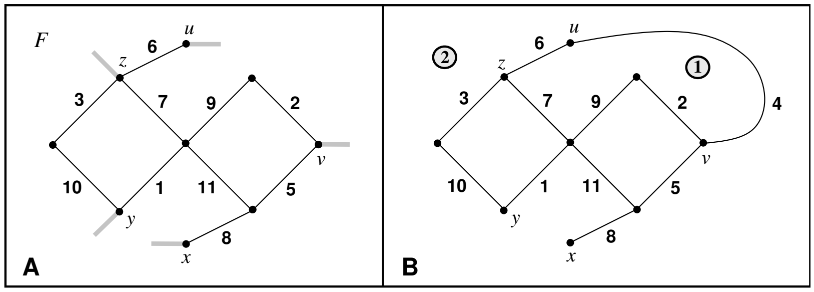

Figure 2.21: Attaching an edge to a map.

Let F be a face of the map $\mathscr{M}$ and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the boundary walk of F. If $e \in \mathbb{R}$ and $(e_{i-1}, e, e_i)^\circ$ is nondecreasing, then we say corner (e_{i-1}, v_i, e_i) **admits** label e. We write $\mathcal{A}_F(e)$ for the set of all corners of F which admit e. Graphically, the condition $(e_{i-1}, v_i, e_i) \in \mathcal{A}_F(e)$ means that if an edge labelled e were attached to v_i in the corner (e_{i-1}, v_i, e_i), then increasing rotators would be maintained.

Let $c = (e_{i-1}, v_i, e_i)$ and $c' = (e_{j-1}, v_j, e_j)$ be distinct corners of F. Then both c and c' belong to $\mathcal{A}_F(e)$ if and only if a new map can be produced by adding an edge $\{v_i, v_j\}$ with label e between these corners. Suppose this is the case, and let $\mathscr{N}$ be the map resulting from the addition of $\{v_i, v_j\}$. Then $\mathscr{N}$ has the same genus as $\mathscr{M}$ but one extra face. Indeed, $\{v_i, v_j\}$ separates F into two faces of $\mathscr{N}$. For our purposes, it will be convenient to assign labels to these faces. We therefore introduce the (admittedly convoluted) notation $\mathscr{M} \oplus (c, c')^e_{s,t}$ to denote the map obtained by assigning labels s and t, respectively, to the faces of $\mathscr{N}$ containing corners (e, v_i, e_i) and (e, v_j, e_j).

Example 2.7.5. Consider the map $\mathscr{M}$ drawn in Figure 2.21A, with face F as indicated. Then $\mathcal{A}_F(4) = \{c_0, c_1, c_2, c_3, c_4\}$, where

$$c_0 = (3, z, 6), \quad c_1 = (6, u, 6), \quad c_2 = (2, v, 5), \quad c_3 = (8, x, 8), \quad c_4 = (1, y, 10). \qquad (2.42)$$

Shaded half-edges extending from these corners into F have been drawn to emphasize how an edge with label 4 could be attached to $\mathscr{M}$. Figure 2.21B shows the four-face map $\mathscr{M} \oplus (c_1, c_2)^4_{1,2}$ resulting from the attachment of the edge $\{u, v\}$ with label 4 between corners c_1 and c_2. Note the face-labelling of this new map. □

Observe that the descents of F are split amongst faces s and t of $\mathscr{M} \oplus (c, c')^e_{s,t}$. That is, if F has d descents, then faces s and t have d_s and d_t descents, respectively, where $(d_s, d_t) \models d$. In fact,

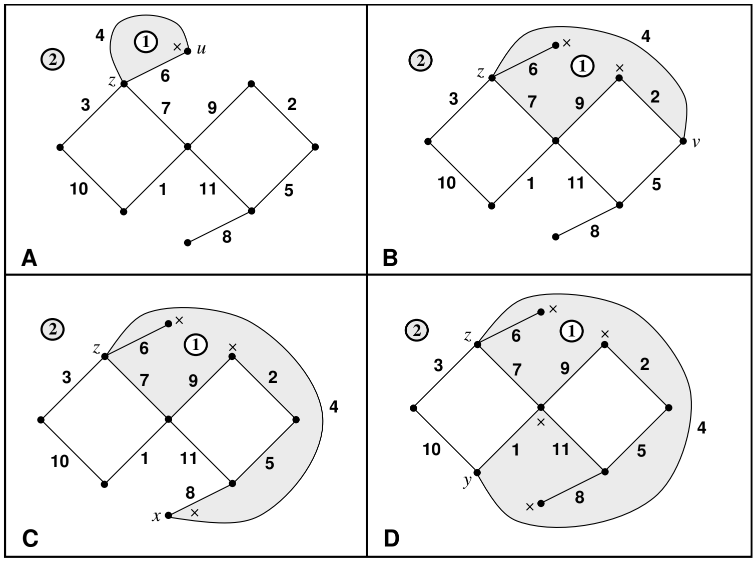

Figure 2.22: An illustration of Lemma 2.7.6.

if label e and corner $c \in \mathcal{A}_F(e)$ are fixed, then the next lemma shows that with every composition $(d_s, d_t) \models d$ there corresponds a *unique* corner $c' \in \mathcal{A}_F(e)$ such that faces s and t of $\mathscr{M} \oplus (c, c')^e_{s,t}$ have d_s and d_t descents, respectively. This enables us to add edges to a map while maintaining complete control of its descent class.

Lemma 2.7.6. *Let $\mathscr{M}$ be a map and let F be a face of $\mathscr{M}$ with d descents and boundary walk $((v_0, e_0), \dots, (v_k, e_k))^\circ$. Let $e \in \mathbb{R}$ be distinct from $e_0, \dots, e_k$. Then $|\mathcal{A}_F(e)| = d$. Moreover, if $\mathcal{A}_F(e) = \{c_0, \dots, c_{d-1}\}$, where $c_j = (e_{i_j-1}, v_{i_j}, e_{i_j})$ and $0 = i_0 < \cdots < i_{d-1} \leq k$, then faces s and t of $\mathscr{M} \oplus (c_0, c_j)^e_{s,t}$ have j and $d - j$ descents, respectively, for $1 \leq j \leq d-1$.*

Proof. That $|\mathcal{A}_F(e)| = d$ follows immediately from the definition of $\mathcal{A}_F(e)$ and Lemma 2.6.5. Suppose $\mathcal{A}_F(e) = \{c_0, \dots, c_{d-1}\}$, where $c_j = (e_{i_j-1}, v_{i_j}, e_{i_j})$ and $0 = i_0 < \cdots < i_{d-1} \leq k$. If $1 \leq j \leq d-1$, then face s of $\mathscr{M} \oplus (c_0, c_j)^e_{s,t}$ has boundary walk $((v_{i_0}, e_{i_0}), \dots, (v_{i_j-1}, e_{i_j-1}), (v_{i_j}, e))^\circ$. Lemma 2.7.4 shows that this face has j descents. Since a vertex is at a descent of faces s or t of $\mathscr{M} \oplus (c_0, c_j)^e_{s,t}$ if and only if it is at a descent of face F of $\mathscr{M}$, face t has $d - j$ descents. □

Example 2.7.7. Reconsider the map $\mathscr{M}$ with face F drawn in Figure 2.21A. In Example 2.7.5 we saw that $|\mathcal{A}_F(4)| = 5$, and plainly F has 5 descents. Define corners $c_0, \dots, c_4$ of F as in (2.42). Then panels A through D of Figure 2.22 illustrate the maps $\mathscr{M} \oplus (c_0, c_1)^4_{1,2}$, $\mathscr{M} \oplus (c_0, c_2)^4_{1,2}$, $\mathscr{M} \oplus (c_0, c_3)^4_{1,2}$ and $\mathscr{M} \oplus (c_0, c_4)^4_{1,2}$, respectively. In each of these maps, face 1 has been highlighted

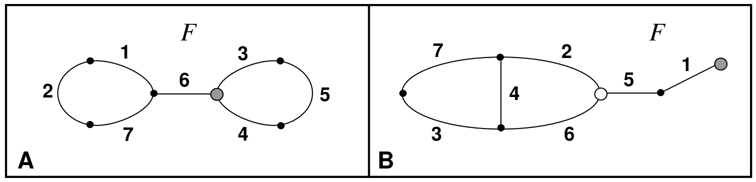

Figure 2.23: Three-face planar maps with a tail.

and its descents have been marked with crosses. Note that faces 1 and 2 of $\mathscr{M} \oplus (c_0, c_j)_{1,2}^4$ have j and $5 - j$ descents, respectively, for $1 \leq j \leq 4$. □

2.7.3 Two-Face Smooth Planar Maps Revisited

We are now ready to make another attempt at counting two-face smooth planar maps, this time through the use of Lemma 2.7.6. We begin by introducing some convenient terminology.

Definition 2.7.8. *An **ordered path** is a planar map with one face and exactly two leaves, one coloured white and the other grey. The leaves are called the **ends** of the path.*

Definition 2.7.9. *A map $\mathscr{M}$ is said to have a **tail** in face F if either (1) $\mathscr{M}$ is smooth and a vertex at a descent of F has been coloured grey, or (2) $\mathscr{M}$ contains only one branch, which is an ordered path in face F whose white end is the base vertex of the branch.*

A diagram reveals the reason for our use of the term *tail*. For example, both maps of Figure 2.23 have a tail in the face marked F. The following lemma shows that it is easy to derive the generating series for maps with a tail from the series for smooth maps.

Lemma 2.7.10. *Let $\theta \models n$ with $\ell(\theta) = m$. Then the number of genus g, face-labelled maps of descent class θ with a tail in face i is*

$$\left[\mathbf{z}^{\boldsymbol{\theta}} \frac{u^{r_g(\theta)}}{r_g(\theta)!}\right] \frac{z_i}{1 - uz_i} \cdot \frac{\partial}{\partial z_i} \Gamma_m^{(g)}(\mathbf{z}, u).$$

Proof. A map with a tail in face i is formed by selecting a smooth map and either distinguishing a descent of face i, or attaching a branch T in face i, where T is a path. A branch can be attached at a vertex if and only if the increasing rotator condition is maintained in the process. Thus, if a face has d descents, then Lemma 2.6.5 implies that a branch can be attached in that face at exactly d possible base vertices.

The series $(z_i\,\partial/\partial z_i)\Gamma_m^{(g)}(\mathbf{z},u)$ counts smooth maps with one descent of face i distinguished, the series $uz_i/(1-uz_i)$ counts paths T to be attached as a branch in face i, and $(z_i\,\partial/\partial z_i)\Gamma_m^{(g)}(\mathbf{z},u)$ counts maps with a distinguished base vertex for attachment of T. Thus the series counting maps with a tail in face i is

$$z_i\frac{\partial}{\partial z_i}\Gamma_m^{(g)}(\mathbf{z},u)+\frac{uz_i}{1-uz_i}\cdot z_i\frac{\partial}{\partial z_i}\Gamma_m^{(g)}(\mathbf{z},u).$$

The result follows upon simplification. □

Attaching a single edge to an ordered path clearly produces a two-face map. Moreover, if the added edge connects one end of the path to some interior vertex, then the two-face map so produced has a tail. In the next theorem, we show how this naïve construction leads to a bijection between ordered paths and maps with tails.

Theorem 2.7.11. *Fix $\theta=(\theta_1,\theta_2)\models n$. There is a bijection between face-labelled planar maps of descent class θ with a tail, and edge-labelled pairs $(\lambda,\mathscr{P})$, where λ is an edge and $\mathscr{P}$ is an ordered path containing n vertices.*

Proof. Let $\mathcal{M}_\theta$ be the set of face-labelled planar maps that have a tail and are of descent class θ, and let $\mathcal{P}_n$ be the set of all pairs $(\lambda,\mathscr{P})$ of the form described in the theorem. We define $\Omega_\theta:\mathcal{P}_n\longrightarrow\mathcal{M}_\theta$ by constructing $\Omega_\theta(\lambda,\mathscr{P})$ as follows.

Let F be the sole face of $\mathscr{P}$, and let $((v_0,e_0),\ldots,(v_k,e_k))^\circ$ be its boundary walk, where v_0 and v_m are the white and grey ends of $\mathscr{P}$, respectively. For $0\le i\le k$, let $c_i=(e_{i-1},v_i,e_i)$. This setup is illustrated below.

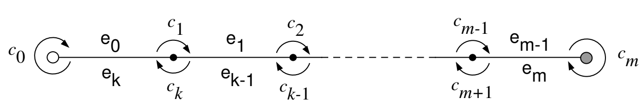

(We remark that we could be more definitive here, as we clearly have $m=n-1$ and $k=2n-3$. However, our more general notation has been chosen with later abstractions in mind, and does not muddy the argument in any case.)

Plainly, F has n descents and $c_0\in\mathcal{A}_F(\lambda)$. Therefore Lemma 2.7.6 guarantees a unique corner $c_r\in\mathcal{A}_F(\lambda)$, with $0<r\le k$, such that the two-face map $\mathscr{P}\oplus(c_0,c_r)^{\lambda}_{1,2}$ is of descent class $(\theta_1,n-\theta_1)=(\theta_1,\theta_2)$. Let $\Omega_\theta(\lambda,\mathscr{P})$ be this new map. Strip v_0 of its colour and, if $r\ne m$, colour vertex v_r white. The construction is illustrated in Figure 2.24. If $r\ne m$, then v_m is the only vertex of $\Omega_\theta(\lambda,\mathscr{P})$ of degree 1, and it is the grey end of a path extending from v_r. If $r=m$, then $\Omega_\theta(\lambda,\mathscr{P})$ is smooth and v_m is grey. In either case, $\Omega_\theta(\lambda,\mathscr{P})\in\mathcal{M}_\theta$.

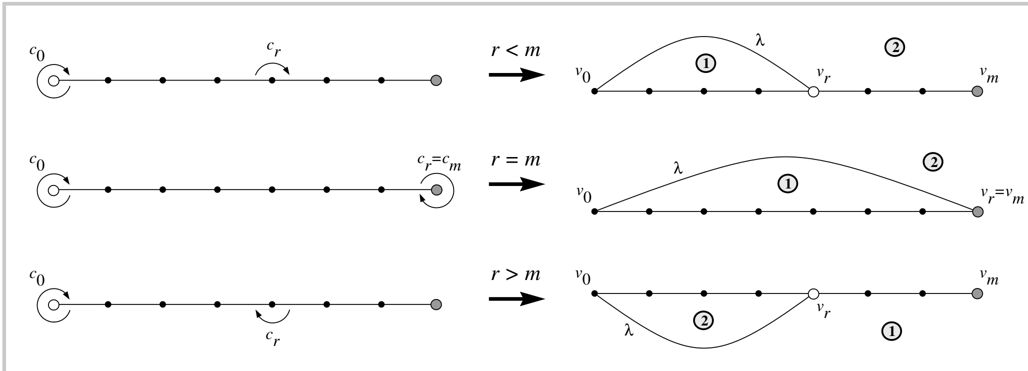

Figure 2.24: The maps produced by $\Omega_{(\theta_1,\theta_2)}$.

We claim that $\Omega_\theta : \mathcal{P}_n \longrightarrow \mathcal{M}_\theta$ is a bijection. It is clearly one-one, since any given map $\mathscr{M} = \Omega_\theta(\lambda, \mathscr{P})$ is of one of the three types shown on the right side of Figure 2.24, and thus both edge λ and vertex v_0 of $\mathscr{M}$ can be uniquely identified, as follows. First let v_r be the grey vertex if $\mathscr{M}$ is smooth, and the white vertex otherwise. Then $\lambda = \{v_0, v_r\}$ is the unique edge such that the vertex-edge pair (v_r, λ) appears in the boundary walk of face 1 of $\mathscr{M}^c$.

To see that Ω_θ is onto, observe that every $\mathscr{M} \in \mathcal{M}_\theta$ belongs to exactly one of the three classes of maps on the right-hand side of Figure 2.24. Thus edge λ and vertices v_0, v_r of $\mathscr{M}$ are uniquely determined, as above. Let $\mathscr{P}$ be the ordered path with white end v_0 that is obtained by deleting edge λ from $\mathscr{M}$, and let F be the sole face of $\mathscr{P}$. Then $\mathscr{M} = \mathscr{P} \oplus (c_0, c_r)^{\lambda}_{1,2}$, where v_0 is at corner $c_0 \in \mathcal{A}_F(\lambda)$ and v_r is at corner $c_r \in \mathcal{A}_F(\lambda)$. Since $\mathscr{M} \in \mathcal{M}_\theta$ is of descent class θ, so also is $\mathscr{P} \oplus (c_0, c_r)^{\lambda}_{1,2}$. But, by definition, $\Omega_\theta(\lambda, \mathscr{P}) = \mathscr{P} \oplus (c_0, c')^{\lambda}_{1,2}$, where $c' \in \mathcal{A}_F(\lambda)$ is the *unique* corner such that this map is of descent class θ. It follows that $c' = c_r$, and thus $\mathscr{M} = \Omega_\theta(\lambda, \mathscr{P})$. Therefore Ω_θ is surjective, and the proof is complete. □

Example 2.7.12. Let $\mathscr{P}$ be the ordered path shown in Panel A of Figure 2.25, and let $\lambda = 5$. The single face F of $\mathscr{P}$ has boundary walk $((v_0, e_0), \ldots, (v_{13}, e_{13}))^\circ$, where

$$(e_0, e_1, \ldots, e_{13}) = (8, 4, 3, 6, 1, 7, 2, 2, 7, 1, 6, 3, 4, 8).$$

Let $c_i = (e_{i-1}, v_i, e_i)$ for $0 \le i \le 13$. Then $\mathcal{A}_F(5) = \{c_{i_0}, \ldots, c_{i_7}\}$, where

$$(i_0, \ldots, i_7) = (0, 2, 3, 5, 7, 8, 10, 13).$$

Except for c_{i_0}, these corners are in indicated with small crosses in Panel B. By Lemma 2.7.6, the

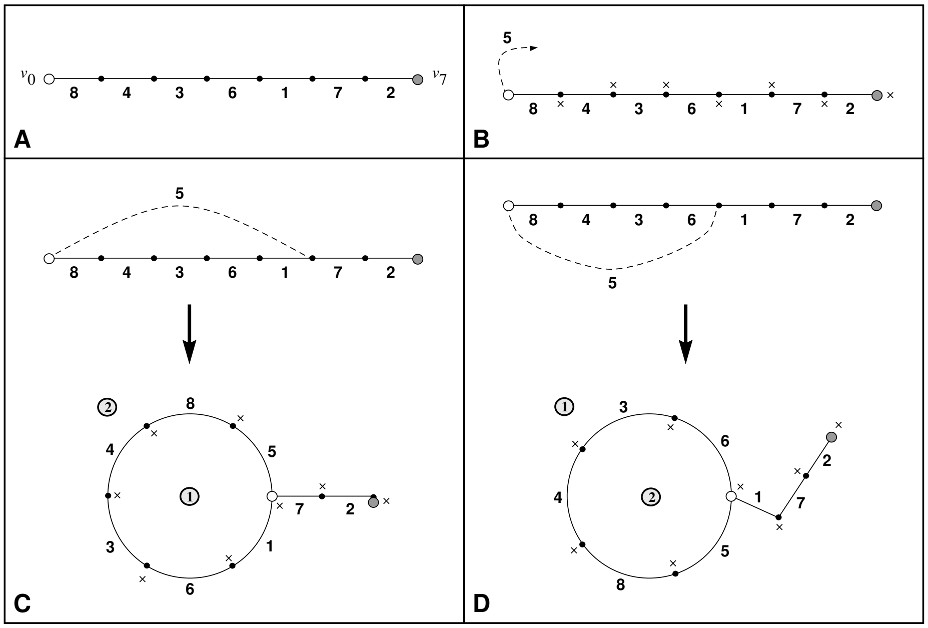

Figure 2.25: Constructing two-face planar maps from ordered paths.

map $\mathscr{P} \oplus (c_0, c_{i_j})^5_{1,2}$ has descent class $(j, 8-j)$ for $1 \leq j \leq 7$. Thus for the bijection Ω_θ determined by $\theta = (j, 8-j)$ we have $\Omega_\theta(\lambda, \mathscr{P}) = \mathscr{P} \oplus (c_0, c_r)^\lambda_{1,2}$, where $r = i_j$.

For example, if $\theta = (3, 5)$ then $\Omega_\theta(\lambda, \mathscr{P})$ is obtained by adding edge $\{v_0, v_{i_3}\} = \{v_0, v_5\}$ with label 5 to $\mathscr{P}$ between corners c_0 and c_5, as illustrated in Panel C. If $\theta = (6, 2)$, then edge $\{v_0, v_{i_6}\} = \{v_0, v_{10}\}$ is instead added to $\mathscr{P}$ between corners c_0 and c_{10} to produce the map of Panel D. Crosses have been drawn at the descents of these maps, showing that they are, indeed, of descent classes (3, 5) and (6, 2), respectively.

To reverse the construction illustrated in Panel C, let v be the white vertex of the final map $\mathscr{M}$ shown there, and observe that the vertex-edge pair $(v, 5)$ appears in the boundary walk of face 1 of the core $\mathscr{M}^c$. This identifies edge 5 as the additional edge. Transfer the white colouring of v to the opposite end of edge 5. Then removal of edge 5 results in the initial path $\mathscr{P}$. □

Together with Lemma 2.7.10, the previous theorem provides another combinatorial proof of Theorem 2.6.11 in the case $m = 2$. This is the content of the following corollary.

Corollary 2.7.13.

$$\left(\frac{z_1}{1-uz_1}\frac{\partial}{\partial z_1} + \frac{z_2}{1-uz_2}\frac{\partial}{\partial z_2}\right)\Gamma_2(z_1, z_2, u) = \frac{u^2 z_1 z_2}{(1-uz_1)(1-uz_2)}. \tag{2.43}$$

Proof. Let $G(z_1, z_2, u)$ be the series on the left-hand side of (2.43). By Lemma 2.7.10, G counts two-face, face-labelled planar maps with a tail. Theorem 2.7.11 gives an edge-preserving bijection between maps of this type of any fixed descent class $(\theta_1, \theta_2) \models n$, and edge-labelled pairs $(\lambda, \mathscr{P})$ such that λ is an edge and $\mathscr{P}$ is an ordered path on n vertices. Thus $G(z_1, z_2, u) = F(z, u) \circ \Delta^+(z\,;\, z_1, z_2)$, where $F(z, u)$ is the series counting pairs $(e, \mathscr{P})$ with respect to vertices of $\mathscr{P}$, marked by z, and labelled edges, marked by u. Clearly $F(z, u) = u \cdot z/(1 - uz)$, so Lemma 1.3.3 gives

$$G(z_1, z_2, u) = \frac{1}{z_1 - z_2}\left(z_2 \frac{uz_1}{1 - uz_1} - z_1 \frac{uz_2}{1 - uz_2}\right)$$

The result follows upon simplification. □

Finally, we mention that Theorem 2.7.11 could be stated in a more simple form than the one we have chosen here. In particular, the theorem obviously gives a bijection between permutations on n symbols and planar maps of fixed descent class $(\theta_1, \theta_2) \models n$ with a tail. Thus there are $n!$ such maps, and Corollary 2.7.13 is obtained immediately by noting that the coefficient of $u^n z_1^{\theta_1} z_2^{\theta_2}/n!$ in the series on the right-hand side of (2.43) is also $n!$. Our presentation has been chosen with the generalizations of Chapter 3 in mind.

2.7.4 Three-Face Smooth Planar Maps

Theorem 2.7.14, below, describes a bijection between ordered paths and three-face smooth planar maps of fixed descent class. The structure of this bijection is very similar to that given in Theorem 2.7.11, but its many cases make it more complicated.

Theorem 2.7.14. *Fix $\theta = (\theta_1, \theta_2, \theta_3) \models n$. There is a bijection between smooth face-labelled planar maps of descent class θ and edge-labelled tuples $(\lambda, \mathscr{P}, \gamma)$, where λ, γ are distinct edges and $\mathscr{P}$ is an ordered path containing n vertices.*

Proof. Let $\mathcal{M}_\theta$ be the set of smooth face-labelled planar maps of descent class θ, and let $\mathcal{P}_n$ be the set of all tuples $(\lambda, \mathscr{P}, \gamma)$ as described in the theorem. We define $\Omega_\theta : \mathcal{P}_n \longrightarrow \mathcal{M}_\theta$ through the following construction. Proofs of claims made within the construction can be found after the main proof.

Given $(\lambda, \mathscr{P}, \gamma) \in \mathcal{P}_n$, let F be the single face of $\mathscr{P}$ and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be its boundary walk, where v_0 and v_m are the white and grey ends of $\mathscr{P}$, respectively. Of course, F has n descents. For $0 \le i \le k$, let $c_i = (e_{i-1}, v_i, e_i)$. The situation is illustrated below.

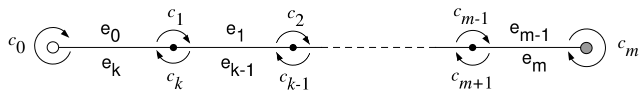

Notice that we trivially have $c_0 \in \mathcal{A}_F(\lambda)$ and $c_m \in \mathcal{A}_F(\gamma)$. Now define

$$I_1 = \{0 < i \leq m-1 \,:\, c_i \in \mathcal{A}_F(\lambda)\} \quad \text{and} \quad \epsilon_1 = \begin{cases} 1 & \text{if } (e_{m-1}, \lambda, \gamma)^\circ \text{ is nondecreasing,} \\ 0 & \text{otherwise,} \end{cases} \tag{2.44}$$

and

$$I_2 = \{m < i \leq k \,:\, c_i \in \mathcal{A}_F(\gamma)\} \quad \text{and} \quad \epsilon_2 = \begin{cases} 1 & \text{if } (e_k, \gamma, \lambda)^\circ \text{ is nondecreasing,} \\ 0 & \text{otherwise.} \end{cases} \tag{2.45}$$

Then $|I_1| + \epsilon_1 + |I_2| + \epsilon_2 = n - 1$ (*Claim 1*). Since $\theta_1 + \theta_2 + \theta_3 = n$, we have only the following two cases to consider.

Case 1: *Suppose* $|I_1| + \epsilon_1 \geq \theta_1$.

Since $c_0 \in \mathcal{A}_F(\lambda)$, Lemma 2.7.6 ensures a unique $c_r \in \mathcal{A}_F(\lambda)$ such that $\mathscr{P} \oplus (c_0, c_r)^\lambda_{1,2}$ is of descent class $(\theta_1, n-\theta_1) = (\theta_1, \theta_2+\theta_3)$. In fact, we have $r \leq m$ (*Claim 2*.) Let $\mathscr{N} = \mathscr{P} \oplus (c_0, c_r)^\lambda_{1,2}$. The construction of $\mathscr{N}$ from $\mathscr{P}$ in cases $r < m$ and $r = m$ is illustrated below.

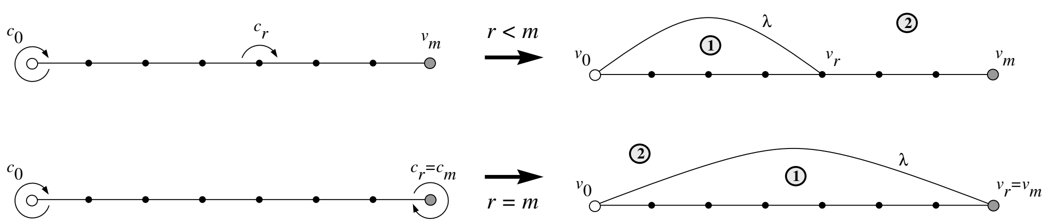

Let G denote face 2 of $\mathscr{N}$. Then G has boundary walk $((v_r, e_r), \ldots, (v_k, e_k), (v_0, \lambda))^\circ$ and $\theta_2 + \theta_3$ descents. We now consider cases $r < m$ and $r = m$ separately.

If $r < m$, then $c_m \in \mathcal{A}_G(\gamma)$ and Lemma 2.7.6 ensures a unique $c \in \mathcal{A}_G(\gamma)$ such that $\mathscr{N} \oplus (c_m, c)^\gamma_{2,3}$ is of descent class $(\theta_1, \theta_2, \theta_3)$. Strip the colour from v_0 and v_m, and let $\Omega_\theta(\lambda, \mathscr{P}, \gamma) = \mathscr{N} \oplus (c_n, c)^\gamma_{2,3}$. This construction is shown below.

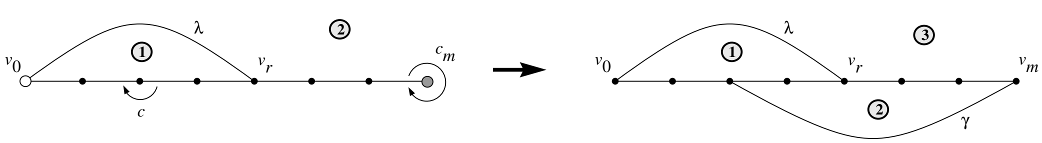

If $r = m$, then $c'_m = (\lambda, v_m, e_m)$ is a corner of G. Moreover, $c'_m \in \mathcal{A}_G(\gamma)$ (*Claim 3.*) Lemma 2.7.6 therefore gives a unique $c \in \mathcal{A}_G(\gamma)$ such that $\mathscr{N} \oplus (c'_m, c)^{\gamma}_{2,3}$ is of descent class $(\theta_1, \theta_2, \theta_3)$. Strip v_0 and v_m of their colour, and let $\Omega_\theta(\lambda, \mathscr{P}, \gamma) = \mathscr{N} \oplus (c'_m, c)^{\gamma}_{2,3}$. The construction is illustrated below.

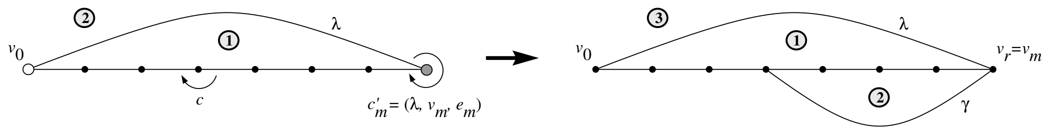

Case 2: *Suppose* $|I_1| + \epsilon_1 < \theta_1$ *and* $|I_2| + \epsilon_2 > \theta_2$.

Since $c_m \in \mathcal{A}_F(\gamma)$, Lemma 2.7.6 guarantees a unique $c_r \in \mathcal{A}_F(\gamma)$ such that $\mathscr{P} \oplus (c_m, c_r)^{\gamma}_{2,1}$ is of descent class $(n - \theta_2, \theta_2) = (\theta_1 + \theta_3, \theta_2)$. In fact, we have $m < r \leq k$ (*Claim 4.*) Let $\mathscr{N} = \mathscr{P} \oplus (c_m, c_r)^{\gamma}_{2,1}$. The construction of $\mathscr{N}$ from $\mathscr{P}$ is illustrated below.

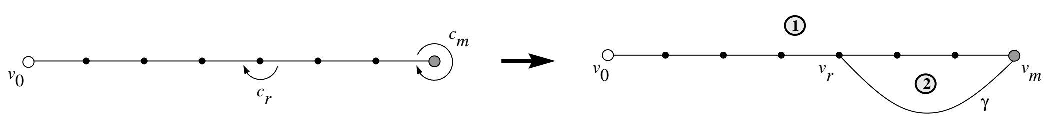

Let G denote face 1 of $\mathscr{N}$. Then G has boundary walk $((v_r, e_r), \ldots, (v_{m-1}, e_{m-1}), (v_m, \gamma))^\circ$ and $\theta_1 + \theta_3$ descents. Since $c_0 \in \mathcal{A}_G(\lambda)$, Lemma 2.7.6 ensures a unique $c \in \mathcal{A}_G(\lambda)$ such that $\mathscr{N} \oplus (c_0, c)^{\lambda}_{1,3}$ is of descent class $(\theta_1, \theta_2, \theta_3)$. In fact, either $c = (\gamma, v_r, e_r)$ or $c = c_s$ for some s with $r < s \leq 2k$ (*Claim 5.*) Strip v_0 and v_m of their colours, and let $\Omega_\theta(\lambda, \mathscr{P}, \gamma) = \mathscr{N} \oplus (c_0, c)^{\lambda}_{1,3}$. The creation of $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$ from $\mathscr{N}$ is shown below.

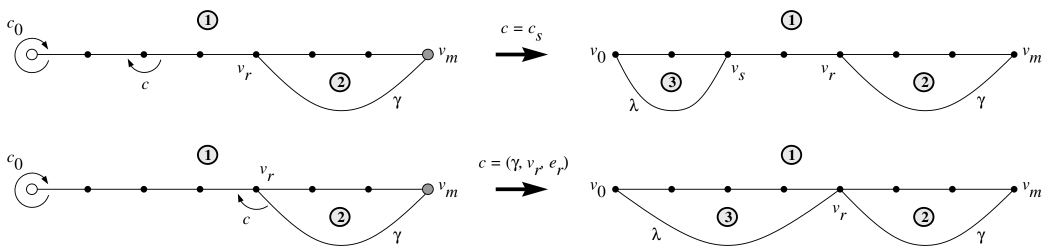

Analysis: Clearly $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$, as constructed above, is an element of $\mathcal{M}_\theta$ for all $(\lambda, \mathscr{P}, \gamma) \in \mathcal{P}_n$. We claim that $\Omega_\theta : \mathcal{P}_n \longrightarrow \mathcal{M}_\theta$ is a bijection. The proof is best described with the aid of Figure 2.26, which shows the ten disjoint classes of three-face smooth planar maps that can be produced by the construction. In particular, maps belonging to classes A through H are obtained through case 1 of the construction, while classes I and J correspond to case 2.

Let $\mathscr{M}$ be a face-labelled map belonging to one of classes A through H. Then the particular class of $\mathscr{M}$ can be determined, and the diagrams of Figure 2.26 unambiguously identify vertices

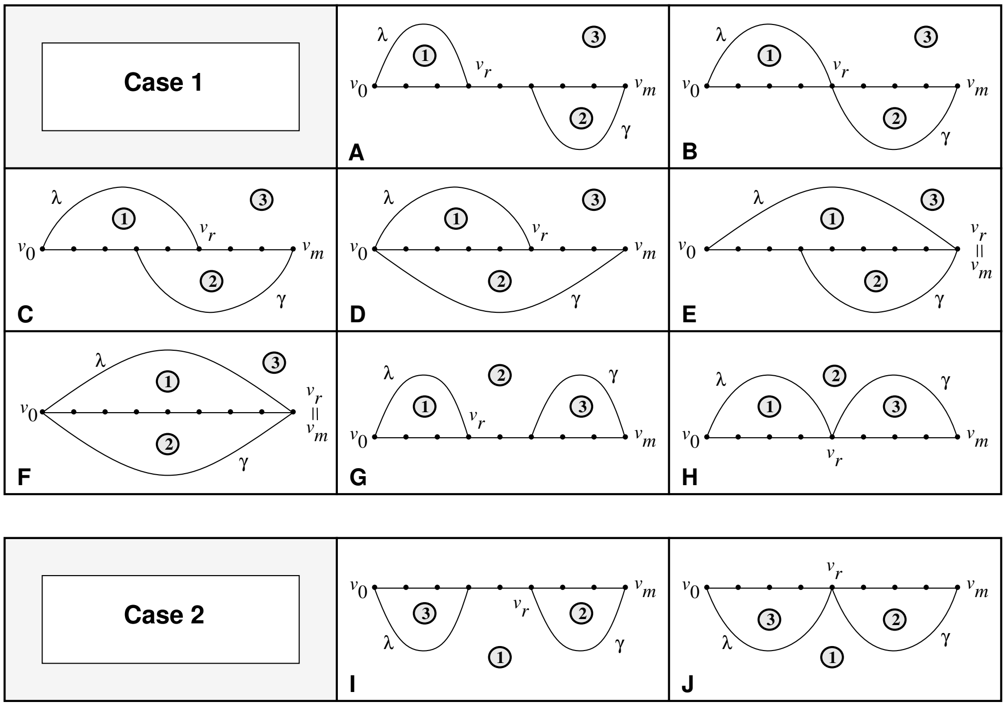

Figure 2.26: The maps produced by $\Omega_{(\theta_1,\theta_2,\theta_3)}$.

v_0, v_m and edges λ, γ of $\mathscr{M}$. This is clear from the fact that such maps lack automorphisms, but the identifications could be carried out practically, as follows. For a vertex v of $\mathscr{M}$, let L_v be the cyclic sequence of alternating face- and edge-labels encountered on a clockwise tour about v. If $\mathscr{M}$ has two vertices of degree 3, say u and v, such that $L_u = (1, a, 2, b, 3, c)^\circ$ and $L_v = (1, d, 3, e, 2, f)^\circ$ for some edges a, b, c, d, e, f, then it belongs to one of classes C through F, and we have $v_r = v$, $\gamma = b = \{u, v_m\}$ and $\lambda = d = \{v, v_0\}$. Classes C through F are now distinguished by equalities between u and v_0, and v and v_m. For example, class C is characterized by the conditions $u \neq v_0$ and $v \neq v_m$, while class D has $u = v_0$ and $v \neq v_m$. Similarly, if $\mathscr{M}$ has a vertex v of degree 4 such that $L_v = (1, a, 3, b, 2, c, 3, d)^\circ$, then $\mathscr{M}$ is of class B, and we have $\lambda = a = \{v, v_0\}$, $\gamma = c = \{v, v_m\}$.

The argument above shows that the construction of $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$ is reversible, hence Ω_θ is one-one. To see that Ω_θ is onto, first observe that every map $\mathscr{M} \in \mathcal{M}_\theta$ belongs to one of the classes of maps shown in Figure 2.26. The particular class of $\mathscr{M}$ is then uniquely determined as above, as are vertices v_0, v_m and edges λ, γ. Let $\mathscr{P}$ be the ordered path obtained by removing λ and γ from $\mathscr{M}$ and colouring v_0 white and v_m grey. Let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the boundary walk of the single face F of $\mathscr{P}$, and set $c_i = (e_{i-1}, v_i, e_i)$ for $0 \leq i \leq k$, so that $c_0 \in \mathcal{A}_F(\lambda)$ and $c_m \in \mathcal{A}_F(\gamma)$. Now define I_1, ϵ_1, I_2 and ϵ_2 as in (2.44) and (2.45).

Suppose $\mathscr{M}$ is of class A through H. Let $\mathscr{N}_0 = \mathscr{M}\backslash\gamma$, and let G be the face of $\mathscr{N}_0$ created by the amalgamation of faces 2 and 3 of $\mathscr{M}$ upon removal of γ. If G is assigned label 2, then $\mathscr{N}_0$ is of descent class $(\theta_1, \theta_2 + \theta_3)$, since $\mathscr{M} \in \mathcal{M}_\theta$. Moreover, face 1 of $\mathscr{N}_0$ has boundary walk $((v_0, e_0), \ldots, (v_{r-1}, e_{r-1}), (v_r, \lambda))^\circ$ for some r with $1 < r \leq m$, and we have $\mathscr{N}_0 = \mathscr{P} \oplus (c_0, c_r)^\lambda_{1,2}$. If $r < m$ then Lemma 2.7.6 implies $|I_1| \geq \theta_1$, whereas $r = m$ implies $|I_1| \geq \theta_1 - 1$. However, in the latter case, observe that $(e_{m-1}, \lambda, \gamma)^\circ$ is increasing, since it is a subsequence of the rotator of v_m in $\mathscr{M}$. Therefore $|I_1| + \epsilon_1 \geq \theta_1$ in either case. By the uniqueness guaranteed by Lemma 2.7.6, $\mathscr{N}_0$ coincides with the intermediary map $\mathscr{N}$ created in (case 1 of) the construction of $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$. Now set $c = c_m$ if $r < m$, and $c = (\lambda, v_m, e_m)$ if $r = m$. Then $\mathscr{M} = \mathscr{N}_0 \oplus (c, c')^\gamma_{2,3}$ for some corner $c' \in \mathcal{A}_G(\gamma)$. But corner c' is unique (by Lemma 2.7.6) and $\mathscr{N} = \mathscr{N}_0$, so comparison with the construction of $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$ reveals that $\mathscr{M} = \Omega_\theta(\lambda, \mathscr{P}, \gamma)$.

If $\mathscr{M}$ is of class I or J, then a similar argument proves $\mathscr{M} = \Omega_\theta(\lambda, \mathscr{P}, \gamma)$. (Here we find that $|I_2| + \epsilon_2 > \theta_2$, so case 2 of the construction is in effect.) Thus Ω_θ is onto, and the main proof is complete. Proofs of the supporting claims follow. □

Proof of Claim 1: Consider the two-face smooth planar map $\mathscr{Q}$ obtained from $\mathscr{P}$ by first augmenting the path with an edge $\{v_m, v\}$ labelled γ, and then attaching an edge $\{v, v_0\}$ labelled λ. This is illustrated below.

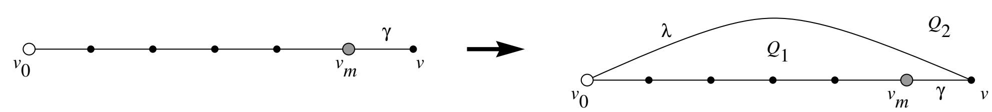

The two faces, Q_1 and Q_2, of $\mathscr{Q}$ have boundary walks

$$((v_0, e_0), \ldots, (v_{m-1}, e_{m-1}), (v_m, \gamma), (v, \lambda))^\circ \qquad \text{and} \qquad ((v_m, e_m), \ldots, (v_k, e_k), (v_0, \lambda), (v, \gamma))^\circ,$$

respectively. Note that $\mathcal{A}_{Q_1}(\lambda) = \{(\lambda, v_0, e_0), (e_{m-1}, v_m, \gamma)\} \cup \{(e_{i-1}, v_i, e_i) : i \in I_1\}$ if $(e_{m-1}, \lambda, \gamma)^\circ$ is nondecreasing, and $\mathcal{A}_{Q_1}(\lambda) = \{(\lambda, v_0, e_0)\} \cup \{(e_{i-1}, v_i, e_i) : i \in I_1\}$ otherwise. Thus $|\mathcal{A}_{Q_1}(\lambda)| = 1 + |I_1| + \epsilon_1$. Similarly, we get $|\mathcal{A}_{Q_2}(\gamma)| = 1 + |I_2| + \epsilon_2$. But, by Lemma 2.7.6, faces Q_1 and Q_2 have $|\mathcal{A}_{Q_1}(\lambda)|$ and $|\mathcal{A}_{Q_2}(\gamma)|$ descents, respectively. Since $\mathscr{Q}$ has $n + 1$ vertices, this gives $|\mathcal{A}_{Q_1}(\lambda)| + |\mathcal{A}_{Q_2}(\gamma)| = n + 1$. Thus $|I_1| + \epsilon_1 + |I_2| + \epsilon_2 = n - 1$. □

Proof of Claim 2: Since $c_m \in \mathcal{A}_F(\lambda)$, Lemma 2.7.6 implies $\mathscr{P} \oplus (c_0, c_m)^\lambda_{1,2}$ has descent class $(d_1, d_2) \models n$, where $d_1 = |I_1| + 1$. If $r > m$, then $d_1 < \theta_1$ and we get the contradiction $|I_1| + \epsilon_1 \leq$

$|I_1| + 1 < \theta_1$. Thus $r \leq m$, as claimed. In fact, if $r = m$ then $j = \theta_1$ and thus $\theta_1 \leq |I_1| + \epsilon_1 \leq |I_1| + 1 = \theta_1$. Therefore $r = m$ implies $\epsilon_1 = 1$.

Proof of Claim 3: In the proof of Claim 2 we showed that $r = m$ implies $\epsilon_1 = 1$. But this implies $(e_{m-1}, \lambda, \gamma)^\circ$ is nondecreasing. We also know that $(e_{m-1}, \gamma, e_m)^\circ$ is nondecreasing, and together these conditions force $(\lambda, \gamma, e_m)^\circ$ to be nondecreasing. Hence $(\lambda, v_m, e_m) \in \mathcal{A}_G(\gamma)$. □

Proof of Claim 4: Since $c_0 \in \mathcal{A}_F(\gamma)$, Lemma 2.7.6 implies $\mathscr{P} \oplus (c_m, c_0)^{\lambda}_{2,1}$ has descent class $(d_1, d_2) \models n$, where $d_2 = |I_2| + 1$. But if $0 \leq s \leq m$, then $d_2 \leq \theta_2$ and we get the contradiction $|I_2| + \epsilon_2 \leq |I_2| + 1 \leq \theta_2$. Thus $m < s \leq k$, as claimed. □

Proof of Claim 5: Clearly c must be one of (e_{m-1}, v_m, γ), or (γ, v_r, e_r), or (e_{s-1}, v_s, e_s) for a unique s with $0 < s < m$ or $r < s \leq k$, as these are all the corners of G aside from c_0. If $c = (e_{m-1}, v_m, \gamma)$, then $c \in \mathcal{A}_G(\lambda)$ implies $(e_{m-1}, \lambda, \gamma)^\circ$ is nondecreasing, so that $\epsilon_1 = 1$. Then since $\mathscr{N} \oplus (c_0, c)^{\lambda}_{1,3}$ has descent class $(\theta_1, \theta_2, \theta_3)$, Lemma 2.7.6 implies $\theta_1 = |I_1| + 1 = |I_1| + \epsilon_1$, which contradicts $|I_1| + \epsilon_1 < \theta_1$. Similarly, if $c = (e_{s-1}, v_s, e_s)$ with $0 < s < m$, then $|I_1| \geq \theta_1$, again a contradiction. Therefore $c = (\gamma, v_r, e_r)$ or $c = (e_{s-1}, v_s, e_s)$ for some s with $r < s \leq k$. □

Example 2.7.15. Let $\mathscr{P}$ be the ordered path shown in Figure 2.27A, and let $\lambda = 6, \gamma = 2$. The single face F of $\mathscr{P}$ has boundary walk $((v_0, e_0), \ldots, (v_{13}, e_{13}))^\circ$, where v_0 and v_7 are the white and grey ends of $\mathscr{P}$, respectively, and

$$(e_0, e_1, \ldots, e_{13}) = (4, 8, 3, 5, 9, 7, 1, 1, 7, 9, 5, 3, 8, 4).$$

Thus $m = 7$ and $k = 13$. Let $c_i = (e_{i-1}, v_i, e_i)$ for $0 \leq i \leq 13$. Then $\mathcal{A}_F(\lambda) = \{c_{i_0}, \ldots, c_{i_7}\}$ and $\mathcal{A}_F(\gamma) = \{c_{j_0}, \ldots, c_{j_7}\}$, where

$$(i_0, \ldots, i_7) = (0, 1, 4, 5, 7, 8, 11, 12) \quad \text{and} \quad (j_0, \ldots, j_7) = (0, 2, 5, 7, 8, 10, 11, 13).$$

Thus $I_1 = \{1, 4, 5\}$ and $I_2 = \{8, 10, 11, 13\}$. Neither $(e_m, \lambda, \gamma)^\circ = (1, 6, 2)^\circ$ nor $(e_k, \gamma, \lambda)^\circ = (4, 2, 6)^\circ$ is nondecreasing, so $\epsilon_1 = \epsilon_2 = 0$. Therefore $|I_1| + \epsilon_1 = 3$ and $|I_2| + \epsilon_2 = 4$.

If $\theta = (3, 2, 3)$, then $|I_1| + \epsilon_1 \geq \theta_1$ and we follow Case 1 to construct $\Omega_\theta(\lambda, \mathscr{P}, \gamma)$. The process is illustrated in Figure 2.27B. First observe that $\mathscr{P} \oplus (c_0, c_r)^{\lambda}_{1,2}$ is of descent class $(3, 5)$ only for $r = 5$. Therefore $\mathscr{N} = \mathscr{P} \oplus (c_0, c_5)^{6}_{1,2}$. Now notice that $r < m$, and that $\mathscr{N} \oplus (c_m, c)^{\gamma}_{2,3}$ is of descent class $\theta = (3, 2, 3)$ precisely when c is corner c_{10} of $\mathscr{N}$. Thus $\Omega_\theta(\lambda, \mathscr{P}, \gamma) = \mathscr{N} \oplus (c_7, c_{10})^{2}_{1,3}$.

If $\theta = (4, 3, 1)$, then $|I_1| + \epsilon_1 < \theta_1$ and $|I_2| + \epsilon_2 > \theta_2$. We therefore follow Case 2 to construct

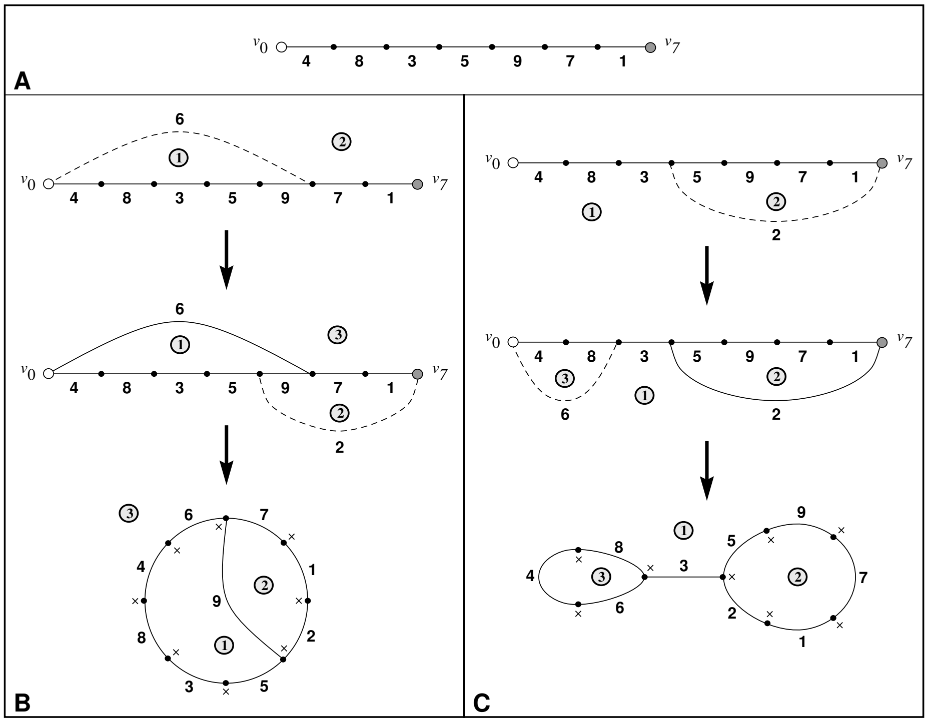

Figure 2.27: Constructing smooth three-face planar maps from ordered paths.

$\Omega_\theta(\lambda, \mathscr{P}, \gamma)$, as shown in Figure 2.27C. Since $\mathscr{P} \oplus (c_m, c_r)^{\gamma}_{2,1}$ is of descent class $(5, 3)$ only for $r = 11$, we have $\mathscr{N} = \mathscr{P} \oplus (c_7, c_{11})^2_{2,1}$. Since $\mathscr{N} \oplus (c_0, c_{12})^{\lambda}_{1,3}$ is of descent class $\theta = (4, 3, 1)$, we have $\Omega_\theta(\lambda, \mathscr{P}, \gamma) = \mathscr{N} \oplus (c_0, c_{12})^6_{1,3}$. □

Let $\mathscr{M}$ be the smooth map constructed in Figure 2.27B, and let F_i denote face i of this map, for $i = 1, 2, 3$. Then the alternating cyclic lists of faces and edges encountered on clockwise tours about the vertices of $\mathscr{M}$ of degree 3 are $(F_1, 6, F_3, 7, F_2, 9)^\circ$ and $(F_1, 9, F_2, 2, F_3, 5)^\circ$. To reverse the construction we compare these statistics with the diagrams of Figure 2.26. This identifies $\mathscr{M}$ as being in class C, with $\lambda = 6$ and $\gamma = 2$. Colour the ends of λ and γ that are not of degree 3 white and grey, respectively. Then removing edges 2 and 6 from $\mathscr{M}$ produces the original ordered path. □

Theorem 2.7.14 provides a combinatorial proof of the following result, which is the $m = 3$ case of Theorem 2.6.11.

Corollary 2.7.16.

$$\Gamma_3(z_1, z_2, z_3, u) = \frac{u^4 z_1 z_2 z_3}{(1 - uz_1)(1 - uz_2)(1 - uz_3)}. \qquad (2.46)$$

Proof. Theorem 2.7.14 shows smooth face-labelled planar maps of fixed descent class $(\theta_1, \theta_2, \theta_3) \models n$ to be in edge-preserving bijection with edge-labelled structures $(\lambda, \mathscr{P}, \gamma)$, where λ, γ are distinct edges and $\mathscr{P}$ is an ordered path on n vertices. Thus $\Gamma_3(z_1, z_2, z_3, u) = F(z, u) \circ \Delta^+(z\,;\, z_1, z_2, z_3)$, where $F(z, u)$ is the series counting tuples $(e_1, \mathscr{P}, e_2)$ with respect to vertices of $\mathscr{P}$, marked by z, and labelled edges, marked by u. Since $F(z, u) = u^2 \cdot z/(1 - uz)$, Lemma 1.3.3 gives

$$\Gamma_3(z_1, z_2, z_3, u) = \sum_{i=1}^{3} \frac{u^2 z_i}{1 - u z_i} \prod_{\substack{1 \le j \le 3 \\ j \ne i}} \frac{z_j}{z_i - z_j}.$$

The result follows upon simplification. □

Again, we note that Theorem 2.7.14 can be stated simply as a bijection between permutations on $n + 1$ symbols and smooth planar maps of fixed descent class $(\theta_1, \theta_2, \theta_3) \models n$. There are therefore $(n + 1)!$ such maps, and Corollary 2.7.16 results by comparing coefficients on both sides of (2.46). The motivation behind our approach will become clear in Chapter 3.

2.7.5 A Differential Decomposition for Smooth Planar Maps

The bijections given in Theorems 2.7.11 and 2.7.14 share a common theme. That is, a smooth map of predetermined descent class is built from an ordered path by attaching labelled edges to its endpoints. Lemma 2.7.6 plays a fundamental rôle in these constructions, guaranteeing that edges of any given label can be attached in a unique manner to create a map of the desired descent class. Unfortunately, we have been unable to extend this method to give similar constructions for smooth maps with more than three faces. Thus a combinatorial proof of Theorem 2.6.11 when $m > 3$ is currently beyond reach.

However, Lemma 2.7.6 can be used to develop a recursive decomposition for smooth planar maps. This is done in the next theorem, where the result is stated in the form of a differential equation satisfied by $\Gamma_m(\mathbf{z}, u)$, for $m \geq 4$. In fact, through the identity $\Gamma_m(\mathbf{w}, u) = \Psi_m(\mathbf{x}, u)$, this result is equivalent to Theorem 2.5.6. A positive genus analogue is also readily obtained by the methods used here.

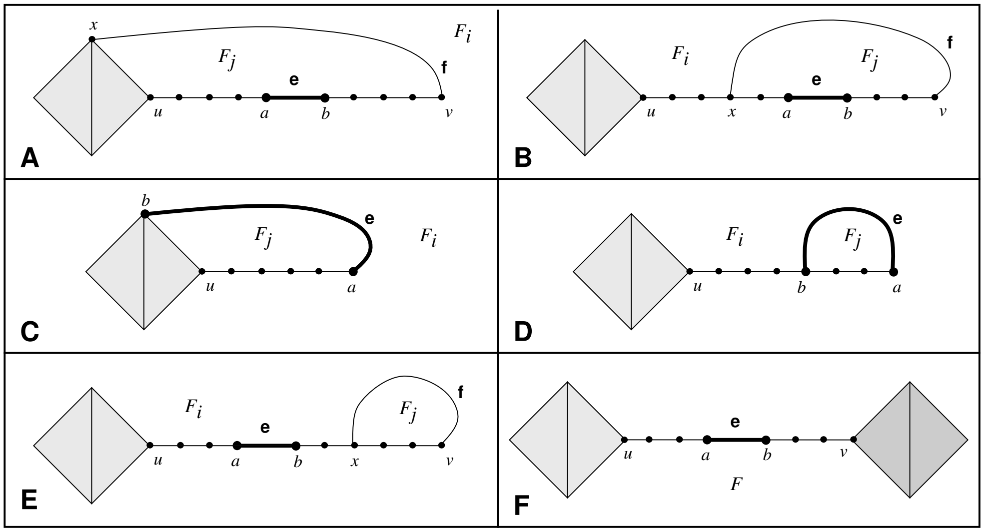

Figure 2.28: Decompositions of a smooth planar map.

Theorem 2.7.17. *Fix $m \geq 4$. For any subset $\lambda = \{\lambda_1, \ldots, \lambda_k\} \subseteq [m]$, where $\lambda_1 < \cdots < \lambda_k$, let $\mathbf{z}_\lambda = (z_{\lambda_1}, \ldots, z_{\lambda_k})$. For $1 \leq i \leq m$, set $\overline{\mathbf{z}}_i = \mathbf{z}_{[m]\setminus\{i\}}$. Also, for each i, let*

$$\partial_i = \frac{z_i}{1 - uz_i}\frac{\partial}{\partial z_i} \tag{2.47}$$

and let $\mathcal{P}_i$ be the set of all pairs $\{\gamma, \lambda\}$ with $\gamma, \lambda \subset [m]$ such that $\gamma \cap \lambda = \{i\}$ and $\gamma \cup \lambda = [m]$. Then

$$\frac{\partial}{\partial u}\Gamma_m(\mathbf{z}, u) = \sum_{i=1}^{m} \sum_{\substack{\{\gamma,\lambda\}\in\mathcal{P}_i \\ |\gamma|,|\lambda|\geq 3}} \partial_i \Gamma_{|\gamma|}(\mathbf{z}_\gamma, u) \cdot \partial_i \Gamma_{|\lambda|}(\mathbf{z}_\lambda, u) + \sum_{\substack{1\leq i,j\leq m \\ i\neq j}} \frac{z_j \partial_i \Gamma_{m-1}(\overline{\mathbf{z}}_j, u)}{(1 - uz_i)(z_i - z_j)}. \tag{2.48}$$

Proof. The series on the left-hand side of (2.48) counts all possible structures obtained by deleting the maximally labelled edge from a smooth face-labelled map with m faces. We show that the series on the right-hand side of (2.48) counts these same structures. To this end, let $\mathscr{M}$ be a smooth face-labelled planar map with $m \geq 4$ faces, and let $e = \{a, b\}$ be its maximal edge.

Since $\mathscr{M}$ is smooth, the various diagrams of Figure 2.28 illustrate the six possible configurations of e within $\mathscr{M}$. In each diagram, the shaded squares represent smooth maps with at least two internal faces (note that this distinguishes case E from F). Vertices u and a, as defined in the diagrams, may coincide in all cases except D. The identity $b = v$ is possible in cases A, B, and F, while $x = b$ is possible in E. With the aid of these diagrams, now consider the effect of deleting e from $\mathscr{M}$.

First examine case F. Let the sole face F incident with e have label i. In this case, removal of e decomposes $\mathscr{M}$ into two planar maps, $\mathscr{M}_a$ and $\mathscr{M}_b$, containing a and b, respectively, whose faces are labelled with $\lambda \subseteq [m]$ and $\gamma \subseteq [m]$, where $\lambda \cap \gamma = \{i\}$, $\lambda \cup \gamma = [m]$, and $|\gamma|, |\lambda| > 2$. Let F_a and F_b be the faces of $\mathscr{M}_a$ and $\mathscr{M}_b$ labelled i. This decomposition is illustrated below. Note that u, v, a, b are distinguished by the deletion process, as has been indicated by colouring these vertices in the diagram.

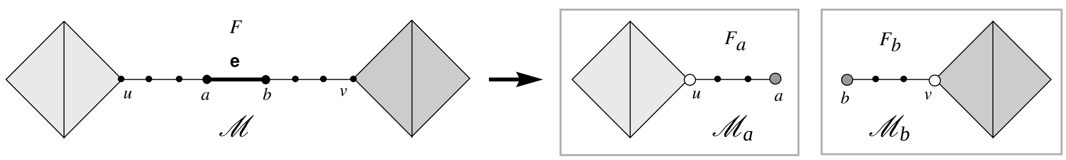

Since e is maximal, both a and b are at descents of F. Therefore u is at a descent of F_a when $u = a$, and v is at a descent of F_b when $v = b$. Thus $\mathscr{M}_a$ has a tail in F_a, and $\mathscr{M}_b$ has a tail in F_b. By Lemma 2.7.10 and (2.47), the series counting all possible pairs $(\mathscr{M}_a, \mathscr{M}_b)$ is

$$\partial_i \Gamma_{|\gamma|}(\mathbf{z}_\gamma, u) \cdot \partial_i \Gamma_{|\lambda|}(\mathbf{z}_\lambda, u).$$

Summing over i and over permissible pairs $\{\lambda, \gamma\}$ gives the first sum on the right-hand side of (2.48).

Now consider cases A, B, and E simultaneously. Assume faces F_i and F_j are labelled i and j, respectively. Let $\mathscr{N}$ be the planar map with $m-1$ faces resulting from the separation of edge $f = \{v, x\}$ from $\mathscr{M}$. Let F_0 be the face of $\mathscr{N}$ created by the merger of F_i and F_j, and assign label 0 to this face. A vertex is at a descent of F_0 if and only if it is at a descent of F_i or F_j. If F_0 has n descents, then F_i and F_j have d and $n-d$ descents, respectively, for some d with $1 \le d \le n-1$. The construction of $\mathscr{N}$ is shown below.

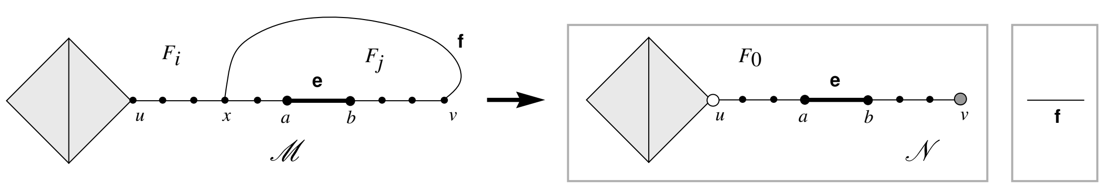

Lemma 2.7.6 implies that the position of edge f within $\mathscr{M}$ is uniquely determined by its label, one of its endpoints (that is, v or x), and the number d of descents of F_i. But v can always be identified as the only leaf of $\mathscr{N}$. Thus $\mathscr{M}$ can be recovered from $\mathscr{M}$, d, and f alone. Deletion of e from $\mathscr{M}$ is therefore equivalent to deletion of e from $\mathscr{N}$, provided that d and f are recorded. Finally, note

that deleting e from $\mathscr{N}$ is a reversible operation that results in a map $\mathscr{N}_a$ and an ordered path $\mathscr{P}$, where $\mathscr{N}_a$ has $m-1$ faces and a tail in face 0, and $\mathscr{P}$ may be of length 0. This is illustrated below.

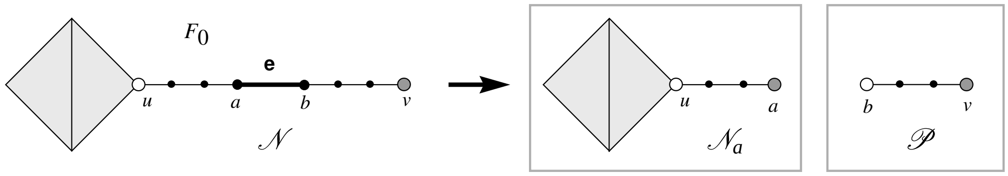

A similar, but less complicated, analysis is valid for cases C and D. Here, deletion of e from $\mathscr{M}$ creates a new map $\mathscr{M}_a$ with $m-1$ faces. Again, faces F_i and F_j (assumed to be labelled i and j, respectively) are merged into a single face F_0 of $\mathscr{M}_a$, and this new face is given label 0. Note that $\mathscr{M}_a$ has a tail in F_0. As before, if F_0 has n descents, then F_i and F_j have d and $n-d$ descents, respectively, where $1 \leq d \leq n-1$. Lemma 2.7.6 shows that $\mathscr{M}$ can be recovered from $\mathscr{M}_a$ together with the number d of descents of F_i.

Let $\overline{\mathbf{z}}_{ij} = \mathbf{z}_{[m]\setminus\{i,j\}}$. By Lemma 2.7.10, the series

$$\frac{z}{1-uz}\frac{\partial}{\partial z}\Gamma_{m-1}(z, \overline{\mathbf{z}}_{ij}, u)$$

counts maps such as $\mathscr{N}_a$ or $\mathscr{M}_a$ that have $m-1$ faces labelled with $\{0, \ldots, m\}\setminus\{i, j\}$ and a tail in face 0. Here z marks descents of face 0, and u marks labelled edges, as usual. The series $z/(1-uz) \cdot u$ counts pairs $(\mathscr{P}, f)$, where $\mathscr{P}$ is an ordered path and f is a labelled edge. Let

$$G(z, \overline{\mathbf{z}}_{ij}, u) = \left(\frac{z}{1-uz}\frac{\partial}{\partial z}\Gamma_{m-1}(z, \overline{\mathbf{z}}_{ij}, u)\right) \cdot \left(\frac{uz}{1-uz} + 1\right).$$

Then, from our above analysis, the series counting all possible structures $\mathscr{M}\setminus e$ arising in cases A through E is obtained from G by replacing x^n with $\sum_{d=1}^{n-1} x_i^d x_j^{n-d}$, for all $n \geq 1$. From (1.1) and Lemma 1.3.3, this gives

$$G(z, \overline{\mathbf{z}}_{ij}, u) \circ \Delta^+(z\,;\, z_i, z_j) = \frac{z_j G(z_i, \overline{\mathbf{z}}_{ij}, u) - z_i G(z_j, \overline{\mathbf{z}}_{ij}, u)}{z_i - z_j}.$$

Observe that $G(z_i, \overline{\mathbf{z}}_{ij}, u) = (1-uz_i)^{-1}\partial_i \Gamma_{m-1}(\overline{\mathbf{z}}_j, u)$ and $G(z_j, \overline{\mathbf{z}}_{ij}, u) = (1-uz_j)^{-1}\partial_j \Gamma_{m-1}(\overline{\mathbf{z}}_i, u)$. Summing over all pairs $\{i, j\} \subseteq [m]$ therefore results in the second summation on the right-hand side of (2.48). □

Notice that the proof of Theorem 2.7.17 sheds some light on the mysterious algebraic cancellation that occurred in the earlier derivation of Theorem 2.5.6. In particular, the cancellation is reflected by our separate treatments of cases E and F (see Figure 2.28). In both cases, e is a bridge, so its deletion separates $\mathscr{M}$ into two maps, but the analysis above shows that case E can be naturally grouped with cases A through D rather than with case F.

2.8 Bijections Between Factorizations and Trees

In §2.4.7 we described a bijection between minimal factorizations of full cycles and vertex-labelled trees. This gave a bijective proof of Dénes result, $H_0((n)) = n^{n-2}$, which is the special case of the Hurwitz formula for $H_0(\alpha)$ when $\ell(\alpha) = 1$. We conclude Chapter 2 with analogous correspondences for minimal factorizations of permutations composed of two or three disjoint cycles, thereby providing combinatorial proofs of the Hurwitz formula when $\ell(\alpha) = 2$ and $\ell(\alpha) = 3$.

The framework for these results is already complete. In fact, the correspondences given in this section are obtained simply by composing a close relative of the tree pruning bijection with the bijections defined in Theorems 2.7.11 and 2.7.14.

2.8.1 Preliminaries

In order to describe the forthcoming correspondences cleanly, we must introduce some minor generalizations of earlier definitions. We begin with the normal indexing of boundary walks. As mentioned in §2.6.2, our earlier definition in the context of vertex-labelled maps was devised with an extension to all maps in mind.

We say a map is **trivial** if it has exactly two vertices and one face. The following technical lemma allows for a well-defined normal indexing of the boundary walks of a nontrivial map.

Lemma 2.8.1. *Let $W = ((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the boundary walk of face F in the map $\mathscr{M}$. For $0 \le i \le k$, define $L_i = (e_i, e_{i+1}, \ldots, e_{i+k}) \in \mathbb{Z}^{k+1}$. If $\prec$ is the strict lexicographic order on $\mathbb{Z}^{k+1}$, then either $\mathscr{M}$ is trivial or there is a unique i with $0 \le i \le k$ such that $L_i = \min_\prec\{L_j : 0 \le j \le k\}$.*

Proof. If there is no such value of i, then $(e_0, \ldots, e_k) = (e_0, \ldots, e_j, e_0, \ldots, e_j, \cdots, e_0, \ldots, e_j)$ for some $j \ge 0$, where there are $m \ge 2$ copies of the sequence $e_0, \ldots, e_j$ in the latter list. This shows W to be of the form

$$((v_0, e_0), \ldots, (v_j, e_j), (v_{j+1}, e_0), \ldots, (v_{2j+1}, e_j), \quad \cdots \quad (v_{(m-1)(j+1)}, e_0), \ldots, (v_{m(j+1)-1}, e_j))^\circ.$$

Thus $e_i = \{v_i, v_{i+1}\} = \{v_{i+j+1}, v_{i+j+2}\}$ for all i. But $v_{i+1} \ne v_{i+j+2}$, since otherwise corners (e_i, v_{i+1}, e_{i+1}) and $(e_i, v_{i+j+2}, e_{i+1})$ of W would be identical. It follows that $v_i = v_{i+j+2}$ and

$v_{i+1} = v_{i+j+1}$. These identities combine to give $v_{i+2} = v_i$ for all i. Since $v_0 \neq v_1$, we conclude that W is incident with exactly two vertices. Moreover, each edge of W is encountered at least twice (hence exactly twice), implying that F is the only face of $\mathscr{M}$. □

In light of the lemma, we say that a boundary walk $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ in a nontrivial map is **normally indexed** if $(e_0, \ldots, e_k) = \min_\prec\{(e_i, \ldots, e_{i+k}) : 0 \le i \le k\}$. Observe that this definition of normal indexing is compatible with Definition 2.6.3, in the sense that they agree for nontrivial vertex- and edge-labelled maps. (No effective definition of normal indexing can be given for trivial maps, since their two vertices are interchangeable.)

Let $\mathscr{M}$ be any nontrivial map, and let vertex v be at one of the d descents of face F of $\mathscr{M}$. Let $((e_0, v_0), \ldots, (e_k, v_k))^\circ$ be the normally indexed boundary walk of F, and let $j_1 < \cdots < j_d$ be the d values of j with $0 \le j \le k$ such that v_j is at a descent of F. Then we have $v_{j_i} = v$ for a unique $i \in \{1, \ldots, d\}$. We call this value of i the **index** of vertex v in F.

The definitions of cores and branches given in §2.6.1 can be generalized, as follows. We define a **submap** of the map $\mathscr{M}$ to be any map that can be obtained from $\mathscr{M}$ by successive removal of leaves (and their incident edges). Thus the **core** of a map is its minimal submap. Let $\mathscr{N}$ be a submap of $\mathscr{M}$, and let $e = \{u, v\}$ be an edge of $\mathscr{M}$ such that $\mathscr{N}$ contains v but not u. Then e is incident with only one face, F, of $\mathscr{M}$. Detaching e from v yields a rooted tree T whose root vertex is incident only with e. We call this tree an $\mathscr{N}$**-branch** of face F. Vertex v is the **base vertex** of T and edge e is its **stem**.

When $\mathscr{N}$ is nontrivial, the *index* of T is defined as follows. Let the face of $\mathscr{N}$ corresponding to F have d descents and normally indexed boundary walk $((v_0, e_0), \ldots, (v_k, e_k))^\circ$. By Lemma 2.6.5, $(e_{j-1}, e, e_j)^\circ$ is increasing for exactly d values of j in the range $0 \le j \le k$. Let these values of j be $j_1 < \cdots < j_d$. Then the **index** of T is the unique value of $i \in \{1, \ldots, d\}$ such that $v = v_{j_i}$.

Finally, we remark that the definition of normal indexing given here is compatible with that given earlier for vertex-labelled maps. That is, the normal indexing of a boundary walk in a nontrivial vertex-labelled map is the same whether our current definition (disregarding vertex labels) or Definition 2.6.3 is used. Of course, this implies that our two definitions of the index of a branch are also compatible.

Example 2.8.2. Consider the face-labelled map $\mathscr{M}$ drawn in Figure 2.29A. Note that the highlighted vertex v is at a descent of face 1. The normally indexed boundary walk of this face is

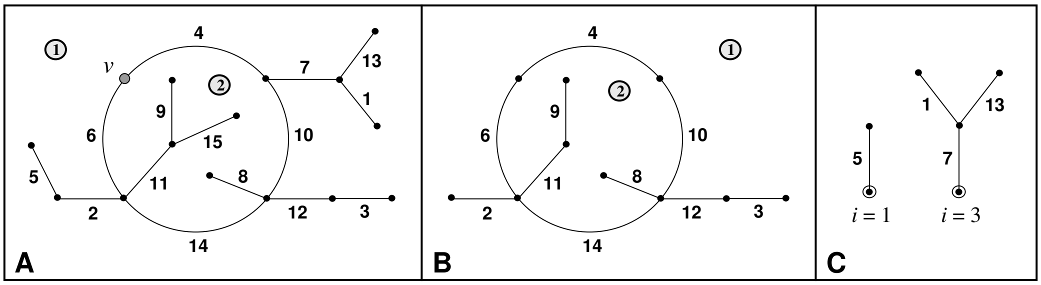

Figure 2.29: (A) A map $\mathscr{M}$, (B) a submap $\mathscr{N}$ of $\mathscr{M}$, and (C) the $\mathscr{N}$-branches of face 1 of $\mathscr{M}$.

$((v_0, e_0), \ldots, (v_{18}, e_{17}))^\circ$, where

$$(e_0, e_1, \ldots, e_{17}) = (1, 1, 7, 10, 12, 3, 3, 12, 14, 2, 5, 5, 2, 6, 4, 7, 13, 13).$$

Note that $e_{j-1} \geq e_j$ for exactly those j in the list $(j_1, \ldots, j_8) = (0, 1, 5, 6, 9, 11, 12, 14, 17)$. Since $v = v_{14} = v_{j_8}$, vertex v has index 8 in face 1.

A submap $\mathscr{N}$ of $\mathscr{M}$ is given in Figure 2.29B. Face 1 of this map has 5 descents, with normally indexed boundary walk $((u_0, f_0), \ldots, (u_8, f_8))^\circ$, where

$$(f_0, f_1, \ldots, f_9) = (2, 2, 6, 4, 10, 12, 3, 3, 12, 14).$$

Thus the $\mathscr{N}$-branches of face 1 of $\mathscr{M}$, and their indices, are as shown in Figure 2.29C. For instance, the branch with stem 7 has index 3, since $(f_{j-1}, 7, f_j)^\circ$ is nondecreasing for j in the list $(j_1, \ldots, j_5) = (1, 3, 4, 7, 8)$, and edge 7 of $\mathscr{M}$ is incident with vertex $u_4 = u_{j_3}$ of $\mathscr{N}$. □

2.8.2 Factorizations of Class (n_1, n_2)

A **dotted factorization** of the permutation $\pi \in \mathfrak{S}_n$ is a factorization of π together with a choice of a distinguished symbol $i \in [n]$. The special symbol is identified by marking it with a dot. For example, $(1\,3)(\dot{2}\,4)(1\,\dot{2})(1\,5)$ is a dotted factorization of $(1\,5\,4\,2\,3)$. Algorithm 2.8.3, below, transforms a minimal transitive dotted factorization of class (n_1, n_2) into a pair of doubly rooted vertex-labelled trees and a certain set partition of $[n_1 + n_2]$. Figure 2.30 on page 95 serves as a running example of the algorithm, illustrating each step as it is applied to the dotted factorization

$$(9\ 10)(8\ 16)(2\ 5)(1\ 12)(5\ 15)(5\ 13)(1\ 8)(8\ 11)(2\ 4)(8\ 10)(6\ 12)(2\ 3)(7\ 16)(1\ 13)(5\ 12)(13\ \dot{14})$$

of the permutation

$$(1\ 2\ 3\ 4\ 5\ 6)(7\ 8\ 9\ 10\ 11\ 12\ 13\ 14\ 15\ 16) \in \mathfrak{S}_{16}.$$

The panels of the figure correspond to similarly labelled steps of the algorithm.

Algorithm 2.8.3.

INPUT: A genus 0 dotted factorization f of $(1\,2\,\cdots\,n_1)(n_1+1\,\cdots\,n_1+n_2)$.

A. Let $\mathscr{M}$ be the face-labelled map with distinguished descents corresponding to f (with its dot ignored) through the bijection of Theorem 2.4.18. The descent class of $\mathscr{M}$ is (n_1, n_2), and the dotted symbol k of f distinguishes a vertex of $\mathscr{M}$, as follows. If $1 \leq k \leq n_1$, then mark the vertex of index k in face 1. If $n_1 < k \leq n_1 + n_2$, then mark the vertex of index $k - n_1$ in face 2.

B. For $j = 1, 2$, let $i_j \in \{1, 2, \ldots, n_j\}$ be the index of the vertex at the distinguished descent of face j of $\mathscr{M}$.

C. There is a unique shortest path in $\mathscr{M}$ from its distinguished vertex to a vertex in its core. Remove all vertices and edges of $\mathscr{M}$ except those belonging to either $\mathscr{M}^c$ or this path to obtain a submap $\mathscr{N}$ of $\mathscr{M}$ with a tail. Let (θ_1, θ_2) be the descent class of $\mathscr{N}$.

D. For $j = 1, 2$, let $\mathcal{B}_j$ be the set of $\mathscr{N}$-branches of face j of $\mathscr{M}$. Calculate the index of each branch in $\mathcal{B}_j$.

E. Let $(\lambda, \mathscr{P})$ be the pair corresponding to $\mathscr{N}$ through the bijection $\Omega_{(\theta_1,\theta_2)}$ of Theorem 2.7.11. Let $l_1, \ldots, l_{n_1+n_2-1}$ be the edge labels of $\mathscr{P}$, in order from white end to grey end.

F. Split $\mathscr{P}$ into ordered paths P_1 and P_2 of lengths $\theta_1 - 1$ and $\theta_2 - 1$, respectively, having edge-labels $l_1, \ldots, l_{\theta_1-1}$ and $l_{\theta_1+1}, \ldots, l_{\theta_1+\theta_2-1}$, as encountered from white end to grey end. Set $e_1 = \lambda$ and $e_2 = l_{\theta_1}$.

G. For $j = 1, 2$, form a doubly rooted tree T_j on n_j vertices by attaching all branches in $\mathcal{B}_j$ of index i to the i-th vertex of P_j. The white and grey vertices of P_j serve as the roots of T_j.

H. For $j = 1, 2$, let E_j be the set of edge labels of T_j. Relabel the edges of T_j with $\{1, \ldots, \theta_j\}$ so that the relative order of the original labels is preserved. Regard T_j as a planar one-face map, and assign label n_j to the vertex with index i_j. (If T_j has exactly two vertices, then its white and grey roots are taken to have indices 1 and 2, respectively.)

I. For $j = 1, 2$, transform T_j into a doubly rooted, vertex-labelled tree by pushing the label of each edge e onto the endpoint of e which is furthest from the vertex labelled n_j.

OUTPUT: Pairs (T_1, T_2), (E_1, E_2), and (e_1, e_2), where each T_j is a doubly rooted, vertex-labelled tree on n_j vertices, and $\{E_1, E_2, \{e_1\}, \{e_2\}\}$ is a set partition of $[n_1 + n_2]$ such that $|E_j| = n_j - 1$ for $j = 1, 2$.

Our previous work shows the algorithm above to be reversible. That is, we have the following theorem.

Theorem 2.8.4. *The correspondence defined by Algorithm 2.8.3 is a bijection between genus 0 dotted factorizations of* $(1\,2\,\cdots\,n)(n_1+1\,\cdots\,n_1+n_2)$ *and tuples* $(T_1, T_2, E_1, E_2, e_1, e_2)$*, where each* T_j *is a doubly rooted, vertex-labelled tree on* n_j *vertices, and where* $\{E_1, E_2, \{e_1\}, \{e_2\}\}$ *is a set partition of* $[n_1 + n_2]$ *such that* $|E_j| = n_j - 1$ *for* $j = 1, 2$. □

Since there are n^n doubly rooted, vertex-labelled trees on n vertices, and $(n_1+n_2)\cdot H_0((n_1, n_2))$ genus 0 dotted factorizations of any permutation of class (n_1, n_2), we have a combinatorial derivation of the following special case of the Hurwitz formula.

Corollary 2.8.5. *There are*

$$\frac{1}{n_1+n_2} n_1^{n_1} n_2^{n_2} \binom{n_1+n_2}{n_1-1,\ n_2-1,\ 1,\ 1} = (n_1+n_2-1)!\,\frac{n_1^{n_1}}{(n_1-1)!}\,\frac{n_2^{n_2}}{(n_2-1)!}$$

minimal transitive factorizations of any fixed permutation $\pi \in \mathscr{C}_{(n_1,n_2)}$. □

2.8.3 Factorizations of Class (n_1, n_2, n_3)

We now describe an algorithm that transforms a minimal transitive factorization of class (n_1, n_2, n_3) into three doubly rooted, vertex-labelled trees and a certain set partition of $[n_1 + n_2 + n_3 + 1]$. The structure of this algorithm is very similar to that of Algorithm 2.8.3. See Figure 2.31 on page 96 for an illustration of the procedure as it is applied to the factorization

$$(11\ 12)(15\ 16)(14\ 17)(3\ 14)(14\ 16)(10\ 12)(7\ 13)(2\ 16)(6\ 13)(13\ 16)(5\ 16)(3\ 16)(9\ 13)(1\ 9)(9\ 12)(7\ 12)(3\ 4)(8\ 13)$$

of the permutation $(1\,2\,3\,4\,5\,6\,7)(8\,9\,10\,11\,12\,13)(14\,15\,16\,17) \in \mathfrak{S}_{17}$.

Algorithm 2.8.6.

INPUT: A genus 0 factorization f of $(1\,2\,\cdots\,n_1)(n_1+1\,\cdots\,n_1+n_2)(n_1+n_2+1\,\cdots\,n_1+n_2+n_3)$.

A. Let $\mathscr{M}$ be the decorated face-labelled map map of descent class (n_1, n_2, n_3) corresponding to f through the bijection of Theorem 2.4.18.

B. For $j = 1, 2, 3$, let $i_j \in \{1, \dots, n_j\}$ be the index of the vertex at the distinguished descent of face j of $\mathscr{M}$.

C. Let $\mathscr{N}$ be the core of $\mathscr{M}$, and let $(\theta_1, \theta_2, \theta_3)$ be its descent class.

D. For $j = 1, 2, 3$, let $\mathcal{B}_j$ be the set of branches of face j of $\mathscr{M}$. Calculate the index of each branch in $\mathcal{B}_j$.

E. Let $(\lambda, \mathscr{P}, \gamma)$ be the tuple corresponding to $\mathscr{N}$ through the bijection $\Omega_{(\theta_1,\theta_2,\theta_3)}$ defined in the proof of Theorem 2.7.14. Let $l_1, \dots, l_{n_1+n_2+n_3-1}$ be the edge labels of $\mathscr{P}$, in order from white end to grey end.

F. Split $\mathscr{P}$ into ordered paths P_1, P_2 and P_3 of lengths $\theta_1 - 1$, $\theta_2 - 1$ and $\theta_3 - 1$, respectively, having edge-labels $l_1, \dots, l_{\theta_1-1}$ and $l_{\theta_1+1}, \dots, l_{\theta_1+\theta_2-1}$, and $l_{\theta_1+\theta_2+1}, \dots, l_{\theta_1+\theta_2+\theta_3-1}$, as encountered from white end to grey end. Set $e_1 = \lambda$, $e_2 = l_{\theta_1}$, $e_3 = l_{\theta_1+\theta_2}$, and $e_4 = \gamma$.

G. For $j = 1, 2, 3$, form a doubly rooted tree T_j on n_j vertices by attaching all branches in $\mathcal{B}_j$ of index i to the i-th vertex of P_j. The white and grey vertices of P_j serve as the roots of T_j.

H. For $j = 1, 2, 3$, let E_j be the set of edge labels of T_j. Now relabel the edges of T_j with $\{1, \dots, \theta_j\}$ so that the relative order of the original labels is preserved. Regard T_j as a planar one-face map, and assign label n_j to the vertex with index i_j. (If T_j has exactly two vertices, then its white and grey roots are taken to have indices 1 and 2, respectively.)

I. For $j = 1, 2, 3$, transform T_j into a doubly rooted, vertex-labelled tree by pushing the label of each edge e onto the endpoint of e which is furthest from the vertex labelled n_j.

OUTPUT: Tuples (T_1, T_2, T_3), (E_1, E_2, E_3), and (e_1, e_2, e_3, e_4), where each T_j is a doubly rooted, vertex-labelled tree on n_j vertices, and where $\{E_1, E_2, E_3, \{e_1\}, \{e_2\}, \{e_3\}, \{e_4\}\}$ is a set partition of $[n_1 + n_2 + n_3 + 1]$ such that $|E_j| = n_j - 1$ for $j = 1, 2, 3$.

Again, this algorithm is clearly reversible. We therefore obtain the following theorem and its enumerative corollary, which is another special case of the Hurwitz formula.

Theorem 2.8.7. *The correspondence defined by Algorithm 2.8.6 is a bijection between genus 0 factorizations of the permutation* $(1\,2\,\cdots\,n)(n_1+1\,\cdots\,n_1+n_2)(n_1+n_2+1\,\cdots\,n_1+n_2+n_3)$ *and tuples* $(T_1, T_2, T_3 E_1, E_2, E_3, e_1, e_2, e_3, e_4)$*, where each* T_j *is a doubly rooted, vertex-labelled tree on* n_j *vertices, and where* $\{E_1, E_2, E_3, \{e_1\}, \{e_2\}, \{e_3\}, \{e_4\}\}$ *is a set partition of* $[n_1+n_2+n_3+1]$ *such that* $|E_j| = n_j - 1$ *for* $j = 1, 2, 3$*.* □

Corollary 2.8.8. *There are*

$$n_1^{n_1} n_2^{n_2} n_3^{n_3} \binom{n_1+n_2+n_3+1}{n_1-1,\ n_2-1,\ n_3-1,\ 1,\ 1,\ 1,\ 1} = (n_1+n_2+n_3+1)!\,\frac{n_1^{n_1}}{(n_1-1)!}\,\frac{n_2^{n_2}}{(n_2-1)!}\,\frac{n_3^{n_3}}{(n_3-1)!}$$

minimal transitive factorizations of any fixed permutation $\pi \in \mathscr{C}_{(n_1,n_2,n_3)}$*.* □

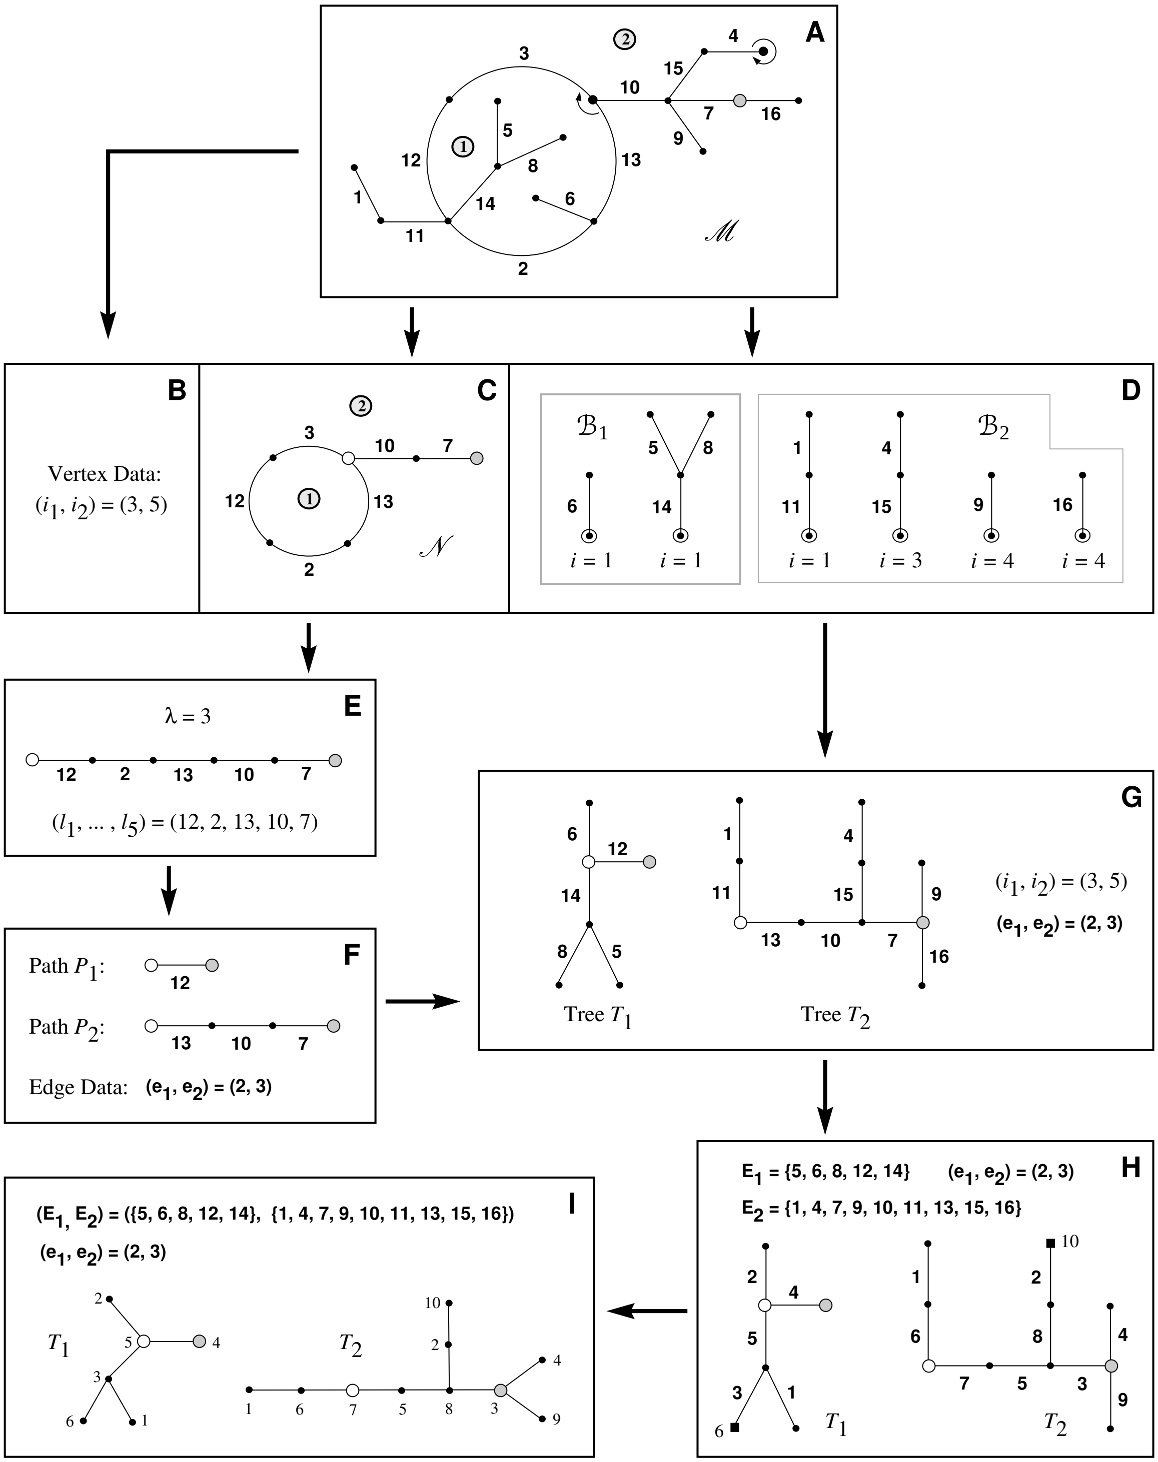

Figure 2.30: An illustration of Algorithm 2.8.3.

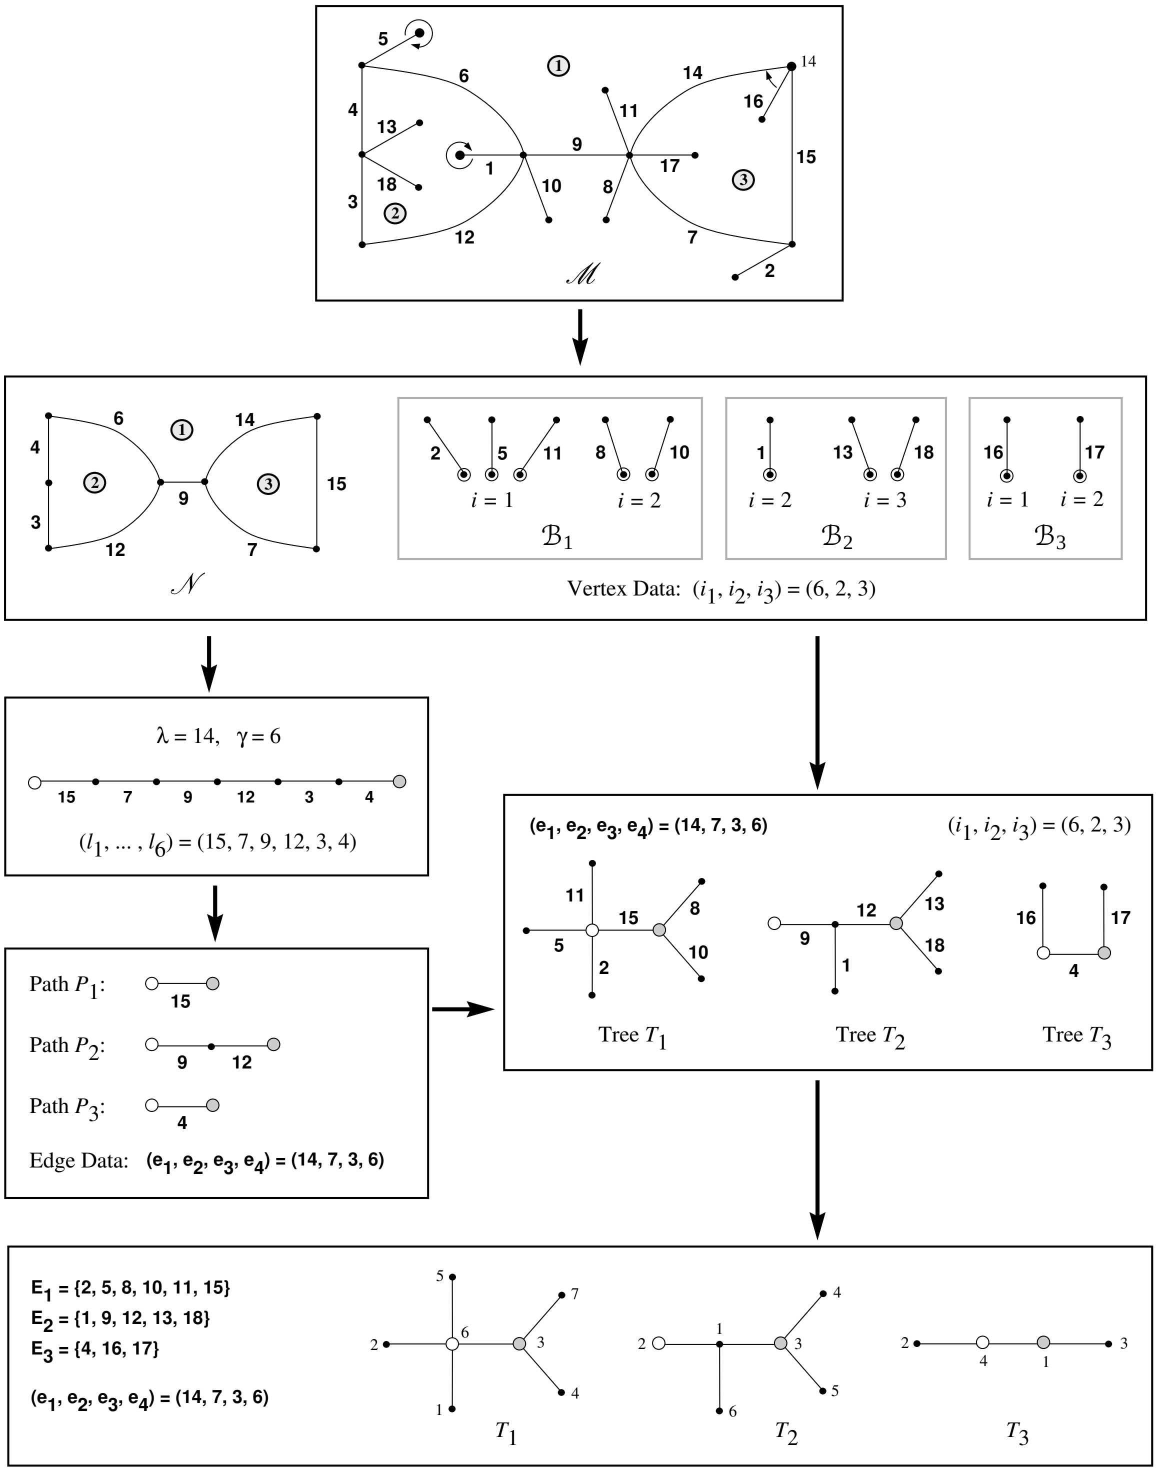

Figure 2.31: An illustration of Algorithm 2.8.6.

Chapter 3

Generalizations

3.1 Introduction

In this chapter we extend some of the results of Chapter 2 to factorizations of permutations in which the factors are not necessarily transpositions. Throughout, we assume the following definitions.

Definition 3.1.1. *A **factorization** of $\pi \in \mathfrak{S}_n$ of **length** r is an r-tuple $(\sigma_r, \ldots, \sigma_1)$ of permutations $\sigma_i \in \mathfrak{S}_n$ such that $\pi = \sigma_r \cdots \sigma_1$. The **class** of a factorization $(\sigma_r, \ldots, \sigma_1)$ of π is the cycle type of π, and its **factor type** is the r-tuple $(\beta_r, \ldots, \beta_1)$, where $\sigma_i \in \mathscr{C}_{\beta_i}$ for $1 \leq i \leq r$.*

For example, $((1\,2)(3\,4)(5), (1\,2\,4)(3)(5), (1\,5)(2\,3)(4))$ is a factorization of $(1\,5)(2\,4)(3)$ of length 3 since

$$(1\,5)(2\,4)(3) = (1\,2)(3\,4)(5) \cdot (1\,2\,4)(3)(5) \cdot (1\,5)(2\,3)(4). \tag{3.1}$$

This factorization has class $[1\,2^2]$ and factor type $([1\,2^2], [1^2\,3], [1\,2^2])$. As in Chapter 2, there is typically no harm in circumventing some formality and referring equation (3.1) itself as a factorization.

The factorization $f = (\sigma_r, \ldots, \sigma_1)$ is **transitive** if the group $\langle \sigma_1, \ldots, \sigma_r \rangle$ generated by its factors acts transitively on $\mathfrak{S}_n$. More generally, if $c = |\operatorname{orb}\langle \sigma_1, \ldots, \sigma_r \rangle|$ is the number of orbits of $[n]$ under the action of $\langle \sigma_1, \ldots, \sigma_r \rangle$, then we say f is a ***c*-component** factorization. Thus 1-component factorizations are synonymous with transitive factorizations.

The following proposition is an extension of Proposition 2.2.9 to arbitrary factorizations.

Proposition 3.1.2. *Let $f = (\sigma_r, \ldots, \sigma_1)$ be a c-component factorization of $\pi \in \mathfrak{S}_n$. Then*

$$\ell(\pi) + \sum_{i=1}^{r} \ell(\sigma_i) \leq n(r-1) + 2c.$$

Proof. Each factor of f can be decomposed into a product $\sigma_i = \tau^i_{r_i} \cdots \tau^i_1$ of $r_i = n - \ell(\sigma_i)$ transpositions. Hence $\pi = \sigma_r \cdots \sigma_1$ can be expressed as a product of the $\sum_{i=1}^{r} r_i = nr - \sum_{i=1}^{r} \ell(\sigma_i)$ transpositions $\{\tau^i_j\}$. Let T be the subgroup of $\mathfrak{S}_n$ generated by these. Proposition 2.2.9 then yields

$$nr - \sum_{i=1}^{r} \ell(\sigma_i) \geq n + \ell(\pi) - 2|\operatorname{orb} T|. \tag{3.2}$$

But $\langle \sigma_1, \ldots, \sigma_r \rangle$ is a subgroup of T, and so we have $|\operatorname{orb} T| \leq |\operatorname{orb}\langle \sigma_1, \ldots, \sigma_r \rangle| = c$. □

Since the quantity $nr - \sum_{i=1}^{r} \ell(\sigma_i)$ on the left-hand side of (3.2) arises as the length of a factorization of π into transpositions, the parity restriction of Proposition 2.2.6 implies $nr - \sum_{i=1}^{r} \ell(\sigma_i) \equiv n - \ell(\pi) \equiv n + \ell(\pi) \pmod 2$. From (3.2), it follows that there is a unique nonnegative integer g such that

$$nr - \sum_{i=1}^{r} \ell(\sigma_i) = n + \ell(\pi) - 2c + 2g.$$

For transitive factorizations (that is, when $c = 1$) we make the following definition.

Definition 3.1.3. *Let f be a transitive factorization of class α and factor type $(\beta_1, \ldots, \beta_r)$. Then*

$$\ell(\alpha) + \sum_{i=1}^{r} \ell(\beta_i) = n(r-1) + 2 - 2g \tag{3.3}$$

for a nonnegative integer g that is called the ***genus*** *of f. Genus 0 factorizations are also referred to as* ***minimal transitive factorizations****.*

3.2 Graphical Representation of General Factorizations

3.2.1 Polymaps

Recall that a map is **2-coloured** if its faces have been painted black and white so every edge is incident with both a black face and a white face. (Thus no two similarly coloured faces are adjacent.). We shall be concerned with a special class of labelled 2-coloured maps, defined as follows.

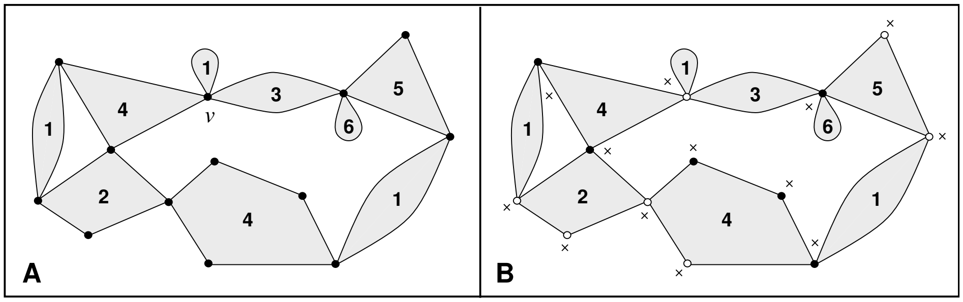

Figure 3.1: (A) A polymap, and (B) its descent structure.

Definition 3.2.1. *A **polygonal map**, or **polymap**, is a 2-coloured map in which the boundary walk of every black face is a cycle.*

We shall find constant need to differentiate between the black faces and white faces of a polymap. The following terminology allows us to do so with minimal effort.

Definition 3.2.2. *The black faces of a polymap are called **polygons**, with an **m-gon** being a black face of degree m. The white faces of a polymap are referred to simply as its **faces**, and a **corner** always refers to a corner of a white face.*

If the polygons of a polymap are labelled, then we define the **rotator** of a vertex v to be the unique cyclic list of black face labels encountered on a clockwise tour of small radius about v. Just as we worked exclusively with edge-labelled maps in Chapter 2, in this chapter we shall be considering only polymaps with labelled polygons. Thus we adopt the following familiar convention throughout:

- The polygons of every polymap are labelled with positive integers in such a way that the rotator of each vertex is increasing.

Notice that, in contrast with our convention for edge-labelled maps, the polygon labels of a polymap need not be distinct.

Example 3.2.3. Figure 3.1A illustrates a polymap with 9 polygons and 3 faces. Note that loops are allowed. The rotator of vertex v is $(1, 3, 4)^{\circ}$. □

We regard the edges of a polymap as being labelled, with an edge inheriting its label from the unique polygon that it borders. This convention allows us to define **descents** of the (white) faces of a polymap exactly as they were defined for edge-labelled maps in §2.4.5. The **descent corners**

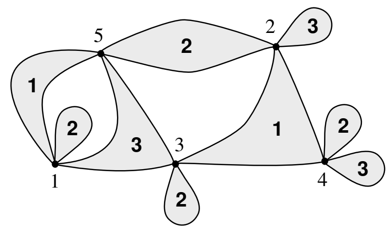

Figure 3.2: A 3-constellation on 5 vertices.

of a face, as well as its **descent set** and **descent cycle**, are defined similarly. Since all rotators are increasing, each vertex of a polymap is at a descent of exactly one face, implying that the descent sets of the faces are disjoint and partition the vertex set. If a polymap has m_i faces containing i descents, then its **descent partition** is $[1^{m_1}2^{m_2}\cdots]$.

Example 3.2.4. Consider the polymap shown in Figure 3.1A. The cyclic list of edge labels encountered along the boundary walk of its outer face is $(1, 4, 1, 3, 5, 5, 1, 4, 4, 2, 2)^{\circ}$. This face therefore has 7 descents. Similarly, the other faces contain 1 and 5 descents. Thus the polymap has descent partition $(7, 5, 1)$. In Figure 3.1B, all descent corners of this polymap are marked with crosses, and the hollow vertices comprise the descent set of the outer face. □

3.2.2 Constellations

Let r be a positive integer. An ***r*-constellation** is a vertex-labelled polymap in which the rotator of every vertex is $(1, 2, \ldots, r)^{\circ}$. Figure 3.2, for example, illustrates a 3-constellation on 5 vertices. Our interest in this special class of polymaps stems from the fact that every transitive factorization of length r corresponds to a unique r-constellation. A formal technical description of this correspondence will be given below, but it is illuminating to begin with a rough outline.

Let $f = (\sigma_r, \ldots, \sigma_1)$ be a transitive factorization. For each cycle $(s_1, \ldots, s_m)$ of σ_i, create a black k-gon labelled i and label its vertices $s_1, \ldots, s_k$ in clockwise order around its perimeter. Doing so for every cycle of each of the factors $\sigma_1, \ldots, \sigma_r$ results in a collection of $\ell(\sigma_1) + \cdots + \ell(\sigma_r)$ labelled black polygons. Now join these polygons by topologically identifying similarly labelled vertices to create a 2-coloured map in which the rotator of every vertex is $(1, \ldots, r)^{\circ}$. This map is the r-constellation associated with f.

Example 3.2.5. The 3-constellation of Figure 3.2 corresponds to the factorization $(\sigma_3, \sigma_2, \sigma_1)$ with $\sigma_1 = (1\,5)(2\,4\,3)$, $\sigma_2 = (1)(2\,5)(3)(4)$, and $\sigma_3 = (1\,5\,3)(2)(4)$. This is a genus 0 factorization of $(1\,2\,4)(3)(5)$. □

Formally describing this correspondence essentially amounts to replacing the somewhat vague reference to "topological identification" with an appeal to Theorem 1.3.4. We now proceed along these lines.

Let $f = (\sigma_r, \ldots, \sigma_1)$ be a transitive factorization of $\pi \in \mathfrak{S}_n$. Let H be the set of all $2nr$ symbols of the form i_j^+ or i_j^- with $1 \leq i \leq r$ and $1 \leq j \leq n$, and consider the pair $(\epsilon, \nu) \in \mathfrak{S}_H \times \mathfrak{S}_H$ defined by

$$\epsilon = \prod_{j=1}^{n} (1_j^- \; 1_{\sigma_1(j)}^+)(2_j^- \; 2_{\sigma_2(j)}^+) \cdots (r_j^- \; r_{\sigma_r(j)}^+) \qquad \text{and} \qquad \nu = \prod_{j=1}^{n} (1_j^- \; 1_j^+ \; 2_j^- \; 2_j^+ \; \cdots \; r_j^- \; r_j^+).$$

For $1 \leq k \leq r$, let π_k denote the partial product $\sigma_k \cdots \sigma_1$. Then, under repeated action of $\epsilon\nu$, the symbol $1_{\sigma_1(j)}^+$ is mapped along the following orbit:

$$1_{\pi_1(j)}^+ \to 2_{\pi_2(j)}^+ \to 3_{\pi_3(j)}^+ \to \cdots \to (r-1)_{\pi_{r-1}(j)}^+ \to r_{\pi_r(j)}^+ \to 1_{\pi_1\pi_r(j)}^+ \to 2_{\pi_2\pi_r(j)}^+ \to \cdots .$$

Since $\pi_r = \pi$, it follows that

$$w_p = (1_{\pi_1(p_1)}^+ 2_{\pi_2(p_1)}^+ \cdots r_{p_2}^+ 1_{\pi_1(p_2)}^+ 2_{\pi_2(p_2)}^+ \; \cdots \; r_{p_{m-1}}^+ 1_{\pi_1(p_m)}^+ 2_{\pi_2(p_m)}^+ \cdots r_{p_1}^+) \tag{3.4}$$

is a cycle of $\epsilon\nu$ whenever $p = (p_1 \; p_2 \; \cdots \; p_m)$ is a cycle of π. Also note that the symbol i_j^- follows the orbit

$$i_j^- \to i_{\sigma_i^{-1}(j)}^- \to i_{\sigma_i^{-2}(j)}^- \to i_{\sigma_i^{-3}(j)}^- \to \cdots$$

under the repeated action of $\epsilon\nu$. Thus

$$b_{s,i} = (i_{s_m}^- \; i_{s_{m-1}}^- \; i_{s_{m-2}}^- \; \cdots \; i_{s_2}^- \; i_{s_1}^-) \tag{3.5}$$

is a cycle of $\epsilon\nu$ whenever $s = (s_1 \, s_2 \; \cdots \; s_m)$ is a cycle of σ_i. In fact, all cycles of $\epsilon\nu$ are of one of the two forms (3.4) or (3.5), as can be seen by observing that every symbol $i_j^{\pm}$ appears in some such cycle. That is, we have the disjoint cycle decomposition $\epsilon\nu = \prod_p w_p \cdot \prod_{i=1}^m \prod_s b_{s,i}$, where the first product extends over all cycles of π, and the last extends over all cycles of σ_i.

Since $\langle \sigma_1, \ldots, \sigma_n \rangle$ acts transitively on $[n]$, the pair (ϵ, ν) defines a transitive rotation system on the set of half-edge symbols H. Let $\mathscr{M}_f$ be the map associated with this system through the correspondence of Theorem 1.3.4. As described in §1.3.8, the half-edges of $\mathscr{M}_f$ are labelled with H, and its vertices, edges, and faces correspond to the cycles of ν, ϵ, and $\epsilon\nu$, respectively.

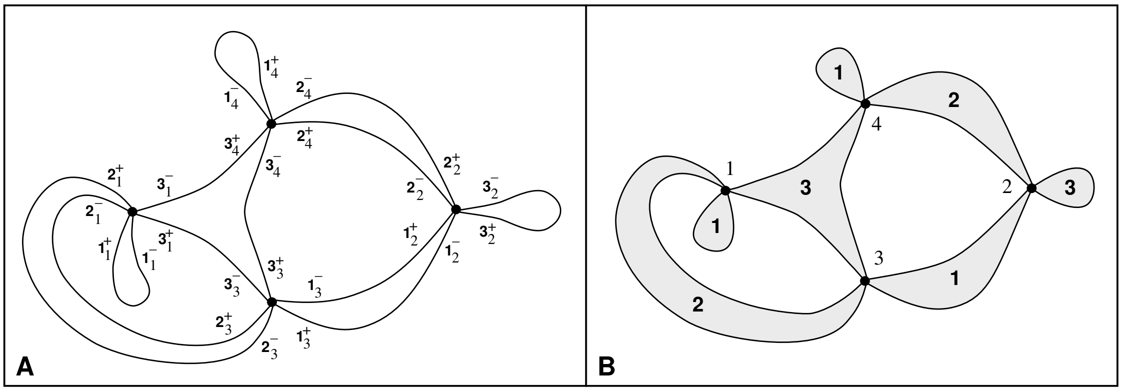

Figure 3.3: The map of the factorization $(1)(2\,4)(3) = (1\,4\,3)(2) \cdot (1\,3)(2\,4) \cdot (1)(2\,3)(4)$.

Example 3.2.6. Consider the factorization $f = (\sigma_3, \sigma_2, \sigma_1)$ of $\pi = (1)(2\,4)(3) \in \mathfrak{S}_4$, where $\sigma_1 = (1)(2\,3)(4)$, $\sigma_2 = (1\,3)(2\,4)$, and $\sigma_3 = (1\,4\,3)(2)$. Here we have

$$
\begin{aligned}
\nu &= (1_1^- \, 1_1^+ \, 2_1^- \, 2_1^+ \, 3_1^- \, 3_1^+)(1_2^- \, 1_2^+ \, 2_2^- \, 2_2^+ \, 3_2^- \, 3_2^+)(1_3^- \, 1_3^+ \, 2_3^- \, 2_3^+ \, 3_3^- \, 3_3^+)(1_4^- \, 1_4^+ \, 2_4^- \, 2_4^+ \, 3_4^- \, 3_4^+) \\
\epsilon &= (1_1^- \, 1_1^+)(2_1^- \, 2_3^+)(3_1^- \, 3_4^+)(1_2^- \, 1_3^+)(2_2^- \, 2_4^+)(3_2^- \, 3_2^+)(1_3^- \, 1_2^+)(2_3^- \, 2_1^+)(3_3^- \, 3_1^+)(1_4^- \, 1_4^+)(2_4^- \, 2_2^+)(3_4^- \, 3_3^+).
\end{aligned}
$$

The half-edge-labelled map $\mathscr{M}_f$ corresponding to the rotation system (ϵ, ν) is shown in Figure 3.3A. Observe that the cycles in the product

$$
\begin{aligned}
\epsilon\nu &= w_{(1)} \cdot w_{(2\,4)} \cdot w_{(3)} \cdot b_{(1),1} \cdot b_{(2\,3),1} \cdot b_{(4),1} \cdot b_{(1\,3),2} \cdot b_{(2\,4),2} \cdot b_{(1\,4\,3),3} \cdot b_{(2),3} \\
&= (1_1^+ \, 2_3^+ \, 3_1^+)(1_3^+ \, 2_1^+ \, 3_4^+ \, 1_4^+ \, 2_2^+ \, 3_2^+)(1_2^+ \, 2_4^+ \, 3_3^+)(1_1^-)(1_3^- \, 1_2^-)(1_4^-)(2_3^- \, 2_1^-)(2_4^- \, 2_2^-)(3_3^- \, 3_4^- \, 3_1^-)(3_2^-)
\end{aligned}
$$

describe the faces of this map. For example, the cycle $w_{(2\,4)} = (1_3^+ \, 2_1^+ \, 3_4^+ \, 1_4^+ \, 2_2^+ \, 3_2^+)$ lists the terminal half-edges encountered along the boundary walk of the outer face. □

Now redecorate $\mathscr{M}_f$ by replacing the half-edge labelling with the following equivalent scheme. For $1 \leq j \leq n$, assign label j to the vertex of $\mathscr{M}_f$ associated with the cycle $(1_j^- \, 1_j^+ \, 2_j^- \, 2_j^+ \, \cdots \, r_j^- \, r_j^+)$ of ν. For $1 \leq i \leq r$, assign label i to each face of $\mathscr{M}_f$ that is associated with a cycle of $\epsilon\nu$ of type $b_{s,i}$. Now paint the faces of $\mathscr{M}_f$ by colouring white all those faces corresponding to cycles of $\epsilon\nu$ of type w_p, and colouring black all those corresponding to cycles of type $b_{s,i}$. Note that the two half-edges i_j^+ and $i_{\sigma_i(j)}^-$ comprising the generic edge $(i_j^- \, i_{\sigma_i(j)}^+)$ appear in $\epsilon\nu$ in cycles of types $b_{s,i}$ and w_p, respectively. Thus no edge occurs in the boundary walk of two distinct, similarly coloured faces. Moreover, the rotator of each vertex of $\mathscr{M}_f$ is $(1, 2, \ldots, r)^\circ$, by construction. Remove the original half-edge labels of $\mathscr{M}_f$, as they are now superfluous. The resulting structure, which we continue to denote by $\mathscr{M}_f$, is an r-constellation naturally corresponding to f.

Example 3.2.7. The 3-constellation $\mathscr{M}_f$ corresponding to the factorization $f = (\sigma_3, \sigma_2, \sigma_1)$ discussed in the previous example is shown in Figure 3.3B. Note that f is easily recovered from $\mathscr{M}_f$, as the disjoint cycles of σ_i are just the cyclic lists of vertex labels encountered along *clockwise* boundary traversals of the polygons of $\mathscr{M}_f$ labelled i. □

In agreement with the terminology of Chapter 2, we refer to the constellation $\mathscr{M}_f$ corresponding to a transitive factorization f as the **polymap of** $\boldsymbol{f}$. The construction of $\mathscr{M}_f$ from f described above is clearly reversible, so the correspondence $f \mapsto \mathscr{M}_f$ is bijective between transitive factorizations and constellations. We denote this bijection by MAP.

3.2.3 A Bijection Between Factorizations and Polymaps

Let $\mathscr{M}$ be an r-constellation on n vertices. For $1 \leq r \leq n$ and $j \geq 1$, let b_{ij} be the number of polygons of $\mathscr{M}$ of degree j with label i. Since each vertex of $\mathscr{M}$ is incident with exactly one polygon labelled i, we have $\sum_j j b_{ij} = n$ for each i. Define partitions $\beta_1, \ldots, \beta_r$ of n by $\beta_i = [1^{b_{i1}}\, 2^{b_{i2}} \cdots] \vdash n$. We call the r-tuple $(\beta_1, \ldots, \beta_r)$ the **polygon type** of $\mathscr{M}$.

The following result is an analogue of Theorem 2.4.11 for generic factorizations.

Theorem 3.2.8. *Let $\alpha, \beta_1, \ldots, \beta_r \vdash n$. The correspondence* MAP $: f \mapsto \mathscr{M}_f$ *restricts to a bijection between genus g factorizations of class α and factor type $(\beta_1, \ldots, \beta_r)$ and r-constellations of genus g with descent partition α and polygon type $(\beta_1, \ldots, \beta_r)$. Moreover, if f is a factorization of $\pi \in \mathfrak{S}_n$, then the descent cycles of $\mathscr{M}_f$ coincide with the cycles of π.*

Proof. Let $f = (\sigma_r, \ldots, \sigma_1)$ be a factorization of $\pi \in \mathscr{C}_\alpha \subset \mathfrak{S}_n$ with factor type $(\beta_1, \ldots, \beta_r)$. From (3.5) it is immediate that $\mathscr{M}_f$ possesses precisely one polygon labelled i for each cycle of σ_i, so the polygon type of $\mathscr{M}_f$ is $(\beta_1, \ldots, \beta_r)$. From (3.4) we see that each cycle of π corresponds to a unique (white) face of $\mathscr{M}_f$. Let F be the face corresponding to the cycle $(p_1, \ldots, p_m)$ of π. Then, in particular, (3.4) indicates that the cyclic list of edge labels encountered along the boundary walk of F is $(1, 2, \ldots, r, 1, 2, \ldots, r, \cdots, 1, 2, \ldots, r)^\circ$, where there are m iterations of the sequence $1, 2, \ldots, r$. Moreover, the descents of F are seen from (3.4) to occur at the vertices of $\mathscr{M}_f$ labelled $p_1, \ldots, p_m$. Thus the descent cycle of F is precisely $(p_1\, p_2 \cdots p_m)$.

Since $\mathscr{M}_f$ has n vertices, each of degree $2r$, it has nr edges. It also has $\ell(\alpha) + \sum_{i=1}^r \ell(\beta_i)$ faces; that is, the number of faces plus the number of polygons of each label. Suppose $\mathscr{M}_f$ is of genus g.

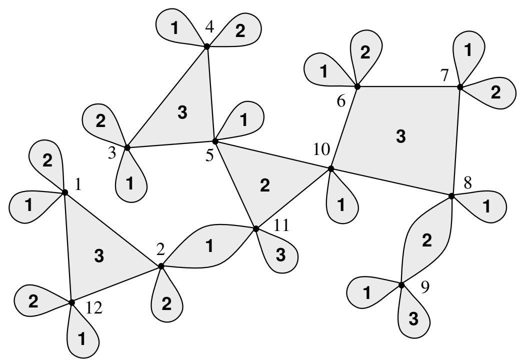

Figure 3.4: The cactus of the factorization (3.6).

Then the Euler-Poincaré formula gives $n - nr + \ell(\alpha) + \sum_{i=1}^{r} \ell(\beta_i) = 2 - 2g$. Equation (3.3) now identifies g with the genus of the factorization f. $\square$

3.2.4 Minimal Transitive Factorizations of Full Cycles

A planar polymap with only one face is called a **cactus**. Thus cacti are natural polymap analogues of trees. What follows is a generalization of the correspondence introduced in §2.4.7 between trees and minimal transitive factorizations of full cycles into transpositions.

For $r \geq 1$, an ***r*-cactus** is a cactus in which every vertex has rotator $(1, 2, \ldots, r)^\circ$ or, equivalently, a planar r-constellation with only one face. By Theorem 3.2.8, minimal transitive factorizations of length r and class (n) are in bijection with vertex-labelled r-cacti on n vertices. For example, Figure 3.4 illustrates the 3-cactus corresponding to the factorization

$$(1\ 2\ 3\ \cdots\ 12) = \overbrace{(1\ 2\ 12)(3\ 4\ 5)(6\ 7\ 8\ 10)}^{\sigma_3} \cdot \overbrace{(5\ 10\ 11)(8\ 9)}^{\sigma_2} \cdot \overbrace{(2\ 11)}^{\sigma_1}. \qquad (3.6)$$

If f is a genus 0 factorization of a fixed full cycle, then observe that the second claim of Theorem 3.2.8 implies all vertex labels of the cactus $\mathscr{M}_f$ are determined by the position of vertex 1. Thus minimal transitive factorizations of $(1\,2\,\cdots\,n)$ with factor type $(\beta_1, \ldots, \beta_r)$ are in one-one correspondence with vertex-rooted r-cacti having polygon type $(\beta_1, \ldots, \beta_r)$. Counting such r-cacti leads to the following result, which originally appears in [27]. See the Additional Notes at the end of this section for further comments.

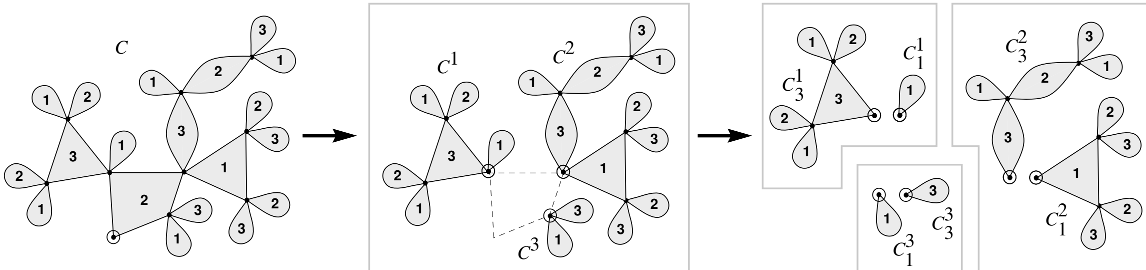

Figure 3.5: Decomposition of a rooted cactus.

Theorem 3.2.9. *Let $\beta_1, \ldots, \beta_r \vdash n$, and set $t_i = \ell(\beta_i)$ for $1 \le i \le r$. If $t_1 + \cdots + t_r = (r-1)n+1$, then there are*

$$n^{r-1} \frac{(t_1-1)!\,(t_2-1)! \cdots (t_r-1)!}{|\operatorname{Aut}(\beta_1)|\,|\operatorname{Aut}(\beta_2)| \cdots |\operatorname{Aut}(\beta_r)|}$$

minimal transitive factorizations of $(1\,2\,\cdots\,n)$ of factor type $(\beta_1, \ldots, \beta_r)$, and there are no such factorizations when this condition is not met.

Sketch proof: Fix $r \ge 1$, and let $\mathcal{C}_i$ be the set of all vertex-rooted cacti in which every vertex except the root has rotator $(1, 2, \ldots, r)^\circ$, while the root vertex itself is incident with a single polygon labelled i. Let $\omega_i = \omega_i(u_i, \mathbf{p}_i)$ be the generating series for $\mathcal{C}_i$, where u_i records the total number of black polygons labelled i, and the component p_{ij} of $\mathbf{p}_i = (p_{i1}, p_{i2}, \ldots)$ marks the number of j-gons labelled i.

Consider any fixed cactus $C \in \mathcal{C}_i$. If its root vertex is incident with a k-gon, then removal of this polygon results in an ordered collection of k rooted cacti, $C^1, \ldots, C^k$, each of whose roots has rotator $(1, \ldots, \widehat{i}, \ldots, r)^\circ$, where the hat indicates that label i is to be suppressed. In turn, each cactus C^j decomposes into $r-1$ rooted cacti, $C^j_1, \ldots, \widehat{C^j_i}, \ldots, C^j_{r-1}$, where $C^j_i \in \mathcal{C}_i$, as is seen by detaching polygons from the root. See Figure 3.5 for an illustration of this decomposition in the case $r = 3$, $k = 4$ and $i = 2$.

For $1 \le i \le r$, define $P_i \in \mathbb{Q}[\mathbf{p}_i][[z]]$ by

$$P_i(z) = \sum_{k \ge 1} p_{ik} z^{k-1}.$$

Then the combinatorial decomposition just described yields

$$\omega_i = u_i P_i(\omega_1 \cdots \widehat{\omega_i} \cdots \omega_r), \qquad \text{for } 1 \le i \le r, \tag{3.7}$$

where again the hat indicates that the factor ω_i is suppressed in the product.

Let $\Omega = \Omega(\mathbf{u}, \mathbf{p})$ be the generating series for the set $\mathcal{C}$ of all vertex-rooted r-cacti, where $\mathbf{u} = (u_1, u_2, \cdots)$, $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \cdots)$, and the various indeterminates mark the same substructures as before. Let $\beta_i = [1^{b_{i1}} 2^{b_{i2}} \cdots]$ and $\mathbf{b}_i = (b_{i1}, b_{i2}, \ldots)$ for $1 \le i \le r$, and set $\mathbf{b} = (\mathbf{b}_1, \mathbf{b}_2, \ldots)$. Then the number of vertex-rooted cacti with polygon type $(\beta_1, \ldots, \beta_r)$ is $[\mathbf{u^t p^b}]\,\Omega$. Now observe that unhinging a rooted r-cactus at its root vertex results in an r-tuple of rooted cacti, one from each of the sets $\mathcal{C}_1, \ldots, \mathcal{C}_r$. Thus $\Omega = \omega_1 \cdots \omega_r$.

The system (3.7) of functional equations implicitly defines $\omega_i \in \mathbb{Q}[\mathbf{p}][[\mathbf{u}]]$ for $1 \le i \le r$. The coefficient $[\mathbf{u^t p^b}]\,\Omega = [\mathbf{u^t p^b}]\,\omega_1 \cdots \omega_r$ can now be evaluated through multivariate Lagrange inversion applied to this system; see [27] for details. Note that the condition $t_1 + \cdots + t_r = (r-1)n+1$ necessary for the desired coefficient to be nonzero is immediate from (3.3) upon setting $g = 0$, $\ell(\beta_i) = t_i$, and $\ell(\alpha) = 1$. □

3.2.5 Suppression of Loops

Let e be a loop in the polymap $\mathscr{M}$. Then e appears in the boundary walk of a unique polygon. In fact, since the boundary walk of this polygon must be a cycle, we see that e bounds a 1-gon. Contracting e to the single vertex with which it is incident has the effect of eliminating this 1-gon from $\mathscr{M}$. Of course, the contraction of loops can be iterated.

Definition 3.2.10. *The loopless polymap obtained from the polymap $\mathscr{M}$ by contracting each of its loops is called the* ***reduction*** *of $\mathscr{M}$, and is denoted by $\mathscr{M}^\dagger$.*

Of importance here is the observation that an r-constellation $\mathscr{M}$ can be recovered from its reduction $\mathscr{M}^\dagger$ provided that r is known. This follows because the location and label of the missing 1-gons are uniquely specified by the fact that the rotator of every vertex of $\mathscr{M}$ is $(1, 2, \ldots, r)^\circ$. For example, Figure 3.6 illustrates the reduction of the map of the factorization

$$f = (\sigma_3, \sigma_2, \sigma_1), \qquad \sigma_3 = (1\,5\,3)(2)(4), \sigma_2 = (1)(2\,5)(3)(4), \sigma_1 = (1\,5)(2\,4\,3). \tag{3.8}$$

To recover $\mathscr{M}_f$ from $\mathscr{M}_f^\dagger$, loops are simply added so as to make each rotator equal $(1, 2, 3)^\circ$.

For the factorization (3.8), notice that the descent structure of the 3-constellation $\mathscr{M}_f$ is effectively unaltered by suppressing its loops. That is, the descent cycle of each face of $\mathscr{M}_f$ is equal to that of the corresponding face of $\mathscr{M}_f^\dagger$. That this is usually the case is a consequence of the increasing rotator condition, as we now demonstrate.

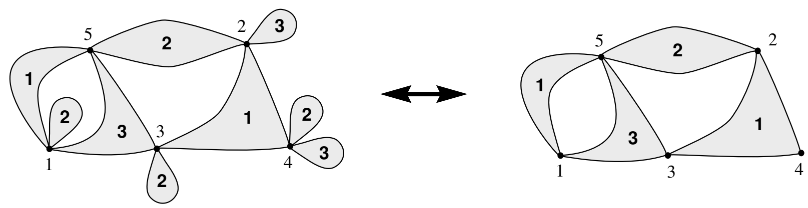

Figure 3.6: The reduction of a 3-constellation.

Lemma 3.2.11. *Let $\mathscr{M}$ be a polymap on at least two vertices, and let e be a loop incident with the face F of $\mathscr{M}$. Let $\mathscr{M}'$ be the polymap obtained from $\mathscr{M}$ by contracting e, and let F' be the face of $\mathscr{M}'$ corresponding to F. Then the descent cycles of F and F' are identical.*

Proof. Suppose e is incident with vertex v. If F were of degree 1, then e would bound both F and a 1-gon, so v would be the only vertex of $\mathscr{M}$. If F were of degree 2, then its boundary walk would be $((v, e), (v, e'))^\circ$ for some loop $e' \neq e$. But, since e' also bounds a 1-gon, no edges aside from e and e' could be incident with v, and again v would be the only vertex of $\mathscr{M}$. Therefore F is of degree at least 2.

The boundary walk of F is therefore $((v, e), (v, e_0), (v_1, e_1), \ldots, (v_k, e_k))^\circ$ for some v_i, e_i, so that the boundary walk of F' is $((v, e_0), (v_1, e_1), \ldots, (v_k, e_k))^\circ$. Thus v is at a descent of F if and only if $e_k \geq e$ or $e \geq e_0$, while v is at a descent of F' if and only if $e_k \geq e_0$.

If $e_0 = e_k$, then one of $e_k \geq e$ or $e \geq e_0$ holds, and obviously $e_k \geq e_0$. Thus v is at a descent of both F and F' in this case.

If $e_0 \neq e_k$, then v is incident with at least three polygons, namely the 1-gon bounded by e and at least two other polygons with labels e_0 and e_k. In fact, $(e_k, e, e_0)^\circ$ is a subsequence of the rotator of v in $\mathscr{M}$, and is therefore increasing. Thus either $e_k < e < e_0$, or $e < e_0 < e_k$, or $e_0 < e_k < e$. It follows that $e_k \geq e_0$ if and only if either $e_k \geq e$ or $e \geq e_0$. That is, v is at a descent of F' if and only if it is at a descent of F. Clearly the cyclic orders in which the descents of F and F' occur along their respective boundary walks are the same, and the result follows. □

Proposition 3.2.12. *Let f be a transitive factorization of a permutation on at least two symbols. Then the descent cycles of corresponding faces of $\mathscr{M}_f$ and $\mathscr{M}_f^\dagger$ are identical.*

Proof. This follows immediately by repeated application of the lemma. □

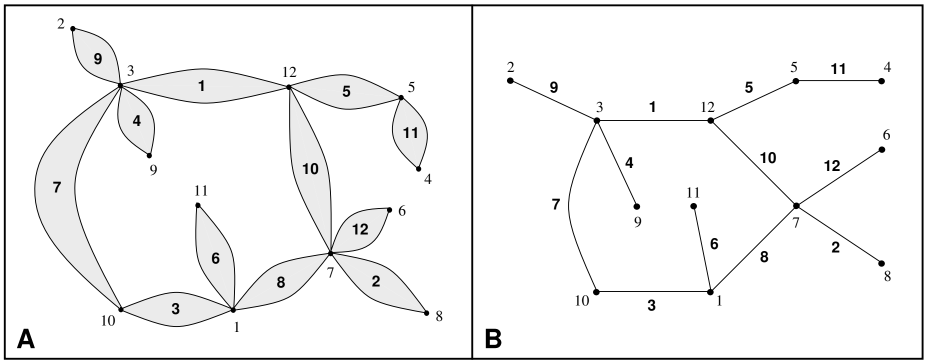

Figure 3.7: Maps corresponding to the factorization (3.9).

3.2.6 Factorizations into Transpositions

Consider the transitive factorization $f = (\tau_r, \ldots, \tau_1)$ of (1 2 3 4 5 6 7 8)(9 10 11 12) into $r = 12$ transpositions given below:

$$\begin{aligned}&(1\ 2\ 3\ 4\ 5\ 6\ 7\ 8)(9\ 10\ 11\ 12)\\ &\qquad = (6\ 7)(4\ 5)(7\ 12)(2\ 3)(1\ 7)(3\ 10)(1\ 11)(5\ 12)(3\ 9)(1\ 10)(7\ 8)(3\ 12).\end{aligned} \tag{3.9}$$

The map $\mathscr{M}_f$ of this factorization is a 12-constellation that contains precisely one 2-gon and ten 1-gons labelled i, for $1 \le i \le 12$. The only polygons in the reduction $\mathscr{M}_f^\dagger$ are therefore 2-gons, which are labelled distinctly with the integers $1, \ldots, 12$. See Figure 3.7A for an illustration of $\mathscr{M}_f^\dagger$. Since the value of r is preserved as the maximal label of these 2-gons, no information has been lost in the reduction $\mathscr{M}_f \mapsto \mathscr{M}_f^\dagger$. Furthermore, "flattening" each 2-gon of $\mathscr{M}_f^\dagger$ into a single edge is clearly a reversible process that results in the vertex- and edge-labelled map drawn in Figure 3.7B. This is, of course, the map we previously called the "map of f", and studied extensively in Chapter 2,

Clearly these same considerations apply more generally to associate with every transitive factorization f into transpositions the vertex- and edge-labelled map previously called the map of f. In this way, Theorem 2.4.11 is seen to be a special case of Theorem 3.2.8. We point out that Proposition 3.2.12 is instrumental in this connection, for it establishes that the descent cycles of $\mathscr{M}_f$ are the same as those of $\mathscr{M}_f^\dagger$, which are, in turn, plainly identical to those of the final "flattened" map. The need for many of the contrivances introduced in the earlier discussion of Theorem 2.4.11 (such as carriers and orbits) is eliminated when the result is established in this more general manner.

3.2.7 Additional Notes

We have borrowed the term *constellation* from [8], where it is used in reference to maps that are dual to the constellations defined here. Schaeffer and Bousquet-Mélou do not consider descent structure in [8]. However, their main result, stated in our language, is a very elegant bijective proof of the following formula for the number of planar r-constellations with descent partition $\alpha = [1^{m_1}\, 2^{m_2}\, \cdots]$:

$$r\frac{((r-1)n-1)!}{((r-1)n-\ell(\alpha)+2)!}\prod_{i=1}^{r}\left[i\binom{ri-1}{i}\right]^{m_i}. \tag{3.10}$$

By Theorem 3.2.8, this formula gives the total number of minimal transitive factorizations $(\sigma_r, \ldots, \sigma_1)$ of class $\alpha \vdash n$. If none of the factors σ_i of such a factorization is the identity, then $\ell(\sigma_i) \le n-1$ for all i, and thus $\sum_{i=1}^{r} \ell(\sigma_i) \le r(m-1)$. With (3.3), this gives $r \le n+\ell(\alpha)-2$. Setting $r = n+\ell(\alpha)-2$ forces these inequalities to be tight, so that $\ell(\sigma_i) = n-1$ for all i. That is to say, a genus 0 factorization of class α and of length $r = n+\ell(\alpha)-2$ in which no factor is the identity is necessarily a minimal transitive factorization into transpositions. This fact is used in [8] to derive the Hurwitz formula from (3.10), by applying inclusion-exclusion to eliminate the contribution of factorizations containing trivial factors.

Since all factorizations of a full cycle are necessarily transitive, Theorem 3.2.9 actually provides an evaluation of the connection coefficient $c^{(n)}_{\beta_1,\ldots,\beta_r}$ of $Z(\mathbb{C}\mathfrak{S}_n)$ in the special case that $\sum_{i=1}^{r} \ell(\beta_i) = (r-1)n+1$. It is in this context that the result first appears, in Goulden and Jackson's extension [27] of previous work of Goupil and Bédard [40]. (See also Farahat and Higman [21].) The cacti considered in [27] are dual to those introduced in this section.

Theorem 3.2.9 is thoroughly generalized in [58], which contains an evaluation of $c^{(n)}_{\beta_1,\ldots,\beta_r}$ for *arbitrary* r and partitions $\beta_1, \ldots, \beta_r \vdash n$. The special case $c^{(2k)}_{\beta,[2^k]}$ is of particular interest because of the following link with geometry. Factorizations of the form $\pi = \sigma\rho$, where σ is a full cycle and ρ is a fixed-point free involution, are seen, by Theorem 1.3.4, to parameterize *monopoles* — that is, maps with a single vertex. They appeared in this guise in the work of Harer and Zagier [43] on the Euler characteristic of the moduli space of curves. These authors obtain explicit enumerative formulae through integration over random matrices, but the same results have since been derived through the character theory of the symmetric group [46, 58] and, recently, by direct bijection [37]. Many other attempts have been made at evaluating particular connection coefficients of $Z(\mathbb{C}\mathfrak{S}_n)$. See, for instance, [5], [7] and [73] for some early efforts.

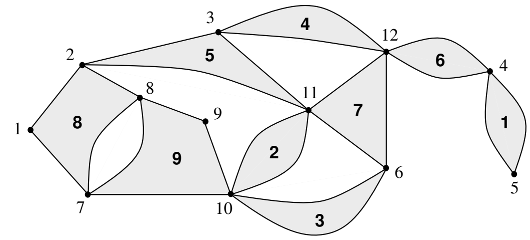

Figure 3.8: The reduced map of the cycle factorization (3.11).

3.3 Cycle Factorizations

We shall now restrict our focus somewhat and explore factorizations whose factors are all cycles, possibly of different lengths. The factorizations (into transpositions) studied in Chapter 2 are of this variety, so that the results of this section naturally generalize our earlier efforts. In particular, we shall find that the method of pruning trees developed in §2.6 remains effective in this more general setting.

3.3.1 Preliminaries

A **cycle factorization** is a factorization whose factors are all cycles of length at least 2. The **cycle index** of such a factorization is the vector $(c_2, c_3, \ldots)$, where c_k is the number of k-cycle factors it contains. For example,

$$\begin{aligned} &(1\ 2\ 3\ 4\ 5\ 6\ 7)(8)(9\ 10)(11)(12) \\ &= (7\ 8\ 9\ 10)\cdot(1\ 2\ 8\ 7)\cdot(6\ 11\ 12)\cdot(4\ 12)\cdot(2\ 3\ 11)\cdot(3\ 12)\cdot(6\ 10)\cdot(10\ 11)\cdot(4\ 5) \end{aligned} \tag{3.11}$$

is a transitive cycle factorization of length 9 with cycle index $(5, 2, 2, 0, 0, \ldots)$.

If f is a transitive cycle factorization of length r with cycle index $(c_2, c_3, \ldots)$, then its reduced map $\mathscr{M}_f^\dagger$ is composed of r polygons distinctly labelled $1, \ldots, r$, with c_k of these being k-gons, for $k \geq 2$. For example, Figure 3.8 shows the reduced map of the factorization (3.11).

As we shall be working exclusively with cycle factorizations in this section, we adopt the following conventions throughout:

- All polymaps are loopless.
- The polygons of every polymap are distinctly labelled.

It will also be convenient to define the **polygon index** of a polymap containing i_k k-gons, for $k \geq 2$, to be the vector $(i_2, i_3, \ldots)$. Of course, cycle factorizations with cycle index $(c_2, c_3, \ldots)$ correspond to polymaps with polygon index $(c_2, c_3, \ldots)$ under the bijection $f \mapsto \mathscr{M}_f^\dagger$.

3.3.2 Properly Labelled Polymaps

Let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be an m-part composition. In accordance with §2.4.10, we say a polymap is of **descent class** α if it contains exactly m faces, these being labelled $1, \ldots, m$ so that the face with label s has exactly α_s descents, for $1 \leq s \leq m$. The **canonical descent sets** $\mathbb{D}_1(\alpha), \ldots, \mathbb{D}_m(\alpha)$ associated with α are defined as before, as is the subset $\mathfrak{S}(\alpha)$ of permutations whose cycles are supported by these sets.

Definition 3.3.1. *A vertex- and face-labelled polymap is said to be* ***properly labelled*** *if it is of descent class α and the face labelled s has descent set $\mathbb{D}_s(\alpha)$, for $1 \leq s \leq m$.*

Theorem 3.3.2, below, is a generalization of Theorem 2.4.18 for cycle factorizations. It is a straightforward consequence of Theorem 3.2.8 and the fact that vertex-labelled polymaps have no nontrivial automorphisms.

Theorem 3.3.2. *Let α be a composition. The set of all genus g cycle factorizations $(\sigma_r, \ldots, \sigma_1)$ satisfying $\sigma_r \cdots \sigma_1 \in \mathfrak{S}(\alpha)$ is in bijection with the set of properly labelled, genus g polymaps that are of descent class α and contain r polygons. Moreover, under this bijection, a factorization with cycle index $(c_2, c_3, \ldots)$ corresponds to a polymap with polygon index $(c_2, c_3, \ldots)$.*

Proof. Let $f = (\sigma_r, \ldots, \sigma_1)$ be a genus g cycle factorization of $\pi \in \mathfrak{S}(\alpha)$. Suppose α has m parts. Then Theorem 3.2.8 and Proposition 3.2.12 together show that $\mathscr{M}_f^\dagger$ is a loopless, vertex-labelled, genus g polymap with m faces, whose descent cycles are supported by $\mathbb{D}_1(\alpha), \ldots, \mathbb{D}_m(\alpha)$, and whose polygons are labelled distinctly with $1, \ldots, r$. Moreover, the cycle index of f coincides with the polygon index of $\mathscr{M}_f^\dagger$. Assigning label s to the face of $\mathscr{M}_f^\dagger$ with descent set $\mathbb{D}_s(\alpha)$, for $1 \leq s \leq m$, therefore produces a loopless, properly labelled, genus g polymap of descent class α with polygon index equal to the cycle index of f. Clearly any such polymap can be constructed in this way and, since $\mathscr{M}_f^\dagger$ admits only the trivial automorphism, two different factorizations cannot lead to the same polymap. □

Example 3.3.3. The properly labelled polymap corresponding to the cycle factorization (3.11) is drawn in Figure 3.9. □

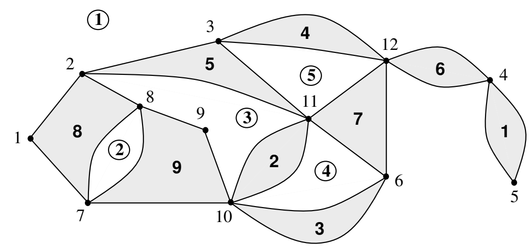

Figure 3.9: The properly labelled planar polymap corresponding to (3.11).

Like its earlier analogue, Theorem 3.3.2 puts us in position to study transitive cycle factorizations solely through the combinatorics of properly labelled polymaps. We begin with the following definitions, which are familiar from §2.4.10.

Definition 3.3.4. *For a vector* $\mathbf{i} = (i_2, i_3, \ldots)$ *of nonnegative integers and a composition* α*, let* $M_g(\alpha\,;\,\mathbf{i})$ *denote the number of properly labelled genus* g *polymaps of descent class* α *that have polygon index* $\mathbf{i}$*. For fixed* $m \geq 1$*, let*

$$\Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u) = \sum_{n\geq 1} \sum_{\mathbf{i}\geq \mathbf{0}} \sum_{\substack{\alpha\models n \\ \ell(\alpha)=m}} M_g(\alpha\,;\,\mathbf{i})\, \mathbf{p}^{\mathbf{i}}\, \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \frac{u^{r(\mathbf{i})}}{r(\mathbf{i})!},$$

be the generating series for the numbers $\{M_g(\alpha\,;\,\mathbf{i}) \,:\, \ell(\alpha) = m,\ \mathbf{i} \geq \mathbf{0}\}$*, where* $\mathbf{x} = (x_1, \ldots, x_m)$ *and* $\mathbf{p} = (p_2, p_3, \ldots)$ *are vectors of indeterminates, and* $r(\mathbf{i}) = i_2 + i_3 + \cdots$*. When considering the genus 0 series, we shall typically write* Ψ_m *in place of* $\Psi_m^{(0)}$*.*

For $1 \leq i \leq m$, the indeterminate x_i in $\Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u)$ is an exponential marker for vertices at descents of face i of a polymap. These vertices are labelled with the i-th canonical descent set. Clearly u is an exponential marker for labelled polygons, and p_k records the number of k-gons, for $k \geq 2$. Observe that the series $\Psi_m^{(g)}(\mathbf{x}, u)$ introduced in §2.4.10 is recovered by setting $p_2 = 1$ and $p_3 = p_4 = \cdots = 0$ in $\Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u)$. Throughout the remainder of this section, the symbol $\mathbf{p}$ will denote the vector $(p_2, p_3, \ldots)$ of indeterminates.

Corollary 3.3.5. *Let* $\alpha \vdash n$ *and fix* $\pi \in \mathscr{C}_\alpha$*. The number of genus* g *cycle factorizations of* π *with cycle index* $\mathbf{c} = (c_2, c_3, \ldots)$ *is given by*

$$r!\, \alpha_1 \cdots \alpha_m \cdot [\mathbf{x}^{\boldsymbol{\alpha}} \mathbf{p}^{\mathbf{c}} u^r]\, \Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u),$$

where $r = c_2 + c_3 + \cdots$*.*

Proof. This is immediate from Theorem 3.3.2 and the fact that $|\mathfrak{S}(\alpha)| = \prod_i (\alpha_i - 1)!$. □

The next proposition is a polymap analogue of Proposition 2.4.24, and plays a similar rôle in our analysis. It implies that the vertices of a planar, face-labelled polymap of descent class $\alpha = (\alpha_1, \ldots, \alpha_m)$ can usually be labelled in $\alpha_1! \cdots \alpha_m!$ ways to obtain distinct properly labelled polymaps. The only exceptions to this rule are polymaps with only one polygon, since the vertices of a k-gon can clearly be labelled in $(k-1)!$ distinct ways. In particular, for $m \geq 2$, we can regard $\Gamma_m(\mathbf{z}, \mathbf{p}, u)$ as the counting series for smooth, planar, face-labelled polymaps with m faces, with respect to descent class and polygon type.

Proposition 3.3.6. *A face-labelled planar polymap with at least two polygons has no nontrivial automorphisms.*

Proof. Suppose ϕ is a nontrivial automorphism of the face-labelled planar polymap $\mathscr{M}$. Clearly a vertex and its image under ϕ are incident with precisely the same polygons. Therefore, since ϕ is nontrivial, it cyclically permutes the vertices of all polygons. Thus either $\mathscr{M}$ consists of a single polygon, or there exist distinct vertices u and v, each incident with at least two polygons, such that $\phi(u) = v$. In the latter case, the cyclic lists of alternating polygon and face labels encountered on clockwise tours about u and v must be the same. In particular, u and v are both incident with distinct polygons labelled a and b, and corners (a, u, b) and (a, v, b) belong to the same face. The situation is illustrated below.

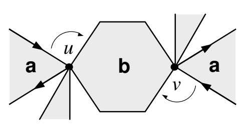

Compatibly directed half-edges in the diagram must be connected to complete polygon a so that the marked corners remain in the same face. Clearly it is not possible to do this in the plane. □

Note that the Proposition 3.3.6 is restricted to *planar* polymaps. Counterexamples in positive genus are given in Figure 3.10A, where both face-labelled polymaps shown are invariant under rotation of all their polygons by 180°. Since non-planar, face-labelled polymaps may have nontrivial automorphisms, vertex-labellings cannot generally be ignored when working with $\Psi_m^{(g)}$ for $g \geq 1$. This problem could be overcome by considering rooted maps, but we shall not need to do so since our attention will generally be restricted to planar polymaps.

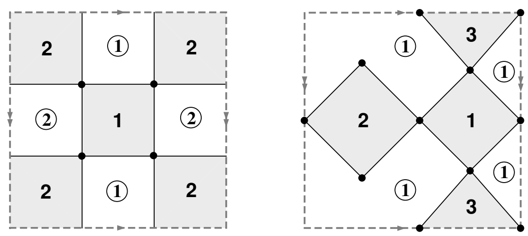

Figure 3.10: Failure of Proposition 3.3.6 in genus 1.

3.3.3 Minimal Transitive Cycle Factorizations of Full Cycles

Recall that a *cactus* is a planar polymap with a single face. By definition, the series $\Psi_1(x, \mathbf{p}, u)$ counts properly labelled cacti, where x is an exponential marker for labelled vertices, p_k is an ordinary marker for k-gons, and u is an exponential marker for polygons. We wish to evaluate this series, and thereby count minimal transitive cycle factorizations of $(1\,2\,\cdots\,n)$.

To this end, first let $w = w(x, \mathbf{p}, u)$ be the generating series for vertex-rooted, properly labelled cacti, with respect to the same markers as above. Then we have

$$w = x \frac{d}{dx} \Psi_1(x, \mathbf{p}, u). \tag{3.12}$$

We now give a decomposition for such cacti that preserves labelled vertices and polygons.

Suppose the root vertex v of a rooted cactus is incident with m polygons. Detaching these polygons from the root results in the single vertex v together with a collection $\{C_1, \ldots, C_m\}$ of rooted cacti with labelled non-root vertices. The root of each C_i is unlabelled and incident with only one polygon. See Figure 3.11. (Labels have been suppressed in the diagram for clarity.) Note that the ordering of $C_1, \ldots, C_m$ around v need not be recorded, as it is can be deduced by virtue of the increasing rotator condition. Thus $w = x \sum_{m \geq 0} (\bar{w})^m / m! = x e^{\bar{w}}$, where $\bar{w} = \bar{w}(x, \mathbf{p}, u)$ is the generating series for cacti such as C_i, and x marks only labelled vertices throughout.

If the root of C_i is incident with a k-gon, then removal of this polygon leaves a $(k-1)$-tuple $C_i^1, \ldots, C_i^{k-1}$ of rooted cacti, as shown in Figure 3.11. This accounts for a contribution $u p_k w^{k-1}$ to the series $\bar{w}$. Summing over k therefore gives

$$\bar{w} = u P(w), \tag{3.13}$$

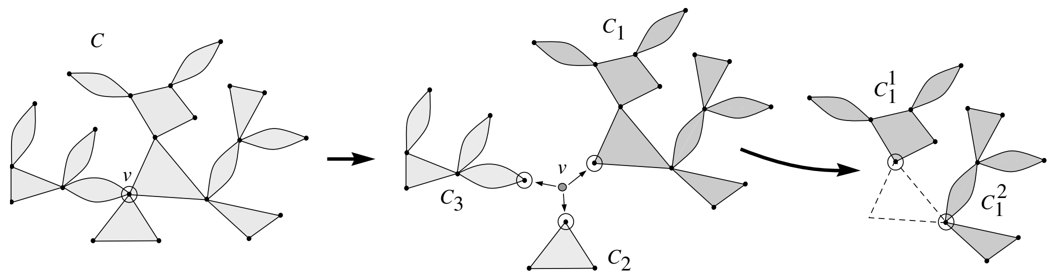

Figure 3.11: Decomposition of a rooted cactus.

where $P \in \mathbb{Q}[\mathbf{p}][[z]]$ is defined by

$$P(z) = \sum_{k\geq 2} p_k z^{k-1}. \tag{3.14}$$

Hence we obtain the recursive definition

$$w = xe^{uP(w)}. \tag{3.15}$$

This functional equation can be solved in $\mathbb{Q}[\mathbf{p}, u][[x]]$ by Lagrange inversion, as is demonstrated in the proof of the next theorem. Since the analysis above is actually a simplification of that used to prove Theorem 3.2.9, this result also follows immediately as a special case of that earlier theorem.

Theorem 3.3.7. *Let* $(i_2, i_3, \ldots)$ *be a sequence of nonnegative integers and set* $r = i_2 + i_3 + \cdots$. *Then there are*

$$\frac{n^{r-1}\, r!}{\prod_{k\geq 2} i_k!}$$

minimal transitive cycle factorizations of $(1\,2\,\cdots\,n)$ *with cycle index* $(i_2, i_3, \ldots)$ *in the case that* $n + r - 1 = \sum_{k\geq 2} k i_k$, *and zero otherwise.*

Proof. From (3.15), Lagrange inversion gives

$$\begin{aligned}
[x^n \mathbf{p}^{\mathbf{i}} u^r]\, w &= \frac{1}{n}\,[\lambda^{n-1}\mathbf{p}^{\mathbf{i}} u^r]\, e^{nuP(\lambda)} \\
&= \frac{n^{r-1}}{r!}\,[\lambda^{n-1}\mathbf{p}^{\mathbf{i}}]\left(\sum_{k\geq 2} p_k \lambda^{k-1}\right)^r \\
&= \frac{n^{r-1}}{r!}\binom{r}{i_2, i_3, \ldots}\,[\lambda^{n-1}]\,\lambda^{\sum_k i_k(k-1)}.
\end{aligned}$$

The result follows by Corollary 3.3.5, since (3.12) implies $[x^n \mathbf{p}^{\mathbf{i}} u^r]\, w = n \cdot [x^n \mathbf{p}^{\mathbf{i}} u^r]\, \Psi_m(x, \mathbf{p}, u)$. □

An interesting special case of this theorem concerns factorizations of a full cycle into cycles of the same length. In general, for $k \geq 2$, we define a ***k*-cycle factorization** to be a factorization whose factors are all k-cycles.

Lemma 3.3.8. *Let $k \geq 2$. A k-cycle factorization of $\pi \in \mathfrak{S}_n$ of genus g has exactly*

$$\frac{n + \ell(\pi) + 2g - 2}{k - 1}$$

factors. If this number is not integral then no such factorization of π exists.

Proof. Let $(\sigma_r, \ldots, \sigma_1)$ be a k-cycle factorization of π of genus g. Then $\ell(\sigma_i) = n - k + 1$ for all $1 \leq i \leq r$, so that $\ell(\pi) + \sum_{i=1}^{r}(n - k + 1) = n(r - 1) + 2 - 2g$. Solving for r produces the result. □

Notice that the lemma identifies $\frac{1}{k-1}(n + \ell(\alpha) - 2)$ as the minimal number of factors in a transitive k-cycle factorization of class α, with this minimum attained for minimal transitive (*i.e.* genus 0) factorizations. We now have the following corollary of Theorem 3.3.7.

Corollary 3.3.9. *Fix $n \geq 1$ and $k \geq 2$. If $n = 1 + r(k - 1)$ for a positive integer r, then there are n^{r-1} minimal transitive k-cycle factorizations of the full cycle $(1\,2\,\cdots\,n)$.* □

3.3.4 Differential Equations for Planar Polymaps

Having introduced properly labelled polymaps, we should now look for a polymap analogue of Theorem 2.5.1. Such a result would, at least, provide us with a recursive computational scheme for evaluating the series Ψ_m for all $m \geq 1$.

To prove Theorem 2.5.1, we considered the effect of deleting the edge of maximal label from a face-labelled map. Since an edge is incident with at most two faces, its deletion either separates a map into two maps, or merges two faces into one. Our proof of Theorem 2.5.1 came from an analysis of these distinct cases. Analogously, we should now study the effect of deleting polygons from face-labelled planar polymaps. (Vertex labels can be ignored by virtue of Proposition 3.3.6.)

Let $\mathscr{M}$ be a planar polymap with two faces. Clearly no polygon of $\mathscr{M}$ can border more than two faces, so the removal of any polygon leaves either two cacti, or a cactus and a smaller two-face polymap. This decomposition leads to a recursive differential equation involving (3.12) as initial data. Solving this equation yields an expression for $\Psi_2(\mathbf{x}, \mathbf{p}, u)$ that generalizes Corollary 2.5.4. The derivation is similar to that of the earlier corollary, but is not included here because we shall obtain the result through different methods later. (See Corollary 3.3.15.)

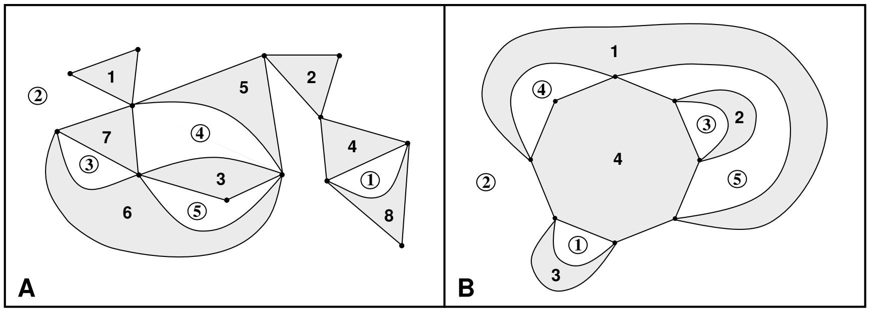

Figure 3.12: Complicated polymaps.

For polymaps with three or more faces, a cursory analysis reveals complications that did not arise in Chapter 2. Consider, for example, the polymap $\mathscr{M}$ shown in Figure 3.12A. Let Δ_i be the triangle (3-gon) of $\mathscr{M}$ with label i, for $1 \le i \le 8$. Observe that some triangles, such as Δ_2, are incident with only one face, some are incident with two faces (*e.g.*, Δ_5), and some are incident with three faces (*e.g.*, Δ_6). Thus deletion of a triangle can result in one, two, or three polymaps, and faces of $\mathscr{M}$ can be merged in complicated ways in the process. This makes it quite difficult to keep track of descents. For example, removal of Δ_5 results in two polymaps; the outer face of one of these inherits descents from faces 2 and 4 of $\mathscr{M}$, while the outer face of the other inherits descents only from face 2.

In general, removal of a k-gon incident with j faces results in a collection of $k-j+1$ polymaps (some of these may consist only of a single vertex), and the possible interactions between a k-gon and its ambient polymap grow more complex for larger values of k. For instance, consider the removal of the octagon from the polymap of Figure 3.12B. This leaves four polymaps (one of these consists of a single vertex), each with only one face. Clearly some description of the incidences between these polymaps and the original must be recorded if the deletion process is to be reversible.

Significant progress has been made on this problem in [31], though the paper is written entirely in terms of minimal transitive k-cycle factorizations, and not their associated polymaps, as is done here. To describe the results contained therein, some notation is required. For $k \ge 2$, let $\Psi_{m,k}(\mathbf{x})$ be the series obtained by setting $u = p_k = 1$ and $p_i = 0$, for $i \ne k$, in $\Psi_m(\mathbf{x}, \mathbf{p}, u)$. For instance, (3.12) and (3.15) imply that

$$x \frac{d}{dx} \Psi_{1,k}(x) = s,$$

where $s \in \mathbb{Q}[[x]]$ is the unique series solution of the functional equation

$$s = xe^{s^{k-1}}. \tag{3.16}$$

Of course, $\Psi_{m,k}$ counts minimal transitive k-cycle factorizations of permutations composed of m disjoint cycles.

The main result of [31] is a recursive differential equation satisfied by the specialized series $\Psi_{m,k}(\mathbf{x})$ for $m \geq 1$, where $k \geq 2$ is fixed. When $k = 2$, the equation coincides with that given in Theorem 2.5.1. However, for general $k \geq 2$, the terms of the equation are indexed by certain two-coloured trees on k edges, which themselves correspond to factorizations of full cycles in $\mathfrak{S}_k$. In our language, these trees parameterize the possible incidences of a k-gon in a m-faced polymap. The equations are solved easily when $m = 1, 2$, but significantly more effort is required to obtain $\Psi_{3,k}$. For $m \geq 4$ the expressions involved appear intractable. Thus the following partial result is currently the best that is known. (The formula for $\Psi_{2,k}(\mathbf{x})$ given in [31] is off by a factor of $k - 1$.)

Theorem 3.3.10. *Fix $k \geq 2$. Let s be defined as in* (3.16)*, let $s_i = s(x_i)$ for $i \geq 1$, and set*

$$\begin{aligned} F_{1,k}(s_1) &= 1, \\ F_{2,k}(s_1, s_2) &= (k-1)\left(h_{k-2}(s_1, s_2)\right)^2, \\ F_{3,k}(s_1, s_2, s_3) &= \left(h_{k-3}(s_1, s_2, s_3) + (k-1)h_{2k-4}(s_1, s_2, s_3)\right)^2, \end{aligned}$$

where $h_j(z_1, \ldots, z_m)$ is the complete symmetric function of total degree j. For $m = 1, 2, 3$, we have

$$\Psi_{m,k}(\mathbf{x}) = \left(\sum_{i=1}^{m} x_i \frac{\partial}{\partial x_i}\right)^{m-3} F_{m,k}(s_1, \ldots, s_m) \prod_{i=1}^{m} x_i \frac{ds_i}{dx_i}. \tag{3.17}$$

□

When $k = 2$, we have $F_{1,k} = F_{2,k} = F_{3,k} = 1$, and the theorem is seen to be a special case of Theorem 2.3.9. It is conjectured that, for suitable symmetric polynomials $F_{m,k}$ dependent on the parameter k, the identity (3.17) holds for all $m \geq 1$.

The methods employed in [31] are generally more transparent when interpreted in the context of polymaps, but no real progress has been made by this change of view. We expect that it should be tedious, but not fundamentally difficult, to extend the proof of Theorem 3.3.10 to obtain an expression for $\Psi_3(\mathbf{x}, \mathbf{p}, u)$. This work has not yet been done. The only higher genus analogues known for any of these results are those that can be obtained through specialization of the arbitrary

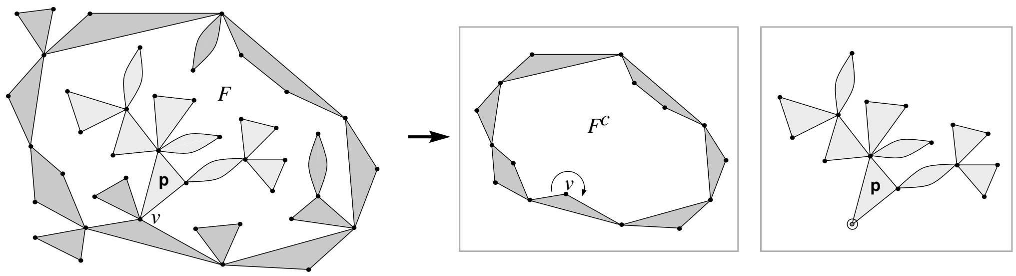

Polymap $\mathscr{M}$ Core $\mathscr{M}^c$ Branch C

Figure 3.13: The core of a polymap and one of its branches.

genus extension [58] of Theorem 3.2.9.

3.3.5 Pruning Cacti

Since cacti are the natural polymap analogues of trees, it comes as no surprise that the tree pruning bijection generalizes to the pruning of cacti from polymaps corresponding to cycle factorizations. The necessary definitions and constructions are essentially the same as those given earlier.

Definition 3.3.11. *A* ***leaf*** *of a polymap is a polygon that shares exactly one vertex with another polygon. A polymap is* ***smooth*** *if it does not have any leaves. If the polymap $\mathscr{M}$ is not a cactus, then iteratively removing the leaves of $\mathscr{M}$ results in a unique smooth polymap that we call the* ***core*** *of $\mathscr{M}$ and denote by $\mathscr{M}^c$.Labels of $\mathscr{M}$ are inherited by $\mathscr{M}^c$ in the obvious way.*

Let $\mathscr{M}$ be any polymap that is not a cactus. Let p be a polygon of $\mathscr{M}$ that shares only one vertex, v, with the core $\mathscr{M}^c$. Let F be the unique face of $\mathscr{M}$ incident with p. Separating p from v results in two components, one of which is a rooted cactus C whose root vertex is incident only with the polygon p. (If $\mathscr{M}$ is vertex-labelled, then the non-root vertices of C are labelled, but its root is not.) The cactus C is called a **branch** of face F, and the polygon p is its **stem**. We say that vertex v is the **base vertex** of this branch. If F^c is the face of $\mathscr{M}^c$ corresponding to F, then the corner of F^c at which p was attached is called the **base corner** of C. See Figure 3.13 for an illustration of these constructions, where the arrow indicates the base corner of branch C.

If the boundary walks of $\mathscr{M}$ can be normally indexed, then the *index* of the branch C is defined exactly as before. Suppose F^c has d descents and let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be its normally indexed boundary walk. Let (e_{b-1}, v, e_b) be the base corner of C, where $0 \leq b \leq k$. Then $(e_{b-1}, p, e_b)^\circ$ is a subsequence of the rotator of v, and is therefore increasing. However, by Lemma 2.6.5,

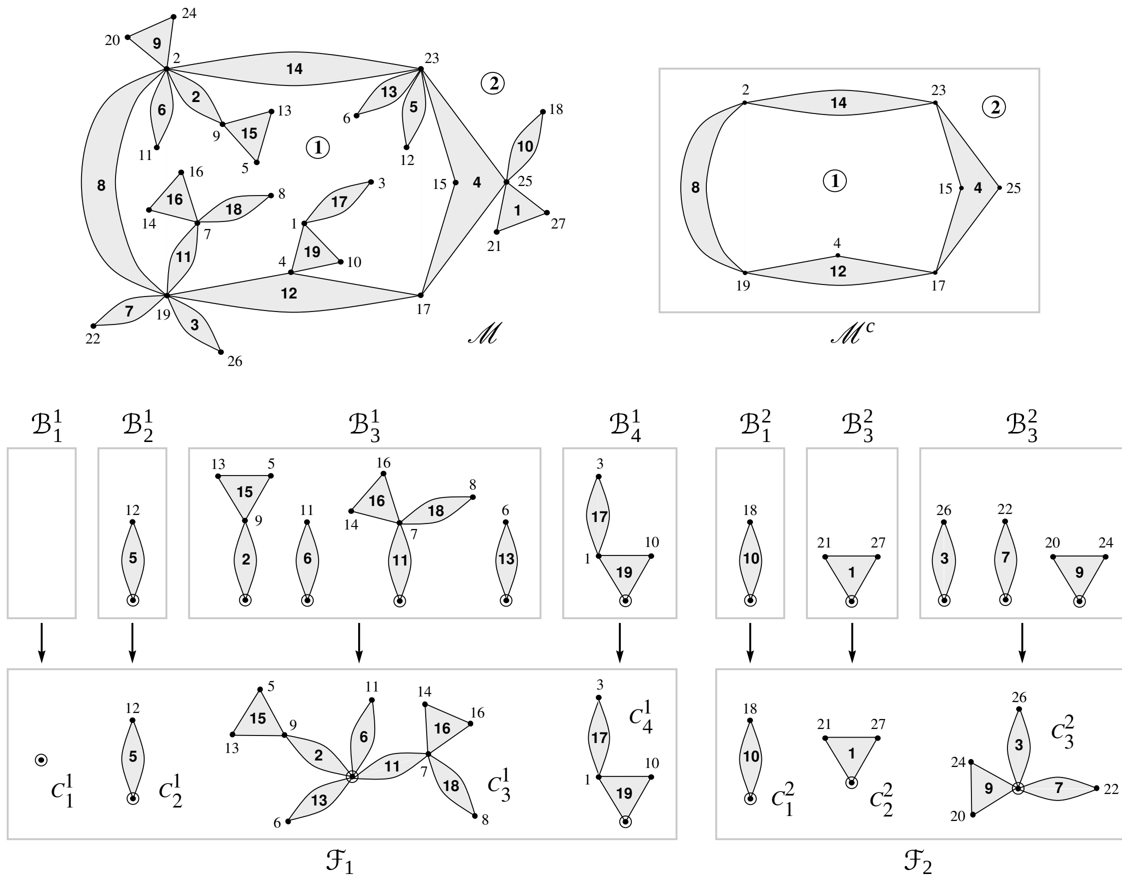

Figure 3.14: Pruning cacti from a polymap.

$(e_{j-1}, p, e_j)^\circ$ is increasing for exactly d values of j with $0 \le j \le k$, say $j_1 < \cdots < j_d$. The **index** of C is the unique value of $i \in \{1, \dots, d\}$ such that $j_i = b$.

With these definitions, the tree pruning bijection (Theorem 2.6.7) is readily extended to a cactus pruning bijection for polymaps, as follows. Let $\mathscr{M}$ be a properly labelled genus g polymap of descent class $\alpha = (\alpha_1, \dots, \alpha_m)$, and let $\theta = (\theta_1, \dots, \theta_m)$ be the descent class of $\mathscr{M}^c$. For $1 \le s \le m$ and $1 \le i \le \theta_s$, let $\mathcal{B}_i^s$ be the set of all branches of index i in face s of $\mathscr{M}$. Assemble all the cacti of $\mathcal{B}_i^s$ into a single rooted cactus C_i^s by identifying their root vertices, and then let $\mathcal{F}_s$ be the ordered forest $(C_1^s, \dots, C_{\theta_s}^s)$. This gives a reversible decomposition of $\mathscr{M}$ into the smooth polymap $\mathscr{M}^c$ of descent class θ and the m-tuple $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ of ordered forests of rooted cacti. See Figure 3.14 for an example of this decomposition.

Definition 3.3.12. *For a vector* $\mathbf{i} = (i_2, i_3, \ldots)$ *of nonnegative integers and a composition* α*, let* $S_g(\theta\,;\,\mathbf{i})$ *denote the number of smooth, properly labelled, genus g polymaps of descent class* α *that have polygon index* $\mathbf{i}$*. For fixed* $m \geq 1$*, let*

$$\Gamma_m^{(g)}(\mathbf{z}, \mathbf{p}, u) = \sum_{k\geq 1} \sum_{\mathbf{i}\geq \mathbf{0}} \sum_{\substack{\theta \models k \\ \ell(\theta)=m}} S_g(\theta\,;\,\mathbf{i})\, \mathbf{p}^{\mathbf{i}} \frac{\mathbf{z}^{\boldsymbol{\theta}}}{\boldsymbol{\theta}!} \frac{u^{r(\mathbf{i})}}{r(\mathbf{i})!},$$

be the generating series for the numbers $\{S_g(\theta\,;\,\mathbf{i}) \;:\; \ell(\theta) = m\}$*, where* $\mathbf{z} = (z_1, \ldots, z_m)$ *and* $r(\mathbf{i}) = i_2 + i_3 + \cdots$*. We typically write* Γ_m *for the genus 0 series* $\Gamma_m^{(0)}$*.*

Proposition 3.3.6 implies the genus 0 series Γ_m can generally be viewed as the generating series for smooth planar polymaps with m labelled faces.

The cacti pruning bijection described above leads to the following polymap generalization of Theorem 2.6.10. Its proof is essentially identical to that of the earlier theorem, the only change being that the tree series of Chapter 2 is now replaced by the generating series for rooted cacti.

Theorem 3.3.13. *For* $i \geq 1$*, let* $w_i = w(x_i, \mathbf{p}, u)$*, where* w *is given by* (3.15)*. Then, for* $g \geq 0$ *and* $m \geq 1$ *with* $(g, m) \neq (0, 1)$*, we have*

$$\Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u) = \Gamma_m^{(g)}(\mathbf{w}, \mathbf{p}, u), \tag{3.18}$$

where $\mathbf{x} = (x_1, \ldots, x_m)$ *and* $\mathbf{w} = (w_1, \ldots, w_m)$*.*

Proof. Let $\mathcal{C}$ be the set of vertex-rooted cacti with labelled non-root vertices. Unhinging the polygons incident with the root of $C \in \mathcal{C}$ leaves a collection of rooted cacti whose root vertices are unlabelled and incident with only one polygon. The series $\bar{w} = \bar{w}(x, \mathbf{p}, u)$ counting such cacti was derived in §3.3.3. In particular, (3.13) implies that $e^{uP(w)} \in \mathbb{Q}[u, \mathbf{p}][[x]]$ is the generating function for $\mathcal{C}$, where x marks labelled vertices.

Each of the forests $\mathcal{F}_1, \ldots, \mathcal{F}_m$ obtained through the pruning bijection is comprised of cacti belonging to $\mathcal{C}$. The proof now proceeds exactly as it did for Theorem 2.6.10, using $(e^{uP(w_i)})^{\theta_i}$ as the generating series of the forest $\mathcal{F}_i$, for $1 \leq i \leq m$. □

For example, note that this theorem anticipates the appearance in Theorem 3.3.10 of the series s, which counts rooted cacti all of whose polygons are k-gons. Indeed, with $\Gamma_{m,k}(\mathbf{x})$ defined in the obvious way, (3.18) gives the identity $\Psi_{m,k}(x_1, \ldots, x_m) = \Gamma_{m,k}(s_1, \ldots, s_m)$ for all $m \geq 2$, $k \geq 1$, where $s_i = s(x_i)$.

In what follows, the symbols w and P are defined as in (3.15) and (3.14). That is, $P \in \mathbb{Q}[\mathbf{p}][[z]]$ is given by $P(z) = \sum_{k\geq 2} p_k z^{k-1}$, and $w = w(x, \mathbf{p}, u)$ is the generating series for rooted cacti, implicity defined through the functional equation $w = xe^{uP(w)}$. Implicitly differentiating this equation yields

$$x\frac{dw}{dx} = \frac{w}{1 - uwP'(w)} = w\,\xi(w, \mathbf{p}, u), \tag{3.19}$$

where we have defined the series $\xi \in \mathbb{Q}[u, \mathbf{p}][[z]]$ by

$$\xi(z, \mathbf{p}, u) = \frac{1}{1 - uzP'(z)}.$$

Dependence of these series on $\mathbf{p}$ and u will be assumed, and we shall henceforth write $w(x)$ and $\xi(z)$ for $w(x, \mathbf{p}, u)$ and $\xi(z, \mathbf{p}, u)$, respectively, whenever it is convenient to do so. Also, for $i \geq 1$, we define the symbols w_i, ξ_i, P_i, and P_i'as follows:

$$w_i = w(x_i), \qquad \xi_i = \xi(w_i), \qquad P_i = P(w_i), \qquad \text{and} \quad P_i' = P'(w_i).$$

Thus, for instance, $x_i(dw_i/dx_i) = w_i\xi_i = w_i/(1 - uw_iP_i')$.

3.3.6 Two-Face Smooth Planar Polymaps

We now show how the methods of §2.7.1 can be extended to enumerate smooth, planar, properly labelled, two-face polymaps. With Theorem 3.3.13, this leads to an expression for $\Psi_2(x_1, x_2, \mathbf{p}, u)$, the series counting minimal transitive cycle factorizations of class (n_1, n_2) with respect to cycle index.

Theorem 3.3.14.

$$\Gamma_2(z_1, z_2, \mathbf{p}, u) = \log\left(\frac{z_1 - z_2}{z_1e^{-uP(z_1)} - z_2e^{-uP(z_2)}}\right) - u\left(\frac{z_1P(z_1) - z_2P(z_2)}{z_1 - z_2}\right).$$

Proof. We interpret Γ_2 as the generating series for smooth planar polymaps with two labelled faces. Let $\mathscr{M}$ be such a map, say with r distinctly labelled polygons. Observe that $\mathscr{M}$ is simply a closed chain of polygons, each incident with exactly two others. Let $(l_1, \ldots, l_r)^\circ$ be the cyclic list of polygon labels encountered along the boundary walk of face 1, and set $\gamma_i = (j_1, j_2)$ if the polygon labelled l_i is a $(j_1 + j_2)$-gon that has $j_s - 1$ vertices incident only with face s, for $s = 1, 2$. Then $\mathscr{M}$ is completely described by the cyclic sequence $(l_1, \gamma_1, \ldots, l_r, \gamma_r)^\circ$. We say that a vertex incident with only one face is *internal* to that face; all other vertices are said to be *extremal*. For example,

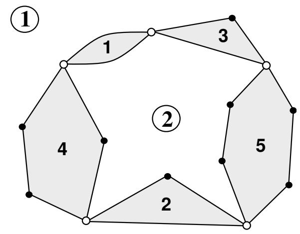

Figure 3.15: A smooth, planar, two-face polymap.

the polymap shown in Figure 3.15 corresponds to the sequence

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

Extremal vertices are coloured white in the diagram.

Clearly all vertices internal to a given face are at descents of that face. Therefore, temporarily ignoring its incident extremal vertices, a polygon with $j_s - 1$ vertices internal to face s, for $s = 1, 2$, contributes $up_{j_1+j_2} z_1^{j_1-1} z_2^{j_2-1}$ to the series $\Gamma_2(z_1, z_2, \mathbf{p}, u)$. Sum over $j_1, j_2 \geq 1$ to define

$$\delta = \sum_{j_1, j_2 \geq 1} up_{j_1+j_2} z_1^{j_1-1} z_2^{j_2-1} = u \sum_{k \geq 2} p_k \left(\frac{z_1^{k-1} - z_2^{k-1}}{z_1 - z_2} \right) = u \frac{P(z_1) - P(z_2)}{z_1 - z_2}.$$

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 fall of the cyclic permutation $(l_1, \ldots, l_r)^\circ$, and at a descent of face 2 otherwise. By Lemma 2.7.1, we therefore have

$$\Gamma_2(z_1, z_2, \mathbf{p}, u) = \log \left(\frac{x - y}{xe^y - ye^x} \right) \bigg|_{x = z_1 \delta,\ y = z_2 \delta} = \log \left(\frac{z_1 - z_2}{z_1 e^{z_2 \delta} - z_2 e^{z_1 \delta}} \right).$$

Rearranging this expression using the identities

$$z_1 \delta = u \left(\frac{z_1 P(z_1) - z_2 P(z_2)}{z_1 - z_2} \right) - uP(z_2), \quad z_2 \delta = u \left(\frac{z_1 P(z_1) - z_2 P(z_2)}{z_1 - z_2} \right) - uP(z_1)$$

gives the desired result. □

Corollary 3.3.15.

$$\left(x_1 \frac{\partial}{\partial x_1} + x_2 \frac{\partial}{\partial x_2} \right) \Psi_2(x_1, x_2, \mathbf{p}, u) = u^2 w_1 w_2 \xi_1 \xi_2 \left(\frac{P_1 - P_2}{w_1 - w_2} \right) \left(\frac{w_1 P_1' - w_2 P_2'}{w_1 - w_2} \right)$$

Proof. Let $\Psi_2 = \Psi_2(x_1, x_2, \mathbf{p}, u)$. Then combining the above result with Theorem 3.3.13 and (3.15) immediately yields

$$\Psi_2 = \log\left(\frac{w_1 - w_2}{x_1 - x_2}\right) - u\frac{w_1 P_1 - w_2 P_2}{w_1 - w_2}.$$

Differentiating, and simplifying using (3.19), gives

$$x_1\frac{\partial \Psi_2}{\partial x_1} = \frac{w_1}{w_1 - w_2}\left(1 + u\xi_1 w_2 \frac{P_1 - P_2}{w_1 - w_2}\right) - \frac{x_1}{x_1 - x_2},$$
$$x_2\frac{\partial \Psi_2}{\partial x_2} = \frac{w_2}{w_2 - w_1}\left(1 + u\xi_2 w_1 \frac{P_1 - P_2}{w_1 - w_2}\right) - \frac{x_2}{x_2 - x_2}.$$

Adding these expressions results in

$$\left(x_1\frac{\partial}{\partial x_1} + x_2\frac{\partial}{\partial x_2}\right)\Psi_2 = u w_1 w_2 \left(\frac{\xi_1 - \xi_2}{w_1 - w_2}\right)\left(\frac{P_1 - P_2}{w_1 - w_2}\right). \tag{3.20}$$

Observe that $\xi_1 - \xi_2 = u\xi_1\xi_2(w_1 P_1' - w_2 P_2')$ to complete the proof. □

As a special case of Corollary 3.3.15, we can quickly derive the $m = 2$ case of Theorem 3.3.10.

Corollary 3.3.16. *Fix $k \geq 2$, and define $s \in \mathbb{Q}[[x]]$ as in* (3.16). *Then*

$$\left(x_1\frac{\partial}{\partial x_1} + x_2\frac{\partial}{\partial x_2}\right)\Psi_{2,k}(x_1, x_2) = (k-1)\,(h_{k-2}(s_1, s_2))^2\; x_1 x_2 \frac{ds_1}{dx_1}\frac{ds_2}{dx_2},$$

where $s_1 = s(x_1)$ and $s_2 = s(x_2)$.

Proof. Set $p_i = 0$ for $i \neq k$, and $u = p_k = 1$ in Corollary 3.3.15. These substitutions reduce w_i to s_i, P_i to s_i^{k-1}, $w_i P_i'$ to $(k-1)s_i^{k-1}$, and $w_i\xi_i$ to $x_i(ds_i/dx_i)$. Thus

$$\left(x_1\frac{\partial}{\partial x_1} + x_2\frac{\partial}{\partial x_2}\right)\Psi_{2,k}(x_1, x_2) = (k-1)\left(\frac{s_1^{k-1} - s_2^{k-2}}{s_1 - s_2}\right)^2 x_1 x_2 \frac{ds_1}{dx_1}\frac{ds_2}{dx_2},$$

and the result follows. □

3.3.7 Attaching Digons to a Polymap

The material of §2.7.2 is readily modified to describe the addition of digons (*i.e.* 2-gons, or fattened edges) to a polymap. Indeed, let F be a face of the polymap $\mathscr{M}$ with boundary walk $((v_0, e_0), \ldots, (v_k, e_k))^\circ$, and let $c_i = (e_{i-1}, v_i, e_i)$ and $c_j = (e_{j-1}, v_j, e_j)$ be distinct corners of F. Let $g \in \mathbb{R}$ be distinct from $e_0, \ldots, e_k$. If $(e_{i-1}, g, e_i)^\circ$ and $(e_{j-1}, v_j, e_j)^\circ$ are both increasing,

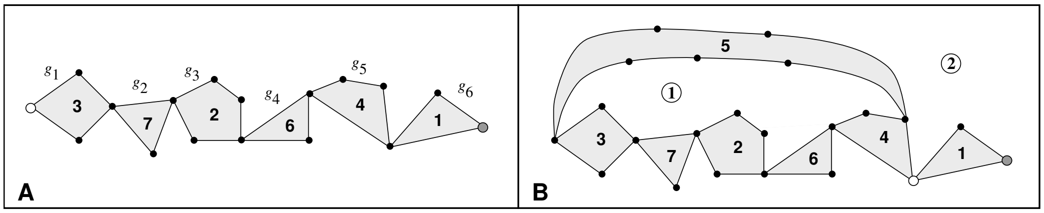

Figure 3.16: (A) A polygonal path. (B) A two-face polymap with a tail.

then the polymap $\mathscr{M} \oplus (c_i, c_j)^g_{s,t}$ is created by first attaching a digon labelled g between c_i and c_j, and then assigning labels s and t, respectively, to the faces of the resulting polymap containing corners (g, v_i, e_i) and (g, v_j, e_j). Since the descents of F are split between these two faces, the proof of Lemma 2.7.6 remains valid in this setting.

These observations can be used to give a bijective proof of Corollary 3.3.15 that exactly mimics our earlier proof of Corollary 2.7.13. The method is outlined below. We emphasize that the bulk of the work has already been done, in §2.7.3. In particular, the general nature of our proof of Theorem 2.7.11 makes it applicable in the current context, essentially without change. We need only define suitable polymap generalizations of *ordered paths* and *tails*.

Let $(g_1, \ldots, g_l)$ be a list of l distinctly labelled polygons. Choose two distinct vertices of each of these polygons, calling one the **top** and the other the **bottom**, and join the polygons in the plane by identifying the top of g_i with the bottom of g_{i+1}, for $1 \leq i < l$. This results in a cactus composed of l polygons, each incident with at most one other. Now distinguish the bottom of g_1 and the top of g_l by colouring them white and grey, respectively. We call the resulting structure an (ordered) **polygonal path** of length l, and refer to the white and grey vertices as its **ends**. For example, Figure 3.16A shows a polygonal path of length 6. Taking cyclic symmetry into account, there are $k-1$ ways of choosing the top and bottom of a k-gon once its head is chosen. Therefore

$$z \sum_{l \geq 1} \left(\sum_{k \geq 2} (k-1) u p_k z^{k-1} \right)^l = z \left(\frac{1}{1 - uzP'(z)} \right) = z(\xi(z) - 1) \tag{3.21}$$

is the generating series for polygonal paths of length at least one, where z marks vertices, and $\mathbf{p}$ and u record polygon index and labelled polygons, as usual.

As a natural extension of Definition 2.7.9, we say a polymap $\mathscr{M}$ has a **tail** in face F if either (1) $\mathscr{M}$ is smooth and a vertex at a descent of F has been coloured grey, or (2) $\mathscr{M}$ contains only one branch, which is a polygonal path in face F whose white end is the base vertex of the branch.

For example, the two-face polymap shown in Figure 3.16B has a tail in face 2. Using (3.21), a derivation similar to that of Lemma 2.7.10 shows the counting series for planar polymaps with m labelled faces and a tail in face i to be

$$\xi(z_i) z_i \frac{\partial}{\partial z_i} \Gamma_m(\mathbf{z}, \mathbf{p}, u). \tag{3.22}$$

Unsurprisingly, polymaps with tails of fixed descent class can be constructed from polygonal paths by the addition of a polygon, as follows. Fix $\theta = (\theta_1, \theta_2) \models n$. Let $\mathscr{P}$ be a polygonal path containing n vertices, let λ be distinct from the polygon labels of $\mathscr{P}$, and let d_1, d_2 be any nonnegative integers. Now let the sole face F of $\mathscr{P}$ have boundary walk $((v_0, e_0), \ldots, (v_k, e_k))^\circ$, where v_0 and v_m are the white and grey ends of $\mathscr{P}$, respectively. For $0 \leq i \leq k$, let $c_i = (e_{i-1}, v_i, e_i)$. This setup is illustrated below.

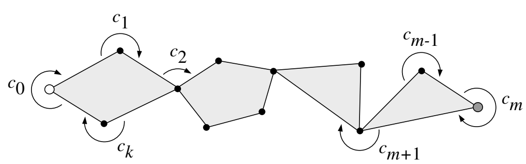

Clearly F has n descents and $c_0 \in \mathcal{A}_F(\lambda)$, so Lemma 2.7.6 guarantees a unique corner $c_r \in \mathcal{A}_F(\lambda)$, with $0 < r \leq k$, such that the two-face map $\mathscr{M} = \mathscr{P} \oplus (c_0, c_r)^{\lambda}_{1,2}$ is of descent class $(\theta_1, n - \theta_1) = (\theta_1, \theta_2)$. Strip v_0 of its colour and, if $r \neq m$, colour the bottom of the polygon incident with edge e_r white. Then $\mathscr{M}$ is of descent class θ and has a tail. Finally, observe that $\mathscr{M}$ can made to be of descent class $(\theta_1 + d_1, \theta_2 + d_2)$ simply by transforming the newly added digon into a $(d_1 + d_2 + 2)$-gon with exactly d_i of its vertices incident only with face i, for $i = 1, 2$.

Example 3.3.17. The polygonal path $\mathscr{P}$ in Figure 3.17A contains 17 vertices, and the crosses mark the corners at which a digon labelled $\lambda = 5$ could be attached to this path. The enlarged cross indicates the unique corner c such that $\mathscr{P} \oplus (c_0, c)^5_{1,2}$ is of descent class $(7, 10) \models 17$. The polymap $\mathscr{P} \oplus (c_0, c)^5_{1,2}$, itself, is drawn in Figure 3.17B, and crosses there mark the 7 descents of face 1. Notice that transforming the additional digon into a 7-gon results in the polymap of descent class $(10, 12)$, with a tail, shown in Figure 3.16B.

Similarly, the polygonal path in Figure 3.17C has 15 vertices. Crosses mark corners at which a digon labelled $\lambda = 3$ could be attached, and the large cross is the unique such corner at which attachment yields a polymap of descent class $(9, 6) \models 15$. This smooth polymap with a tail is shown in Figure 3.17D. □

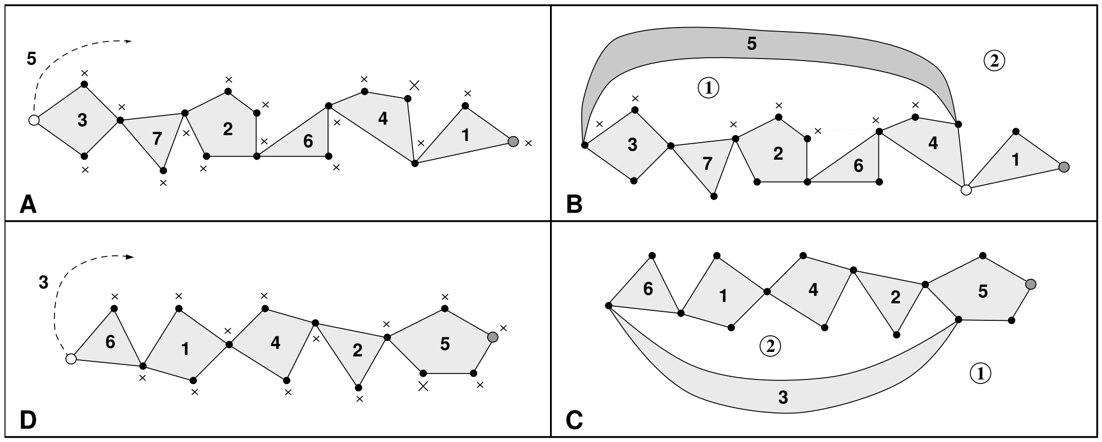

Figure 3.17: Creation of a two-face polymap with a tail.

The correspondence $(\mathscr{P}, \lambda, d_1, d_2) \mapsto \mathscr{M}$ described above is a straightforward adaptation of the bijection Ω_θ defined in Theorem 2.7.11. Modifying that proof in the obvious way shows that this correspondence is a polygon-preserving bijection between planar polymaps of descent class $(\theta_1+d_1, \theta_2+d_2)$ with a tail, and polygon-labelled structures $(\mathscr{P}, \lambda, d_1, d_2)$, where $\mathscr{P}$ is a polygonal path on $\theta_1+\theta_2$ vertices and (λ, d_1, d_2) describes a labelled (d_1+d_2+2)-gon. Therefore, from (3.21), we see that $\mathscr{P}$ contributes the factor

$$z(\xi(z)-1) \circ \Delta^+(z\,;\, z_1, z_2) = z_1 z_2 \frac{\xi(z_1)-\xi(z_2)}{z_1-z_2}$$

to the series for two-face polymaps with a tail, while (λ, d_1, d_2) induces the factor

$$\sum_{d_1, d_2 \geq 0} u p_{d_1+d_2+2} z_1^{d_1} z_2^{d_2} = u \sum_{k \geq 2} p_k \frac{z_1^{k-1}-z_2^{k-1}}{z_1-z_2} = u \frac{P(z_1)-P(z_2)}{z_1-z_2}.$$

From (3.22) there follows

$$\left(z_1 \xi(z_1) \frac{\partial}{\partial z_1} + z_2 \xi(z_2) \frac{\partial}{\partial z_2} \right) \Gamma_2(z_1, z_2, \mathbf{p}, u) = u z_1 z_2 \frac{P(z_1)-P(z_2)}{z_1-z_2} \cdot \frac{\xi(z_1)-\xi(z_2)}{z_1-z_2}.$$

Finally, observe that (3.19) implies $w_i \xi(w_i) \frac{\partial}{\partial w_i} = x_i \frac{\partial}{\partial x_i}$. Hence (3.20) is obtained by replacing z_i with w_i in the expression above and applying Theorem 3.3.13, and our alternate proof of Corollary 3.3.15 is complete. Of course, the advantage of this proof is that it assigns simple combinatorial meaning to each of the factors appearing in (3.20).

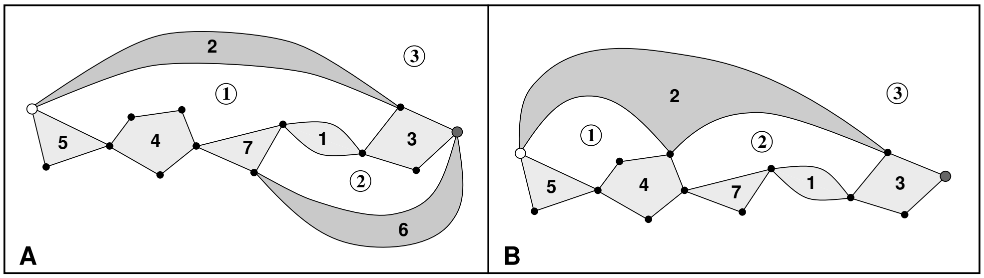

Figure 3.18: Creation of three-face smooth polymaps.

Smooth three-face polymaps can similarly be built by attaching digons to the ends of a polygonal path, as shown in Figure 3.18A. However, not all smooth three-face polymaps can be created in this way. For example, this construction does not account for those maps containing a polygon incident with all faces, such as the one drawn in Figure 3.18B. This is the same complication that was discussed in §3.3.4, and it makes the quest for a polymap analogue of Corollary 2.7.16 technically involved.

3.3.8 Additional Notes

The failure of Proposition 3.3.6 in positive genus is reflected by the lack of a cycle factorization analogue of Proposition 2.4.15. In fact, there exist transitive k-cycle factorizations $f = (\sigma_r, \ldots, \sigma_1)$ such that $f = (\rho\sigma_r\rho^{-1}, \ldots, \rho\sigma_1\rho^{-1})$ for permutations ρ other than the identity. For instance, consider the transitive 4-cycle factorization

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

Notice that each of the three factors on the right-hand side is invariant under conjugation by $\rho = (1\,2)(3\,4)(5\,6)(7\,8)$. Incidentally, this factorization is of genus 1, and was obtained by assigning vertex labels to the rightmost polymap in Figure 3.10. Proving the genus 0 cycle factorization analogue of Proposition 2.4.15 directly (*i.e.* working only within the symmetric group) seems to be tedious.

Theorem 3.3.7 has also appeared in [65], where Springer derives it by using the same decomposition of factorizations into cacti as we do here. He counts these cacti by way of a bijection that generalizes Prüfer's [59] encoding of trees.

3.4 The Double Hurwitz Problem

Recall that the *Hurwitz Enumeration Problem*, studied extensively in Chapter 2, asks for the number of genus g factorizations of a fixed permutation into transpositions. In this final section of Chapter 3, we discuss a particular generalization of this question known as the *Double Hurwitz Problem*. Formal definitions will be given below, but the distinction between the problems is quite simply put as follows.

Rather than counting factorizations $\pi = \tau_r \dots \tau_1$ of a permutation π into transpositions τ_i, we now count factorizations $\pi = \sigma\tau_r \cdots \tau_1$, where again the τ_i are transpositions, but the last factor σ is forced to be of some fixed, but arbitrary, cycle type. Factorizations of this sort have geometrical significance in terms of ramified coverings of the sphere. As outlined in §2.3.6, a genus g factorization $\pi = \sigma\tau_r \cdots \tau_1$ in $\mathfrak{S}_n$ corresponds to an n-sheeted branched covering of the sphere by a Riemann surface of of genus g, with $r + 2$ branch points $\{0, \infty, P_1, \dots, P_r\}$ having simple branching over the P_i, and branching over 0 and ∞ specified by the cycle types of π and σ. See also the Additional Notes at the end of this section.

3.4.1 β-Factorizations

Let $\pi \in \mathfrak{S}_n$ and $\beta \vdash n$. A ***β*-factorization** of π is a transitive factorization $(\sigma, \tau_r, \dots, \tau_1)$ of π such that $\sigma \in \mathscr{C}_\beta$ and each τ_i is a transposition. We allow $r = 0$, but note that only the trivial factorization $\pi = \sigma$ of a full cycle π is possible in this case. As a nontrivial example, consider

$$(1\,2\,3\,4\,5)(6\,7)(8)(9\,10) = (1\,6\,5)(2\,7)(3\,4\,10\,9)(8) \cdot (4\,8)(1\,7)(2\,8)(4\,6)(2\,9)(6\,8). \tag{3.23}$$

This is a $(4, 3, 2, 1)$-factorization of $(1\,2\,3\,4\,5)(6\,7)(8)(9\,10)$ of class $(5, 2, 2, 1)$. Notice that (3.3) gives $r = \ell(\pi) + \ell(\beta) + 2g - 2$ for a β-factorization of π of genus g. Thus the factorization above is of genus 0.

Definition 3.4.1. *Let α, β be partitions, and let $g \geq 0$. Then we write $H_g(\alpha, \beta)$ for the number of β-factorizations of genus g and class α, and we let $r_g(\alpha, \beta)$ denote the number $\ell(\alpha)+\ell(\beta)+2g-2$ of transposition factors any such factorization contains. The numbers $H_g(\alpha, \beta)$ are known as* ***double Hurwitz numbers****.*

The double Hurwitz numbers are symmetrical, in the sense that $H_g(\alpha, \beta) = H_g(\beta, \alpha)$ for all partitions α, β. This is immediate from the fact that $\pi = \sigma\tau_r \cdots \tau_1$ is equivalent to $\sigma = \pi\tau_1 \cdots \tau_r$. Also, since $\iota = \sigma\tau_r \cdots \tau_1$ if and only if $\sigma^{-1} = \tau_r \cdots \tau_1$, the double Hurwitz number $H_g([1^n], \alpha)$

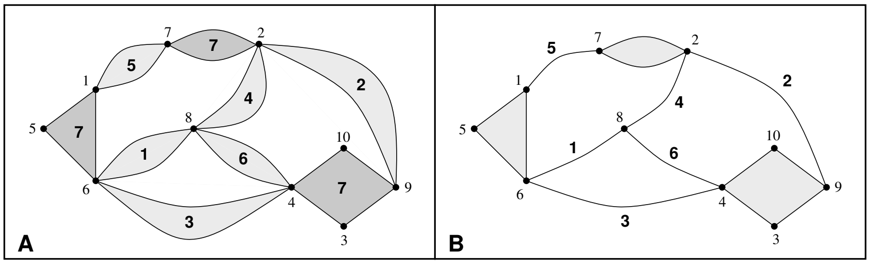

Figure 3.19: The hybrid map of a double Hurwitz factorization.

is equal to the (single) Hurwitz number $H_g(\alpha)$, for which we have the simple formula of Theorem 2.3.3. (Note that genus is preserved because $r_g(\alpha, [1^n]) = \ell(\alpha) + n + 2g - 2 = r_g(\alpha)$.) Thus far, we have seen how to evaluate only one other special class of double Hurwitz numbers. In particular, we have the formula

$$H_0((n), \beta) = n^{\ell(b)-1} \frac{(\ell(\beta) - 1)!}{|\operatorname{Aut}(\beta)|} \tag{3.24}$$

for arbitrary $\beta \vdash n$, which comes from specializing Theorem 3.2.9.

3.4.2 Hybrid Maps

Let $f = (\sigma, \tau_r, \ldots, \tau_1)$ be a β-factorization, where $\beta = [1^{b_1} 2^{b_2} \cdots] \vdash n$. Then the reduced polymap $\mathscr{M}_f^\dagger$ has r digons distinctly labelled $1, \ldots, r$, and, for each $i \geq 1$, exactly b_i i-gons labelled $r + 1$, where a 1-gon is interpreted as a vertex not incident with any polygon. (This caveat is an artifact of our elimination of loops.) We call a polymap with this structure a **hybrid map** of **polygon type $\boldsymbol{\beta}$**. The rationale behind this terminology will be made clear momentarily.

The hybrid map of polygon type $\beta = (4, 3, 2, 1)$ corresponding to the β-factorization (3.23) is shown in Figure 3.19A. Notice that no information in the diagram is lost by suppressing the labels of the polygons with maximal label $r+1 = 7$, and flattening the remaining labelled digons to edges, as illustrated in Figure 3.19B. This is true in general, and we shall always draw hybrid maps in this simplified manner. Thus a hybrid map of polygon type β appears to consist of $\ell(\beta)$ unlabelled polygons (including 1-gons) joined together by distinctly labelled edges. That is, hybrid maps are essentially hybrids of polymaps and the edge-labelled maps of Chapter 2.

We slightly alter our terminology to reflect this view of hybrid maps. By a **polygon** of a hybrid map $\mathscr{M}$ we always mean either a polygon with maximal label, or a vertex not incident with any such polygon. Vertices of this sort are also called **1-gons**. We refer to the remaining digons of $\mathscr{M}$

as its **simple edges**. Thus every vertex is incident with a unique polygon, and every simple edge is incident with one or two (white) faces. Finally, it is both convenient and suggestive to regard the polygons of a hybrid map, and their bounding edges, as having label ∞. This reflects the fact that the polygons of hybrid maps are always maximally labelled, and suggests that our interest lies in the labelling of simple edges. For the purposes of descent structure, ∞ is interpreted as some fixed integer larger than all other labels. For example, the **rotator** of vertex 4 in the hybrid map of Figure 3.19B is $(3, 6, \infty)^\circ$, which is, of course, increasing. The list of edge labels encountered along the boundary walk of the outer face of this map is $(2, \infty, \infty, 3, \infty, \infty, 5, \infty)^\circ$, so the outer face has 5 descents.

The following result is an immediate consequence of Theorem 3.2.8 and these conventions.

Theorem 3.4.2. *Let $\alpha, \beta \vdash n$ and $g \geq 0$. There is a bijection between β-factorizations of genus g and class α, and vertex-labelled hybrid maps of genus g with descent partition α and polygon type β. Under this bijection, a factorization $(\sigma, \tau_r, \ldots, \tau_1)$ of π corresponds with a hybrid map with n vertices, $\ell(\beta)$ polygons, and r simple edges, whose descent cycles coincide with the cycles of π.* □

If the faces of a hybrid map are labelled, then we speak of its *descent class* rather than its descent partition. A vertex- and face-labelled hybrid map $\mathscr{M}$ is **properly labelled** if it is of descent class α and the face labelled s has descent set $\mathbb{D}_s(\alpha)$, for $1 \leq s \leq m$. The next definition should, by now, be familiar.

Definition 3.4.3. *For a partition β and a composition α, let $M_g(\alpha, \beta)$ denote the number of properly labelled genus g hybrid maps of polygon type β and descent class α. For fixed $m \geq 1$, let*

$$\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u) = \sum_{n \geq 1} \sum_{\beta \vdash n} \sum_{\substack{\alpha \models n \\ \ell(\alpha)=m}} M_g(\alpha, \beta)\, q_\beta \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \frac{u^{r_g(\alpha,\beta)}}{r_g(\alpha, \beta)!},$$

be the generating series for the numbers $\{M_g(\alpha, \beta) \,:\, \alpha, \beta \vdash n,\ \ell(\alpha) = m\}$, where $\mathbf{q} = (q_1, q_2, \ldots)$ and $\mathbf{x} = (x_1, \ldots, x_m)$, and where $q_\beta = q_{\beta_1} q_{\beta_2} \cdots$ for the partition $\beta = (\beta_1, \beta_2, \ldots)$. We typically write Θ_m in place of $\Theta_m^{(0)}$.

Notice that $\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u)$ is naturally exponential in the indeterminates $x_1, \ldots, x_m$, which mark labelled vertices in faces $1, \ldots, m$ of a properly labelled hybrid map, and also in u, which marks labelled simple edges. The symmetrized Hurwitz series $\Psi_m^{(g)}(\mathbf{x}, u)$ of Chapter 2 is acquired from $\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u)$ by setting $q_1 = 1$, and $q_i = 0$ for $i \geq 1$.

Vertex-labelled hybrid maps have no non-trivial automorphisms, so their faces can be labelled arbitrarily without obtaining duplicate maps. The proof of Theorem 3.3.2 therefore applies, almost verbatim, to connect β-factorizations with properly labelled hybrid maps. The result is the following analogue of Corollary 3.3.5.

Theorem 3.4.4. *For any composition $\alpha = (\alpha_1, \ldots, \alpha_m)$ and partition β, we have*

$$H_g(\alpha, \beta) = \alpha_1 \cdots \alpha_m \, r_g(\alpha, \beta)! \, [\mathbf{x}^{\boldsymbol{\alpha}} q_\beta u^{r_g(\alpha,\beta)}] \, \Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u).$$

□

Since face-labelled hybrid maps with at least two faces do not admit nontrivial automorphisms, their vertices can be labelled arbitrarily without fear of duplication. Thus $\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u)$ can, for $m \geq 2$, be regarded as the counting series for face-labelled hybrid maps with respect to descent class, polygon type, and labelled simple edges.

3.4.3 Hybrid Cacti

A **hybrid cactus** is a planar hybrid map with only one face. In what follows, we shall be concerned exclusively with vertex-rooted, vertex-labelled hybrid cacti. To avoid redundancy, we refer to these maps simply as *cacti* throughout. We further introduce the term **simple cactus** to describe a cactus whose root vertex is incident only with simple edges (*i.e.* the root vertex is a 1-gon). Let $\vartheta = \vartheta(x, \mathbf{q}, u)$ and $w = w(x, \mathbf{q}, u)$, respectively, be the generating series for cacti and simple cacti with respect to labelled vertices (marked by x), polygon type (marked by $\mathbf{q}$), and labelled simple edges (marked by u). The following combinatorial decomposition of cacti closely resembles the one given in §3.3.3.

Let C be a cactus whose root is incident with an k-gon. Then deletion of this polygon results in an ordered list $(C_1, \ldots, C_k)$ of simple cacti, as depicted in Figure 3.20. Thus we have $\vartheta = Q(w)$, where the series $Q \in \mathbb{Q}[\mathbf{q}][[z]]$ is defined by

$$Q(z) = \sum_{k \geq 1} q_k z^k. \tag{3.25}$$

It follows directly from the definition of Θ_1 that

$$x \frac{d}{dx} \Theta_1(x, \mathbf{q}, u) = \vartheta = Q(w). \tag{3.26}$$

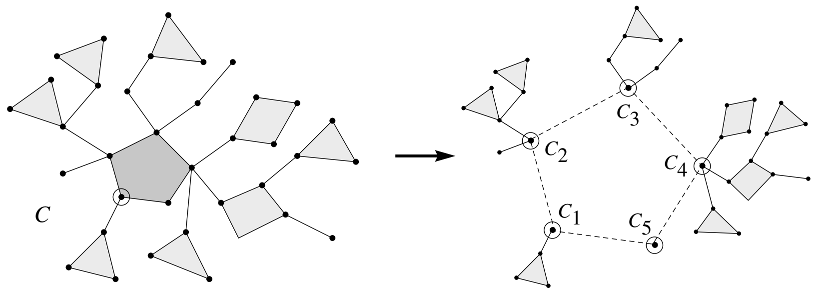

Figure 3.20: Decomposition of a hybrid cactus.

Now observe that if the root of a simple cactus is incident with m simple edges, then removing these edges leaves the root itself, together with a collection of m rooted cacti. Therefore $w = x\sum_{m\geq 1} u^m \vartheta^m/m! = xe^{u\vartheta}$. Thus we find that w is the unique series solution of the functional equation

$$w = xe^{uQ(w)}. \tag{3.27}$$

Using (3.26) and (3.27), we may apply Lagrange inversion to obtain the coefficients of the series $\Theta_1(x, \mathbf{q}, u)$, thereby obtaining a formula for $H_0((n), \beta)$ for arbitrary $\beta \vdash n$. In fact, this work was essentially carried out in Example 1.3.2, where we found that

$$[q_\beta u^r x^n]\, Q(w) = \frac{n^{\ell(\beta)-1}}{|\operatorname{Aut}(\beta)|}.$$

Combining this result with Theorem 3.4.4 does indeed yield (3.24).

3.4.4 Pruning Cacti

A **leaf** of a hybrid map is a polygon incident with at most one simple edge, and a hybrid map is **smooth** if it does not contain any leaves. As usual, leaves (and their incident simple edges) can be iteratively removed from any hybrid map $\mathscr{M}$ which is not a cactus to produce its unique smooth **core**, $\mathscr{M}^c$.Branches of $\mathscr{M}$ and their indices are now defined as before. (That is, branches are simple cacti whose roots are unlabelled and incident with exactly one simple edge.) This permits the pruning of cacti from a properly labelled hybrid map $\mathscr{M}$, with similarly indexed branches of each face being removed from $\mathscr{M}$ and joined at their roots to produce forests of simple cacti. These forests, together with the smooth map $\mathscr{M}^c$, completely specify $\mathscr{M}$.

This familiar process is illustrated once more in Figure 3.21. Notice that face 1 of the hybrid

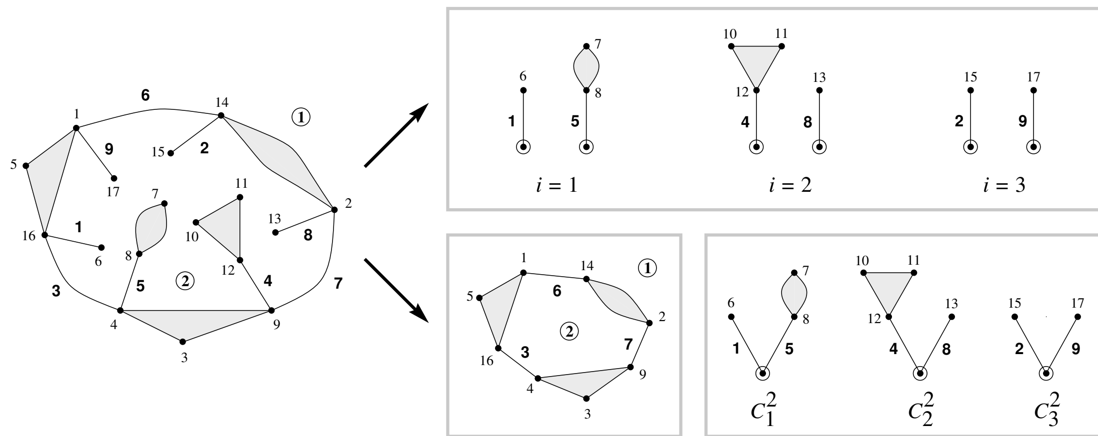

Figure 3.21: Pruning cacti from a properly labelled hybrid map.

map shown there has no branches, so all focus on face 2. Its branches are displayed along with their indices. Also shown are the three simple cacti C_1^2, C_2^2, C_3^2 that are formed when similarly indexed branches are joined together.

Mimicking the proof of Theorem 3.3.13 leads to the following expected result. (Of course, the rooted cactus series used in the proof of Theorem 3.3.13 is to be replaced with the simple cactus series w defined through (3.27).)

Theorem 3.4.5. *Let $S_g(\alpha, \beta)$ be the number of smooth properly labelled genus g hybrid maps of polygon type β and descent class α. For fixed $m \geq 1$, set*

$$\Lambda_m^{(g)}(\mathbf{x}, \mathbf{q}, u) = \sum_{n\geq 1} \sum_{\beta\vdash n} \sum_{\substack{\alpha\models n \\ \ell(\alpha)=m}} S_g(\alpha, \beta)\, q_\beta \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \frac{u^{r_g(\alpha,\beta)}}{r_g(\alpha, \beta)!}.$$

For $i \geq 1$, let $w_i = w(x_i, \mathbf{q}, u)$, where w is given by (3.27). *Then, for $g \geq 0$ and $m \geq 1$ with $(g, m) \neq (0, 1)$, we have*

$$\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u) = \Lambda_m^{(g)}(\mathbf{w}, \mathbf{q}, u),$$

where $\mathbf{w} = (w_1, \ldots, w_m)$. □

We shall be concerned exclusively with the genus 0 series, and write Λ_m in place of $\Lambda_m^{(0)}$. Throughout the remainder of this section, we use the symbols w and Q as they are defined in (3.27)

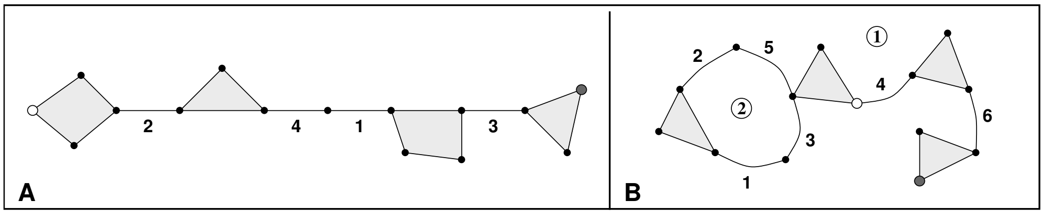

Figure 3.22: (A) A hybrid path. (B) A two-face hybrid map with a tail.

and (3.25). Implicit differentiation of (3.27) gives

$$x\frac{dw}{dx} = \frac{w}{1 - uwQ'(w)} = w\,\xi(w, \mathbf{q}, u), \tag{3.28}$$

where the series $\xi \in \mathbb{Q}[u, \mathbf{q}][[z]]$ is defined by

$$\xi(z, \mathbf{q}, u) = \frac{1}{1 - uzQ'(z)}. \tag{3.29}$$

Dependence of these series on $\mathbf{q}$ and u is assumed, and generally suppressed. Note the close similarity between these definitions and those of the same symbols in §3.3.5.

3.4.5 Combinatorial Constructions for Smooth Hybrid Maps

The material of §2.7.2 is again easily extended to allow for the addition of simple edges to a hybrid map, and hybrid map analogues of all the results recorded in §2.7 are readily obtained. These are described briefly below. As was the case in §3.3.7, the work reduces to giving suitable hybrid map analogues of ordered paths and tails.

Let $n \geq 0$ and let $(l_1, \ldots, l_n)$ be a list of n distinct positive integers. Let $(p_1, \ldots, p_{n+1})$ be a list of $n + 1$ polygons in the plane, possibly including 1-gons (*i.e.* single vertices). Distinguish one vertex of each polygon as its **top**, and a second as its **bottom**. Here we allow the top and bottom of a polygon to coincide. For $1 \leq i \leq n$, attach the top of p_i to the bottom of p_{i+1} by a simple edge with label l_i. This results in a hybrid cactus containing $n + 1$ polygons, each incident with at most two simple edges. Distinguish the bottom of p_1 and the top of p_{n+1} by colouring them white and grey, respectively. We call the resulting structure an (ordered) **hybrid path** of length n. For example, Figure 3.22A shows a hybrid path of length 4, created from the list $(2, 4, 1, 3)$. Because

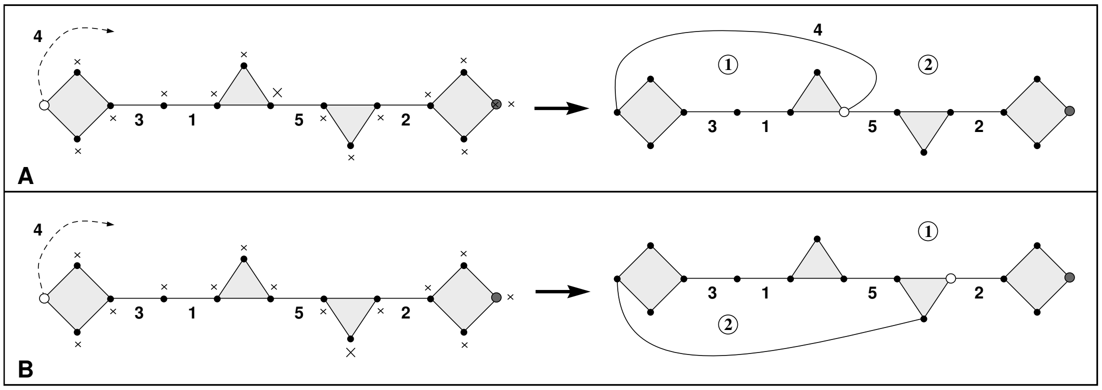

Figure 3.23: Construction of two-face hybrid maps with tails.

of cyclic symmetry, there are k ways of choosing the head and foot of a k-gon. Therefore the series

$$\sum_{n\geq 0} u^n \left(\sum_{k\geq 1} k q_k z^k \right)^{n+1} = \frac{zQ'(z)}{1 - uzQ'(z)} = \frac{1}{u}(\xi(z) - 1) \tag{3.30}$$

counts hybrid paths, with z marking vertices, u marking labelled simple edges, and $\mathbf{q}$ recording polygon type.

We say the hybrid map $\mathscr{M}$ has a **tail** in face F if either (1) $\mathscr{M}$ is smooth and a vertex at a descent of F has been distinguished, or (2) $\mathscr{M}$ contains only one branch, which is in face F, and is a hybrid path of length at least one whose white end is the base vertex of the branch. For example, the two-face hybrid map shown in Figure 3.22B has a tail in face 1. Note that the white end of a hybrid path forming a tail is always a 1-gon, and that such a path cannot be of length 0. Therefore, by (3.30), we find that the series $zu \cdot \frac{1}{u}(\xi(z) - 1) = z(\xi(z) - 1)$ counts hybrid paths of this type. Thus the series counting planar hybrid maps with m labelled faces and a tail in face i is

$$\xi(z_i) z_i \frac{\partial}{\partial z_i} \Lambda_m(\mathbf{z}, \mathbf{q}, u). \tag{3.31}$$

See the proof of Lemma 2.7.10 for further details regarding this derivation.

Attaching a labelled simple edge from the white vertex of a hybrid path to any of its other vertices plainly results in a planar two-face hybrid map with a tail. This is illustrated in Figure 3.23, where a simple edge labelled 4 is attached to a hybrid path containing 15 vertices to obtain hybrid maps with tails of descent classes $(5, 10)$ and $(10, 5)$. The following generalization of Theorem 2.7.11 ensures that this process faithfully produces all such maps.

Theorem 3.4.6. *Fix $\theta = (\theta_1, \theta_2) \models n$. There is a polygon- and edge-preserving bijection between planar face-labelled hybrid maps of descent class θ with a tail, and edge-labelled pairs $(\lambda, \mathscr{P})$ where λ is a simple edge and $\mathscr{P}$ is a hybrid path containing n vertices.*

Proof. Let $(\lambda, \mathscr{P})$ be a pair as described in the theorem. Let $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ be the boundary walk of the sole face, F, of $\mathscr{P}$, where v_0 is its white end and $(e_k, \lambda, e_0)^\circ$ is increasing. Notice that defining v_0 as the white end of $\mathscr{P}$ is not enough to uniquely define an indexing of the boundary walk of F, since this vertex may appear more than twice in the walk. However, the additional condition that $(e_k, \lambda, e_0)^\circ$ be increasing makes the indexing well-defined. Now let m be the minimal positive integer such that v_m is the grey end of $\mathscr{P}$, and follow the proof of Theorem 2.7.11. With very few (trivial) modifications, it remains valid in this setting. □

The following corollary appears in [36], though with a different proof from that given here. See the Additional Notes for further details..

Corollary 3.4.7. *Let $w_i = w(x_i, \mathbf{q}, u)$ and $\xi_i = \xi(w_i)$ for $i = 1, 2$. Then*

$$\left(x_1 \frac{\partial}{\partial x_1} + x_2 \frac{\partial}{\partial x_2}\right) \Theta_2(x_1, x_2, \mathbf{q}, u) = \frac{w_2}{w_1 - w_2}(\xi_1 - 1) + \frac{w_1}{w_2 - w_1}(\xi_2 - 1).$$

Proof. By (3.31), the series on the left-hand side counts two-face hybrid maps with a tail. Now follow the proof of Corollary 2.7.13, replacing the series $zu/(1 - uz)$ for ordered paths used there with the series (3.30) for hybrid paths. The outcome is

$$\left(z_1 \xi(z_1) \frac{\partial}{\partial z_1} + z_2 \xi(z_2) \frac{\partial}{\partial z_2}\right) \Lambda_2(z_1, z_2, \mathbf{q}, u) = u \cdot \left[\frac{1}{u}(\xi(z) - 1) \circ \Delta^+(z\,;\, z_1, z_2)\right].$$

Apply Lemma 1.3.3 to expand the umbral composition, and then substitute w_i for z_i, for $i = 1, 2$. The result follows from (3.28) and Theorem 3.4.5. □

We remark that this corollary can also be deduced by extending the methods of §2.7.1 and §3.3.6 to two-face hybrid maps. Doing so yields

$$\Lambda_2(z_1, z_2, \mathbf{q}, u) = \log\left(\frac{z_1 - z_2}{z_1 e^{-uQ(z_1)} - z_2 e^{-uQ(z_2)}}\right) - u\left(Q(z_1) + Q(z_2)\right),$$

from which the corollary is immediately obtained by differentiation and an appeal to Theorem 3.4.5.

Of course, three-face smooth hybrid maps can be built by adding two labelled simple edges to a hybrid path, one extending from each of its ends. Figure 3.24 illustrates this process, with simple

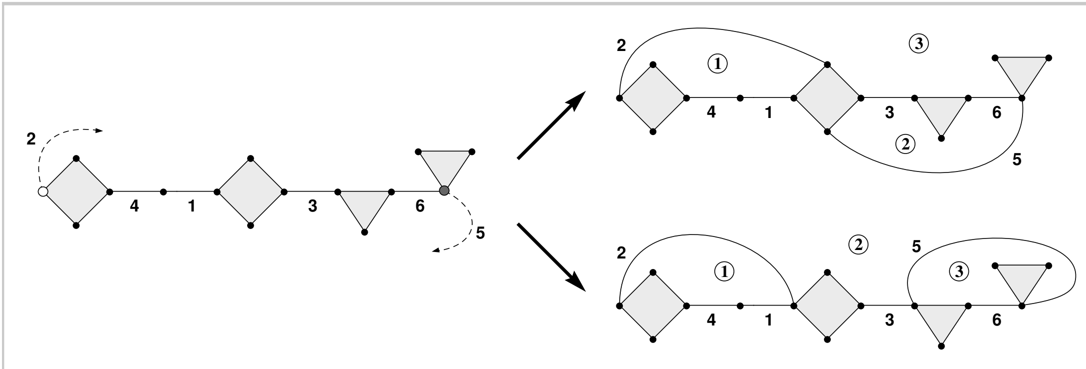

Figure 3.24: Construction of three-face smooth hybrid maps.

edges labelled 2 and 5 being attached to a path on 15 vertices in two different ways to obtain smooth hybrid maps of descent classes $(4, 3, 8)$ and $(3, 8, 4)$. The next theorem is surely expected; it shows that the process just described creates all possible smooth hybrid maps of a given descent class, and does so uniquely.

Theorem 3.4.8. *For any fixed $\theta = (\theta_1, \theta_2, \theta_3) \models n$, there is a polygon- and edge-preserving bijection between smooth, face-labelled, planar hybrid maps of descent class θ, and edge-labelled tuples $(\lambda, \mathscr{P}, \gamma)$, where λ, γ are distinct simple edges and $\mathscr{P}$ is a hybrid path containing n vertices.*

Proof. Modulo some obvious minor modifications, the proof of Theorem 2.7.14 remains valid in this context. The most subtle alteration is that the indexing $((v_0, e_0), \ldots, (v_k, e_k))^\circ$ of the single face of $\mathscr{P}$ should be chosen here so that v_0 is the white end of $\mathscr{P}$, and $(e_k, \lambda, e_0)^\circ$ is increasing. This is clearly always possible, and uniquely specifies the symbols v_i, e_i. We now let m be the unique index, with $0 < m \leq k$, such that v_m is the grey vertex of $\mathscr{P}$ and $(e_{m-1}, \gamma, e_m)^\circ$ is increasing. Note that these conditions guarantee simple edges labelled λ and γ can, indeed, be attached at corners $c_0 = (e_k, v_0, e_0)$ and $c_m = (e_{m-1}, v_m, e_m)$, respectively. □

We quickly derive the following corollary, which again can be found in [36].

Corollary 3.4.9. *Let $w_i = w(x_i, \mathbf{q}, u)$ and $\xi_i = \xi(w_i)$ for $i = 1, 2, 3$. Then*

$$\Theta_3(x_1, x_2, x_3, \mathbf{q}, u) = u^2 \sum_{i=1}^{3} \prod_{\substack{1 \leq j \leq 3 \\ j \neq i}} \frac{w_j}{w_i - w_j} (\xi_i - 1).$$

Proof. Copy the proof of Corollary 2.7.16, replacing the series $zu/(1-uz)$ for ordered paths used there with the series (3.30) for hybrid paths. This results in

$$\Lambda_3(z_1, z_2, z_3, \mathbf{q}, u) = u^2(\xi(z) - 1) \circ \Delta^+(z\,;\, z_1, z_2, z_3).$$

Now apply Lemma 1.3.3, replace z_i with w_i, for $i = 1, 2, 3$, and apply Theorem 3.4.5 to complete the proof. □

Finally, we mention that Theorem 2.7.17 also has a natural hybrid map analogue. We only state the result here, since the proof is virtually identical with the one given earlier. An algebraic proof can be found in [36]. Together with Theorem 3.4.5 and Corollary 3.4.9, this theorem shows that $\Theta_m(\mathbf{x}, \mathbf{q}, u)$ is a rational series in $w_1, \ldots, w_m$, with no explicit dependence on $x_1, \ldots, x_m$. As with Theorem 2.7.17, positive genus analogues are readily obtained.

Theorem 3.4.10. *Fix $m \geq 4$. For any subset $\lambda = \{\lambda_1, \ldots, \lambda_k\} \subseteq [m]$, where $\lambda_1 < \cdots < \lambda_k$, define $\mathbf{z}_\lambda = (z_{\lambda_1}, \ldots, z_{\lambda_k})$. For $1 \leq i \leq m$, set $\overline{\mathbf{z}}_i = \mathbf{z}_{[m]\setminus\{i\}}$. Also, for each i, let $\xi_i = \xi(z_i)$ and $\partial_i = z_i \xi_i \frac{\partial}{\partial z_i}$, and let $\mathcal{P}_i$ be the set of all pairs $\{\gamma, \lambda\}$ with $\gamma, \lambda \subset [m]$ such that $\gamma \cap \lambda = \{i\}$ and $\gamma \cup \lambda = [m]$. Then*

$$\frac{\partial}{\partial u}\Lambda_m(\mathbf{z}, \mathbf{q}, u) = \sum_{i=1}^{m} \sum_{\substack{\{\gamma,\lambda\}\in\mathcal{P}_i \\ |\gamma|,|\lambda|\geq 3}} \partial_i \Lambda_{|\gamma|}(\mathbf{z}_\gamma, u) \cdot \partial_i \Lambda_{|\lambda|}(\mathbf{z}_\lambda, u) + \sum_{\substack{1\leq i,j\leq m \\ i\neq j}} \frac{z_j \xi_i \partial_i \Lambda_{m-1}(\overline{\mathbf{z}}_j, u)}{(z_i - z_j)}.$$

□

3.4.6 A Final Bijection

For $\pi \in \mathfrak{S}_n$ and $g \geq 0$, let $\mathcal{F}_g(\pi)$ be the set of all transitive factorizations of π into transpositions. We conclude this chapter with a bijection on $\mathcal{F}_g(\pi)$ that can be described nicely in terms of duals of hybrid maps.

Let $\pi \in \mathfrak{S}_n$ be of cycle type β, and let $f \in \mathcal{F}_g(\pi)$. Then f is naturally associated with the β-factorization $f' = (\pi^{-1}, \tau_r, \ldots, \tau_1)$ of the identity in $\mathfrak{S}_n$, and hence also to the hybrid map $\mathscr{M}$ corresponding to f' through Theorem 3.4.2. Note that $\mathscr{M}$ is of descent class $[1^n]$, of polygon type β, and contains r simple edges. Now construct the dual map $\mathscr{M}^*$ in the usual way, by placing a vertex in each face and polygon of $\mathscr{M}$ and attaching two of these new vertices by an edge if their corresponding faces are incident with a common edge in $\mathscr{M}$ (simple or not). The faces and edges of $\mathscr{M}^*$ inherit labels from the vertices and edges of $\mathscr{M}$ in the process. See Example 3.4.11, below, for an illustration.

Let F be a face of $\mathscr{M}$ with boundary walk $((v_0, e_0), \ldots, (v_m, e_m))^\circ$, and let P be a k-gon of $\mathscr{M}$, where $k \geq 2$. Let v^F and v^P be the vertices of $\mathscr{M}^*$ corresponding to F and P, respectively. Then the rotator of v^F is $(e_m, \ldots, e_0)^\circ$. But $(e_0, \ldots, e_m)^\circ$ is increasing, since $\mathscr{M}$ is of descent class $[1^n]$, so the rotator of v^F is decreasing. Also note that no edge can appear twice in the boundary walk of F, and the same is obviously true of P. This holds for all faces and polygons of $\mathscr{M}$, so every one of its edges is incident with two distinct faces. Hence $\mathscr{M}^*$ is a loopless edge- and face-labelled map in which the rotator of every vertex corresponding to a (white) face of $\mathscr{M}$ is decreasing.

Transform $\mathscr{M}^*$ into a hybrid map, as follows. First replace edge label l with $r - l + 1$, for $1 \leq l \leq r$, so that decreasing rotators are made increasing. Delete v^P from $\mathscr{M}^*$ and form a polygon from its neighbours in the obvious way; repeat this process for each polygon P of $\mathscr{M}$ with at least two vertices. Let $\mathscr{M}^\#$ be the resulting face-labelled hybrid map of polygon type β. Note that $\mathscr{M}^\#$ contains n vertices and n (white) faces, and is therefore of descent class $[1^n]$.

Now shift the label of each face of $\mathscr{M}^\#$ to the unique vertex that is at a descent of that face. This transforms $\mathscr{M}^\#$ into a vertex-labelled hybrid map whose polygons are labelled identically with those of $\mathscr{M}$. Let $(\pi^{-1}, \tau'_r, \ldots, \tau'_1)$ be the β-factorization of the identity corresponding to $\mathscr{M}^\#$ through Theorem 3.4.2 and, finally, let $f^\#$ be the associated factorization $(\tau'_r, \ldots, \tau'_1)$ of π into transpositions.

Example 3.4.11. Let f be the factorization of $\pi = (1\,2\,3\,4)(5\,6\,7\,8\,9)(10)(11)$ given below:

$$\pi = (7\ 11)(5\,6)(9\ 10)(6\ 11)(7\ 11)(3\,4)(9\ 10)(1\,6)(7\,9)(2\,4)(7\,8)(1\,4)(4\,6). \tag{3.32}$$

The hybrid map $\mathscr{M}$ corresponding to the associated (5, 4, 1, 1)-factorization of the identity is shown in Panel A of Figure 3.25, on page 3.25. Panel B of the figure illustrates the construction of the dual $\mathscr{M}^*$. The hybrid map $\mathscr{M}^\#$ is shown in Panel C, and from it we see that

$$\pi = (4\,6)(1\,6)(7\,8)(1\,2)(8\,9)(1\,4)(8\ 10)(2\,3)(9\ 11)(1\,9)(8\ 10)(5\,9)(6\ 11) \tag{3.33}$$

is the factorization $f^\#$ of π associated with $\mathscr{M}^\#$. □

The transformation $f \mapsto f^\#$ defined above is clearly invertible and genus-preserving, and is therefore a bijection from $\mathcal{F}_g(\pi)$ to itself. We now introduce two statistics on $\mathcal{F}_g(\pi)$ that behave well with respect to this transformation.

Definition 3.4.12. *Let $f = (\tau_r, \ldots, \tau_1) \in \mathcal{F}_g(\pi)$, and let $j \in [n]$. We say that j* ***appears k times*** *in f if exactly k of the transposition factors τ_i are of the form $(j\ l)$ for some $l \in [n]$. We say that j is* ***moved k times*** *by f if there are exactly k factors τ_i such that $\tau_i \tau_{i-1} \cdots \tau_1(j) \neq \tau_{i-1} \cdots \tau_1(j)$, where τ_0 is understood to be the identity. We write $\alpha_k(f)$ for the set of symbols which appear k times in f, and $\mu_k(f)$ for the set of symbols it moves only once.*

For example, we have $\alpha_1(f) = \{2, 3, 5, 8\}$ and $\mu_2(f) = \{1, 2, 4, 10, 11\}$ for the factorization f given in (3.32). From (3.33), also note that $\alpha_2(f^\#) = \{1, 2, 3, 10, 11\}$ and $\mu_1(f^\#) = \{2, 3, 5, 8\}$. The fact that $\alpha_1(f) = \mu_1(f^\#)$ and $|\mu_2(f)| = |\alpha_2(f^\#)|$ is, of course, no coincidence.

Theorem 3.4.13. *We have $\alpha_k(f) = \mu_k(f^\#)$ for all $f \in \mathcal{F}_g(\pi)$ and $k \geq 1$.*

Proof. Let $f^\# = (\tau'_r, \ldots, \tau'_1)$, and let $\mathscr{M}$ and $\mathscr{M}^\#$ be the hybrid maps corresponding to f and $f^\#$, respectively. Fix $j \in [n]$, and let v and $v^\#$, respectively, denote the vertices of $\mathscr{M}$ and $\mathscr{M}^\#$ labelled j. Let $v^\#$ be at a descent of face F of $\mathscr{M}^\#$, so that v is the vertex of $\mathscr{M}$ dual to F. Let $((v_0, e_0), \ldots, (v_m, e_m))^\circ$ be the boundary walk of F, where e_0 is maximal amongst $e_0, \ldots, e_m$. We now prove that $j \in \alpha_k(f)$ is equivalent to $j \in \mu_k(f^\#)$.

First note that $j \in \alpha_k(f)$ if and only if v is incident with exactly k simple edges in $\mathscr{M}$. But the rotator of v is $(r - e_m + 1, \ldots, r - e_0 + 1)^\circ$, so this occurs precisely when either (1) $e_0 = \infty$ and $k = m$, or (2) $e_0 \neq \infty$ and $k = m + 1$. Since $e_0 \geq e_1$ is the single descent of F, we have $v_1 = v^\#$. Hence $\tau'_i \tau'_{i-1} \cdots \tau'_1(j) \neq \tau'_{i-1} \cdots \tau'_1(j)$ holds precisely when $i \in \{e_1, \ldots, e_m\}$ in the case $e_0 = \infty$, whereas the condition holds for $i \in \{e_0, e_1, \ldots, e_m\}$ otherwise. In the former case, $j \in \mu_m(f^\#)$, while $j \in \mu_{m+1}(f^\#)$ in the latter. Therefore $j \in \mu_k(f^\#)$ if and only if either (1) or (2), above, is satisfied. □

A similar proof shows that we also have $\alpha_k(f^\#) = \{\pi(j) : j \in \mu_k(f)\}$, and therefore $|\alpha_k(f^\#)| = |\mu_k(f)|$, for $f \in \mathcal{F}_g(\pi)$. This is reflective of the obvious near-duality between f and $f^\#$ induced by our constructions. In fact, we have the general identity $(f^\#)^\# = \pi f \pi^{-1}$, where the notation on the right indicates that each of the factors of f is to be conjugated by π.

Since $f \mapsto f^\#$ is a bijection on $\mathcal{F}_g(\pi)$, we get the following corollary of Theorem 3.4.13.

Corollary 3.4.14. $|\{f \in \mathcal{F}_g(\pi) : \alpha_k(f) = S\}| = |\{f \in \mathcal{F}_g(\pi) : \mu_k(f) = S\}|$ *for all $S \subseteq [n]$.* □

A factor of $f \in \mathcal{F}_g(\pi)$ is called a **consecutive pair** if it is of the form $(j\ \pi(j))$ for some j. Consider the case where $f = (\tau_{n-1}, \ldots, \tau_1) \in \mathcal{F}_0(\pi)$ is a minimal transitive factorization of the full cycle $\pi \in \mathscr{C}_{(n)}$. Since all factors of f are joins, by Lemma 2.2.5, the consecutive pairs of f are distinct. When $n > 2$, we claim that the mapping $j \mapsto (j\ \pi(j))$ is a bijection between $\mu_1(f)$ and the set of consecutive pairs of f.

Proof of claim: If j is moved only once by f, then the unique factor τ_i such that $\tau_i \tau_{i-1} \cdots \tau_1(j) \neq \tau_{i-1} \cdots \tau_1(j)$ is clearly the consecutive pair $\tau_i = (j\ \pi(j))$. If $\pi(j)$ were also moved only once by f, we would have $\pi^2(j) = j$, which is impossible since π is a full cycle on $n > 2$ symbols. Hence each $j \in \mu_1(f)$ corresponds with the unique consecutive pair $(j\ \pi(j))$ of f.

As noted above, all factors of f are joins. So if $\tau_i = (j\ \pi(j))$, then j does not appear in factors $\tau_1, \ldots, \tau_{i-1}$, and $\pi(j)$ does not appear in $\tau_{i+1}, \ldots, \tau_{n-1}$. (Otherwise, j and $\pi(j)$ would not appear consecutively, in that order, in the single cycle of $\pi = \tau_{n-1} \cdots \tau_1$.) Thus each consecutive pair $(j\ \pi(j))$ of f corresponds with $j \in \mu_1(f)$. □

We remark that this claim is also evident through graphical considerations. If $f \in \mathcal{F}_g(\pi)$, then the symbols moved once by f correspond with faces in the associated hybrid map $\mathscr{M}$ whose boundaries contain only one simple edge. Such an edge necessarily joins consecutive vertices of a polygon, and therefore corresponds to a consecutive pair in f. So the mapping $j \mapsto (j\ \pi(j))$ defined above is generally one-one, but not always onto. However, when f is a minimal factorization of a full cycle, then $\mathscr{M}$ is a planar map with only one polygon, and a simple edge connecting two consecutive vertices of this polygon clearly borders a face of degree 2 whose other boundary edge is not simple. Thus, in this case, consecutive pairs of f correspond with symbols moved only once. For example, the hybrid map corresponding to the factorization $(1\,2\,3\,4\,5\,6\,7\,8) = (2\,6)(4\,5)(1\,7)(2\,5)(7\,8)(2\,3)(1\,6)$ is illustrated below. Its consecutive pairs are $(2\,3)$, $(4\,5)$, and $(7\,8)$, and it moves symbols 2, 4, and 7 only once.

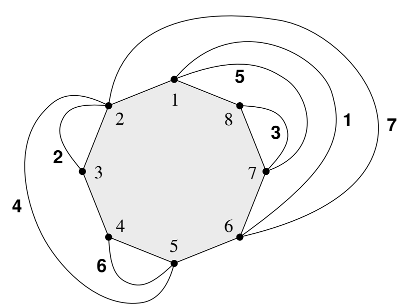

Proposition 3.4.15. *Let π be a full cycle of $\mathfrak{S}_n$, where $n > 2$. The number of minimal transitive factorizations of π containing the consecutive pairs $(j_1\ \pi(j_1)), \ldots, (j_k\ \pi(j_k))$ is equal to the number of such factorizations in which each of the symbols $j_1, \ldots, j_k$ appears only once.*

Proof. This follows directly from the claim and Corollary 3.4.14. □

Proposition 3.4.15 appears in [39], albeit in a somewhat different form than that given here. It is stated in terms of a bijection between vertex-labelled trees and minimal transitive factorizations of a full cycle, through which leaves of a tree are matched with consecutive pairs of a factorization. The bijection is essentially a composition of a specialization of the correspondence between factorizations and hybrid maps given here, and the bijection between trees and factorizations described in §2.4.7.

3.4.7 Additional Notes

A great deal of information on the double Hurwitz problem is contained in [36], including a description of some conjectural links with intersection theory. For more on the deep connections with geometry, also see [35]. Using localisation theory and certain results of [36], it is shown there that Faber's intersection number conjecture [20] can be reduced to a statement concerning genus 0 double Hurwitz numbers. The conjecture concerns intersection theory of the moduli space of genus g smooth curves, and it is hoped that a combinatorial viewpoint will lead to a more direct proof.

Corollaries 3.4.7 and 3.4.9, and Theorem 3.4.10, are proved in [36] through the methods discussed in §2.5.2. (In fact, we simplified our description of those methods by restricting to the single Hurwitz case.) It is also shown there how to extract the desired coefficients from $\Theta_2(x_1, x_2, \mathbf{q}, u)$ and $\Theta_3(x_1, x_2, x_3, \mathbf{q}, u)$ so as to obtain an explicit evaluation of the double Hurwitz number $H_0(\alpha, \beta)$ in the case that $\beta \vdash n$ is arbitrary and α has two or three parts. Finally, the representation theory of $\mathbb{C}\mathfrak{S}_n$ is used, as in §2.3.1, to give closed form expressions for $H_g((n), \beta)$, for all $g \geq 0$ and $\beta \vdash n$.

The bijection of [39] is phrased in terms of what are referred to there as *circle chord diagrams*. A circle chord diagram consists of a circle with n points on it, labelled $1, \dots, n$ in clockwise order around the perimeter, and $n-1$ chords labelled $2, \dots, n$ that connect these points to form a tree. Thus a circle chord diagram is simply an edge-labelled *non-crossing tree* on the circle; see [14, 55], and also §4.4, for more on non-crossing trees. Notice that if a circle chord diagram is turned "inside-out", so that its chords lie on the outside of the circle, then the resulting structure can be viewed as a planar hybrid map with a single n-gon (the circle) and $n-1$ simple edges (the chords).

A glimpse at the connection between factorizations and circle chord can be found in the work of Cohn and Lempel [11]. They use a chord diagram induced by a collection $\tau_1, \dots, \tau_n$ of transpositions to determine a matrix whose rank is directly related to the number of cycles in the product $\tau_1 \cdots \tau_n$. Beck [3] later extended these ideas.

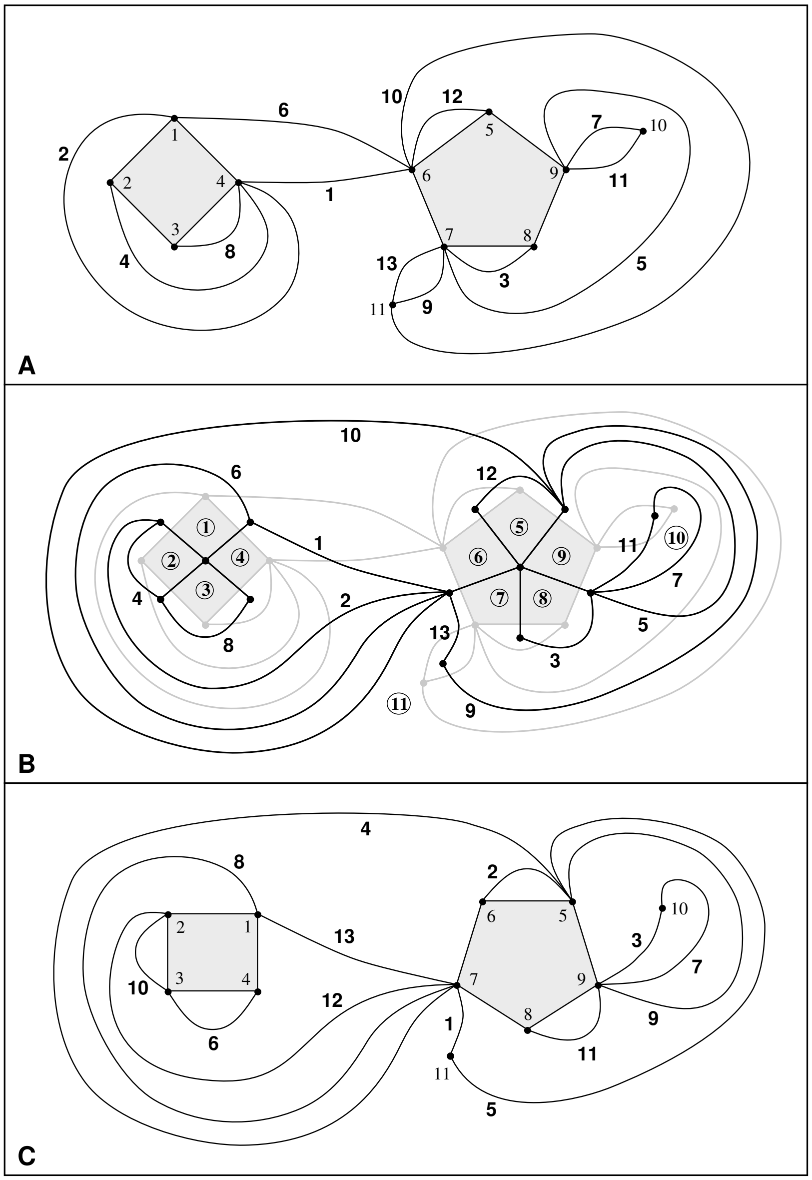

Figure 3.25: The construction of $\mathscr{M}^{\#}$.

Chapter 4

Inequivalent Factorizations

4.1 Introduction

The definition of factorizations as ordered tuples of permutations may be viewed as somewhat rigid, as it distinguishes between factorizations that differ only trivially in the order of their factors. The factorizations $(3\,4)(2\,5)(3\,5)(1\,2)$ and $(2\,5)(3\,4)(3\,5)(1\,2)$ of $(1\,2\,3\,4\,5)$, for instance, are considered distinct, despite the fact that they share the same sets of factors, and the second is obtained from the first simply by swapping the two leftmost (commuting) factors. We now relax the notion of sameness and consider the enumeration of factorizations up to an equivalence relation induced by commutation. Under this relation, the factorizations above will be deemed equivalent.

For the most part, we shall confine our discussions to cycle factorizations. Since two cycles commute only if they are disjoint or equal, we adopt the following definition of equivalence.

Definition 4.1.1. *Two cycle factorizations are* ***equivalent*** *if one can be obtained from the other by a sequence of interchanges of adjacent, disjoint factors. We write $f \sim g$ to indicate that factorizations f and g are equivalent.*

Clearly, $\sim$ is an equivalence relation on the set of cycle factorizations. For example, one equivalence class under this relation consists of the factorizations

$$\{(3\,4)(2\,5)(3\,5)(1\,2),\ (2\,5)(3\,4)(3\,5)(1\,2),\ (3\,4)(2\,5)(1\,2)(3\,5),\\ (2\,5)(3\,4)(1\,2)(3\,5),\ (2\,5)(1\,2)(3\,4)(3\,5)\}.$$

We shall occasionally use the symbol $\tilde{f}$ to denote the equivalence class containing the representative factorization f, but we typically speak of "inequivalent factorizations" rather than "equivalence classes under $\sim$".

The broad goal is to determine the number of inequivalent cycle factorizations of a fixed permutation, subject to a variety of constraints, such as minimality, transitivity, genus, *etc.* Relatively little work has been done on this problem in comparison with the vast amount of literature on ordered factorizations. After describing in moderate detail the few results that are known, we show how the methods of the previous two chapters can be modified so they are relevant in this new context.

4.2 Inequivalent Factorizations into Transpositions

As the title suggests, this section concerns only factorizations into transpositions. The term *factorization* is used exclusively with this meaning throughout.

4.2.1 Minimal Factorizations

The study of inequivalent factorizations began with the work of Eidswick [16] and Longyear [49]. Both authors determined, through different methods, the number of inequivalent minimal factorizations of a full cycle into transpositions. Longyear's analysis relied on a direct decomposition of such factorizations to a canonical form, while Eidswick employed an inclusion-exclusion argument. Here we briefly describe only the work of Longyear, since it seems the more natural of the two approaches. (In fact, our description more closely follows a refined treatment of the method, found in [32], than it does the original paper [49].)

Let $\big[\pi\big]$ denote a minimal transitive factorization of the permutation π. It is shown in [49] that, for any $f = \big[(1\,2\,\cdots\,n)\big]$, where $n \geq 2$, there are unique a, b with $1 < a \leq b \leq n$ such that

$$f \sim \big[(2\,3\,\cdots\,a)\big] * \big[(1\;(b+1)\;\cdots\;n)\big] * (1\,a) * \big[(a\;(a+1)\;\cdots\;b)\big], \tag{4.1}$$

where '$*$' indicates that the factorizations on the right are to be concatenated in the given order. Moreover, the three factorizations on the right-hand side are clearly unique up to equivalence, and all values of a, b with $1 < a \leq b \leq n$ are attainable. Thus (4.1) provides a canonical form for inequivalent minimal factorizations of a full cycle. For example, we have

$$\begin{aligned}(1\,2\,3\,4\,5\,6\,7\,8\,9) &= (2\,4)(6\,7)(1\,8)(5\,7)(1\,4)(2\,3)(4\,7)(8\,9)\\ &\sim (2\,4)(2\,3) * (1\,8)(8\,9) * (1\,4) * (6\,7)(5\,7)(4\,7)\\ &= \big[(2\,3\,4)\big] * \big[(1\,8\,9)\big] * (1\,4) * \big[(4\,5\,6\,7)\big],\end{aligned}$$

so that $a = 4$ and $b = 7$ in this case. In general, the parameters a and b corresponding to the factorization $f = \tau_{n-1} \cdots \tau_1$ are given by $a = \tau_k(1)$ and $b = \tau_1 \cdots \tau_k(1)$, where k is the minimal index such that $\tau_k(1) \neq 1$. Note that this identifies $(1\,a)$ as the rightmost factor of f that moves 1, as is clearly the case from (4.1).

Let $\tilde{h}_n$ be the number of inequivalent minimal factorizations of $(1\,2\,\cdots\,n)$, taking $\tilde{h}_1 = 1$ for the empty factorization, and consider the generating series

$$h(x) = \sum_{n\geq 1} \tilde{h}_n x^{n-1}.$$

The canonical form (4.1) leads to the cubic functional equation

$$h(x) = 1 + xh(x)^3, \tag{4.2}$$

which is solved routinely with Lagrange inversion to yield

$$\tilde{h}_n = \frac{1}{2n-1}\binom{3n-3}{n-1}. \tag{4.3}$$

Thus $\tilde{h}_n$ is a generalized Catalan number. This ubiquitous form implies the existence of myriad bijections between inequivalent factorizations of full cycles and other well-known combinatorial objects.

Indeed, since the publication of [16] and [49], other derivations of (4.3) have been found. Postnikov, for instance, has given a bijection between inequivalent factorizations and non-crossing trees on the circle. This can be found in [68, pg. 139], as can a bijection between non-crossing trees and plane cubic trees. Together, these results establish (4.2). In fact, we shall derive Postnikov's bijection later, in §4.4, as a special case of the graphical interpretation of inequivalent β-factorizations. Springer [65], and also Goulden and Jackson [28], have generalized these results. Their work is described in §4.3.1.

We have treated here only the case of inequivalent minimal factorizations of full cycles. However, the ostensibly more general problem of determining the number of inequivalent minimal factorizations of a permutation π composed of disjoint cycles $\pi_1, \ldots, \pi_m$ is no more difficult. If f is such a factorization, then we clearly have $f \sim [\pi_1] * \cdots * [\pi_m]$. By (4.3), the number of inequivalent minimal factorizations of a permutation of cycle type $\alpha = (\alpha_1, \ldots, \alpha_m)$ is therefore $\prod_{i=1}^{m} \frac{1}{2\alpha_i - 1}\binom{3\alpha_i - 3}{\alpha_i - 1}$. As we shall soon see, the structure of inequivalent minimal *transitive* factorizations is far more complex.

4.2.2 Factorizations of a Prescribed Length

Let $\pi \in \mathfrak{S}_n$ be any permutation of cycle type α. Following the notation of §2.3, let $\widetilde{F}_r(\alpha)$ denote the number of inequivalent factorizations (not necessarily transitive) of π into exactly r transpositions, and let $\widetilde{H}_g(\alpha)$ denote the number of inequivalent genus g (transitive) factorizations of π. Note that $\widetilde{F}_{n-1}((n)) = \widetilde{H}_0((n)) = \tilde{h}_n$ is given by (4.3). We introduce the generating series

$$\widetilde{\Upsilon}(z, \mathbf{p}, u) = \sum_{n,r \geq 1} \sum_{\alpha \vdash n} |\mathscr{C}_\alpha| \widetilde{F}_r(\alpha) \frac{z^n}{n!} u^r \, p_\alpha \tag{4.4}$$

and

$$\widetilde{\Phi}^{(g)}(z, \mathbf{p}, u) = \sum_{n \geq 1} \sum_{\alpha \vdash n} |\mathscr{C}_\alpha| \widetilde{H}_g(\alpha) \frac{z^n}{n!} u^{r_g(\alpha)} p_\alpha,$$

where $\mathbf{p} = (p_1, p_2, \ldots)$ and $p_\alpha = p_{\alpha_1} p_{\alpha_2} \cdots$ for $\alpha = (\alpha_1, \alpha_2, \ldots)$.

As in §2.3.1, an expression for $\widetilde{\Upsilon}$ can be given in terms of the irreducible characters of $\mathfrak{S}_n$, and through the standard logarithmic connection this leads to an expression for $\widetilde{\Phi}^{(g)}$. The derivation of these formulae is based on the commutation monoid of Cartier and Foata [9], which we now introduce.

Let $\mathcal{A}$ be a finite alphabet, and let $\mathcal{C}$ be a set of unordered pairs from $\mathcal{A}$. The elements of $\mathcal{C}$ are to be understood as *commutative pairs*, and two (finite) words on $\mathcal{A}$ are **$\mathcal{C}$-equivalent** if one can be transformed into the other by iteratively exchanging adjacent symbols a and a' for which $\{a, a'\} \in \mathcal{C}$. For example, if $\mathcal{A} = \{a, b, c, d\}$ and $\mathcal{C} = \{\{a, b\}, \{a, c\}, \{a, d\}, \{b, c\}\}$, then the words $a\,b\,c\,d$ and $c\,b\,d\,a$ are $\mathcal{C}$-equivalent, but neither is $\mathcal{C}$-equivalent to $d\,a\,b\,c$.

For $1 \leq k \leq |\mathcal{A}|$, define the formal sum

$$c_k = \sum_{A} a_1 a_2 \cdots a_k,$$

which extends over all subsets $A = \{a_1, \ldots, a_k\} \subseteq \mathcal{A}$ of size k such that every pair $\{a_i, a_j\} \subseteq A$

belongs to $\mathcal{C}$. More precisely, we mean that c_k should be taken to be any sum *of this form*, since the word $a_1 \cdots a_k$ corresponding to $A = \{a_1, \ldots, a_k\}$ depends on the order in which the elements of A have been indexed, and is therefore well-defined only up to $\mathcal{C}$-equivalence. Thus we always have $c_1 = \sum_{a \in \mathcal{A}} a$, whereas $c_2 = \sum_{\{a,b\} \in \mathcal{C}} ab$ depends on the ordering chosen for each pair of $\mathcal{C}$.

It is shown in [9] that

$$\frac{1}{1 - c_1 + c_2 - c_3 + \cdots} = \sum w \tag{4.5}$$

where the summation on the right extends over a complete list of $\mathcal{C}$-inequivalent words w on the alphabet $\mathcal{A}$. Of course, this expression is purely formal. It simply indicates that exactly one word from each $\mathcal{C}$-equivalence class appears in $\sum_j (c_1 - c_2 + c_3 - \cdots)^j$, with unit coefficient, upon formal expansion and simplification of the sum. Products here are interpreted as the usual (noncommutative) concatenations of words. The particular words that appear in the expansion depend on the choice of c_j for $j \geq 2$.

Following a suggestion of Goulden [24], we apply the result of Cartier-Foata in the following context. Fix $n \geq 1$, and let $\mathcal{A}$ be the set of all transpositions in $\mathfrak{S}_n$. Let $\mathcal{C}$ be the set of all pairs of *disjoint* elements of $\mathcal{A}$. Then $\widetilde{F}_r(\alpha)$ is equal to the number of $\mathcal{C}$-inequivalent words of length r on $\mathcal{A}$ that evaluate to a permutation of class $\mathscr{C}_\alpha$ when interpreted as a product in $\mathfrak{S}_n$. Since we clearly have $c_j = \mathsf{K}_{[1^{n-2j}\,2^j]}$, for $j \geq 1$, it follows from (4.5) that

$$\widetilde{F}_r(\alpha) = [\mathsf{K}_\alpha u^r] \left(1 + \sum_{j \geq 1} (-1)^j u^j \mathsf{K}_{[1^{n-2j}\,2^j]}\right)^{-1}. \tag{4.6}$$

Let $\beta_j = [1^{n-2j}\,2^j]$ for $j \geq 0$. Then, since $\mathsf{K}_{[1^n]}\mathsf{K}_\theta = \mathsf{K}_\theta$ for all $\theta \vdash n$, we have

$$\begin{aligned}
\widetilde{F}_r(\alpha) &= [\mathsf{K}_\alpha u^r]\, \mathsf{K}_{[1^n]} \left(1 + \sum_{j \geq 1} (-u)^j \mathsf{K}_{\beta_j}\right)^{-1} \\
&= [\mathsf{K}_\alpha u^r] \left(\sum_{\theta \vdash n} \mathsf{F}_\theta\right) \left(1 + \sum_{j \geq 1} (-u)^j\, |\mathscr{C}_{\beta_j}| \sum_{\theta \vdash n} \frac{\chi^\theta_{\beta_j}}{f^\theta} \mathsf{F}_\theta\right)^{-1} \\
&= [\mathsf{K}_\alpha u^r] \sum_{\theta \vdash n} \mathsf{F}_\theta \left(1 + \sum_{j \geq 1} (-u)^j\, |\mathscr{C}_{\beta_j}| \frac{\chi^\theta_{\beta_j}}{f^\theta}\right)^{-1} \\
&= [u^r] \sum_{\theta \vdash n} \frac{f^\theta}{n!} \chi^\theta_\alpha \left(\frac{1}{f^\theta} \sum_{j \geq 0} (-u)^j\, |\mathscr{C}_{\beta_j}| \chi^\theta_{\beta_j}\right)^{-1},
\end{aligned} \tag{4.7}$$

where the second and final equalities follow by (1.2), and the third by the idempotency of the F_θ. For $\theta \vdash n$, define $s^*_\theta(u) = \frac{1}{n!} \sum_{j \geq 0} (-u)^j\, |\mathscr{C}_{\beta_j}| \chi^\theta_{\beta_j}$. This notation is intended to be suggestive, as (2.8) implies that $s_\theta^*(u)$ is obtained from the Schur function s_θ (viewed as a function of the power sums) through the restrictions $p_1 = 1$, $p_2 = -u$, and $p_3 = p_4 = \cdots = 0$. Then, from (4.4), (4.7), and (2.8), we have

$$\widetilde{\Upsilon}(z, \mathbf{p}, u) = \sum_{n\geq 1}\sum_{\alpha\vdash n} \frac{z^n}{n!}|\mathscr{C}_\alpha| p_\alpha \sum_{\theta\vdash n}\left(\frac{f^\theta}{n!}\right)^2 \frac{\chi_\alpha^\theta}{s_\theta^*(u)} = \sum_{n\geq 1}\frac{z^n}{(n!)^2}\sum_{\theta\vdash n}(f^\theta)^2\frac{s_\theta}{s_\theta^*(u)}. \tag{4.8}$$

Moreover, if f is a factorization with components $f_1, \ldots, f_m$, then the factors of distinct components f_i and f_j clearly commute pairwise, so the class $\tilde{f}$ can be viewed as the unordered collection of classes $\{\tilde{f}_1, \ldots, \tilde{f}_m\}$. Hence we have the following connection between $\widetilde{\Upsilon}$ and the series $\widetilde{\Phi}^{(g)}$ counting transitive factorizations:

$$1 + \widetilde{\Upsilon}(z, \mathbf{p}, u) = \exp\left(\sum_{g\geq 0}\widetilde{\Phi}^{(g)}(z, \mathbf{p}, u)\right). \tag{4.9}$$

As was the case with the analogous expressions (2.9) and (2.11) for ordered factorizations, equations (4.8) and (4.9) do not shed much light on the nature of inequivalent factorizations. In particular, they do not simplify in any obvious way even for restricted cases where simple results are known. For example, it is unclear how one would derive (4.3) from these expressions.

4.2.3 Transitive Factorizations

Formula (4.3) for the number of inequivalent, minimal transitive factorizations of a full cycle into transpositions is strikingly simple. However, for partitions α with two or more parts, far less is known about the numbers $\widetilde{H}_0(\alpha)$. In particular, no analogue of the Hurwitz formula (2.12) is known for $\widetilde{H}_0(\alpha)$ when $\ell(\alpha) \geq 2$. Moreover, the existence of large factors in numerical data (obtained through computer search) implies that these numbers are not of a simple multiplicative form. Some of this data is reproduced in Table 4.1. We remark that the Lagrangian structure of the Hurwitz series (see §2.3.3) was discovered in hindsight, with the Hurwitz formula already conjectured from numerical evidence. We have not surmised any similar structure in the series $\widetilde{\Phi}^{(g)}$ of any genus.

Notice that the simple cut-and-join analysis that led to the differential equation (2.13) is not applicable to the study of inequivalent factorizations, as there is no canonical factor whose behaviour can be analyzed. Recently, however, Goulden, Jackson and Latour [32] have combined elementary cut-join analysis, reduction to a canonical form, and an inclusion-exclusion argument to determine a generating series for the numbers $\widetilde{H}_0((n_1, n_2))$, where $n_1, n_2 \geq 1$ are arbitrary. In fact, the series

α	$\widetilde{H}_0(\alpha)$	α	$\widetilde{H}_0(\alpha)$	α	$\widetilde{H}_0(\alpha)$
(1, 1)	1	(1, 1, 1)	24	(1, 1, 1, 1)	1578
(2, 1)	8	(2, 1, 1)	300	(2, 1, 1, 1)	24000
(2, 2)	74	(2, 2, 1)	3792	(2, 2, 1, 1)	357312
(3, 1)	54	(3, 1, 1)	2754	(3, 1, 1, 1)	258606
(3, 2)	540	(4, 1, 1)	22704	(1, 1, 1, 1, 1)	183120
(4, 1)	352	(3, 2, 1)	35028		
(3, 3)	4134	(2, 2, 2)	48288		
(4, 2)	3696	(3, 2, 2)	447984		
(5, 1)	2275	(3, 3, 1)	324756		
(4, 3)	29232	(4, 2, 1)	289920		
(5, 2)	24700	(5, 1, 1)	177450		
(6, 1)	14688				

Table 4.1: Numbers of inequivalent minimal transitive factorizations.

they obtain is a familiar symmetrization of $\widetilde{\Phi}^{(0)}$.

For $m \geq 1$, let $\widetilde{\Psi}_m(\mathbf{x}, u)$ be the image of $\widetilde{\Phi}^{(0)}(1, \mathbf{p}, u)$ under the symmetrization operator Π_m of (2.19). Then we have

$$\widetilde{\Psi}_m(\mathbf{x}, u) = \sum_{n \geq 1} \sum_{\substack{\alpha \models n \\ \ell(\alpha)=m}} \widetilde{H}_0(\alpha) \frac{x_1^{\alpha_1}}{\alpha_1} \cdots \frac{x_m^{\alpha_m}}{\alpha_m} u^{r_0(\alpha)}. \tag{4.10}$$

With $h(x)$ defined by the functional equation (4.2), the main result of [32] is the identity

$$\widetilde{\Psi}_2(x_1, x_2, 1) = \log\left(1 + x_1 x_2 h(x_1) h(x_2) \frac{h(x_1) - h(x_2)}{x_1 - x_2}\right). \tag{4.11}$$

The proof given there proceeds roughly as follows.

For $n_1, n_2 \geq 1$, let $\mathcal{S}^1_{n_1} = \{1^1, \ldots, n_1^1\}$ and $\mathcal{S}^2_{n_2} = \{1^2, \ldots, n_2^2\}$. We consider factorizations of permutations on the set $\mathcal{S}^1_{n_1} \cup \mathcal{S}^2_{n_2}$. In particular, let $\mathcal{F}(n_1, n_2)$ be the set of all minimal transitive factorizations of permutations π on this set that are composed of a n_1-cycle on $\mathcal{S}^1_{n_1}$ and an n_2-cycle on $\mathcal{S}^2_{n_2}$. Notice that every $f \in \mathcal{F}(n_1, n_2)$ is of length $r_0((n_1, n_2)) = n_1 + n_2$ and has a unique cut, by Lemma 2.2.5. In fact, the unique cut of f is the leftmost factor that is composed of one element from each of from $\mathcal{S}^1_{n_1}$ and $\mathcal{S}^2_{n_2}$.

A transposition factor is called a **possible cut** of $f \in \mathcal{F}(n_1, n_2)$ if it is the unique cut of some factorization equivalent to f. That is, τ is a possible cut of f if the factors of f can be commuted so that τ is a cut of the resulting factorization. For $k \geq 0$, let $\mathcal{D}_k(n_1, n_2)$ be the set of all factorizations

$f \in \mathcal{F}(n_1, n_2)$ in which k of the possible cuts have been distinguished. Thus a factorization with l possible cuts appears $\binom{l}{k}$ times in $\mathcal{D}_k(n_1, n_2)$. Let $\tilde{d}_k(n_1, n_2)$ denote the number of inequivalent factorizations in $\mathcal{D}_k(n_1, n_2)$. Then, since every element of $\mathcal{F}(n_1, n_2)$ has at least one possible cut, a straightforward inclusion-exclusion argument gives

$$\sum_{k \geq 0} (-1)^k \tilde{d}_k(n_1, n_2) = 0.$$

But clearly $\tilde{d}_0(n_1, n_2) = (n_1 - 1)!\,(n_2 - 1)!\,\widetilde{H}_0((n_1, n_2))$, so from (4.10) there follows

$$\widetilde{\Psi}_2(x_1, x_2, 1) = \sum_{k \geq 1} (-1)^{k-1} \widetilde{D}_k(x_1, x_2), \tag{4.12}$$

where we have put

$$\widetilde{D}_k(x_1, x_2) = \sum_{n_1, n_2 \geq 1} \tilde{d}_k(n_1, n_2) \frac{x_1^{n_1}}{n_1!} \frac{x_2^{n_2}}{n_2!}. \tag{4.13}$$

In [32], determination of the series $\widetilde{D}_k(x_1, x_2)$ hinges on a subtle combinatorial decomposition that the authors call a *switching algorithm*. It is first shown that any two given possible cuts $\tau = (a^1\, b^2)$ and $\rho = (c^1\, d^2)$ of a factorization $f \in \mathcal{F}(n_1, n_2)$ can be commuted so that they are adjacent in some factorization f' that is equivalent to f. The **switch** of f, denoted $\vartheta(f)$, is then obtained by replacing the consecutive pair of factors $\tau\rho = (a^1\, b^2)(c^1\, d^2)$ in f' with the pair $(a^1\, d^2)(c^1\, b^2)$. Of course, $\vartheta(f)$ depends on the cuts τ and ρ, but, given these, it is unique up to equivalence.

Note that $\vartheta(f)$ is never an element of $\mathcal{F}(n_1, n_2)$. In fact, the combinatorial significance of the switch is that it "splits" f into two smaller, disjoint factorizations, f_1 and f_2. That is, we have $\vartheta(f) \sim f_1 * f_2$, where each f_i is a minimal transitive factorization of a permutation composed of two cycles. The cycles of f_i are supported by subsets $C_i^1 \subset \mathcal{S}_{n_1}^1$ and $C_i^2 \subset \mathcal{S}_{n_2}^2$ such that $C_1^1 \cup C_2^1$ and $C_1^2 \cup C_2^2$ are set partitions of $\mathcal{S}_{n_1}^1$ and $\mathcal{S}_{n_2}^2$, respectively. Moreover, both f_1 and f_2 contain fewer possible cuts than f. The process can be iterated by choosing possible cuts of f_1 and f_2 and constructing the corresponding switches $\vartheta(f_1)$ and $\vartheta(f_2)$, *etc.*, and it naturally terminates when factorizations with only one possible cut are produced.

If applied only to distinguished possible cuts, the switching algorithm gives a decomposition of a factorization $f \in \mathcal{D}_k(n_1, n_2)$ into a collection of k elements of $\mathcal{D}_1(n_1, n_2)$. Moreover, it can be shown that the algorithm is reversible, up to equivalence. It transpires that $\widetilde{D}_k(x_1, x_2) = \frac{1}{k}\widetilde{D}_1(x_1, x_2)^k$, where the factor $1/k$ comes by taking ordering of the output into account. Thus (4.12) gives

$$\widetilde{\Psi}_2(x_1, x_2, 1) = \log\big(1 + \widetilde{D}_1(x_1, x_2)\big). \tag{4.14}$$

A canonical form akin to (4.1), but for special elements of $\mathcal{D}_1(n_1, n_2)$, is now introduced. Let $\mathcal{D}_\star(n_1, n_2) \subset \mathcal{D}_1(n_1, n_2)$ be the set of factorizations of $(1^1 \ldots n_1^1)(1^2 \ldots n_2^2)$ whose only distinguished possible cut is $(1^1\,1^2)$. A tedious argument shows that, for any $f \in \mathcal{D}_\star(n_1, n_2)$, there exist unique p_1, p_2 with $1 \le p_1 \le n_1$ and $1 \le p_2 \le n_2$ such that

$$f \sim \big[(1^2\,(p_2+1)^2 \cdots n_2^2)\big] * \big[(1^1\,(p_1+1)^1 \cdots n_1^1)\big] * (1^1\,1^2) * \big[(1^1 \cdots p_1^1\,1^2 \cdots p_2^2)\big], \tag{4.15}$$

where $[\pi]$ again denotes a minimal transitive factorization of π. Let $\tilde{d}_\star(n_1, n_2)$ be the number of inequivalent factorizations in $\mathcal{D}_\star(n_1, n_2)$. Then (4.15) implies

$$\begin{aligned}\tilde{d}_\star(n_1, n_2) &= \sum_{p_1=1}^{n_1}\sum_{p_2=1}^{n_2} \tilde{h}_{n_1-p_1+1}\tilde{h}_{n_2-p_2+1}\tilde{h}_{p_1+p_2} \\ &= [x_1^{n_1}x_2^{n_2}]\, x_1x_2h(x_1)h(x_2)\frac{h(x_1)-h(x_2)}{x_1-x_2},\end{aligned} \tag{4.16}$$

where $\tilde{h}_i$ and h are as defined in §4.2.2. Finally, observe that the symbols of any factorization in $\mathcal{D}_\star(n_1, n_2)$ can be relabelled in $n_1!n_2!$ ways to obtain distinct elements of $\mathcal{D}_1(n_1, n_2)$. Therefore $\tilde{d}_1(n_1, n_2) = n_1!n_2!\,\tilde{d}_\star(n_1, n_2)$. Equations (4.13), (4.14), and (4.16) now combine to give (4.11).

We have investigated the extension of this method to the enumeration of inequivalent minimal transitive factorizations of class (n_1, n_2, n_3), but our attempts have met with little success. Factorizations of this type have $r_0((n_1, n_2, n_3)) = n_1 + n_2 + n_3 + 1$ factors and exactly two cuts, which we call the *left* and *right* with obvious meaning. The existence of two cuts introduces complications that were not encountered in the derivation of $\widetilde{\Psi}_2$, above. For instance, note that the cuts, themselves, may commute. This makes the analysis of factorizations of $(1^1 \cdots n_1^1)(1^2 \cdots n_2^2)(1^3 \cdots n_3^3)$ quite intricate. It is unclear whether one should focus on a single cut at a time, trying to devise some sort of shelling scheme, or whether one should instead consider possible pairs of simultaneous left and right cuts.

Let $\mathcal{D}_\star(n_1, n_2, n_3)$ be the set of minimal transitive factorizations of $(1^1 \cdots n_1^1)(1^2 \cdots n_2^2)(1^3 \cdots n_3^3)$ that have a single distinguished pair of possible left and right cuts, namely $(1^1\,1^2)$ on the left and $(1^3\,a^1)$ on the right, where $1 \le a \le n_1$. Notice that this choice of cuts is completely general, since any minimal transitive factorization of $(1^1 \cdots n_1^1)(1^2 \cdots n_2^2)(1^3 \cdots n_3^3)$ can be relabelled to be of this form. By arguments similar to those used in [32] to obtain (4.15), we have found canonical forms for factorizations in $\mathcal{D}_\star(n_1, n_2, n_3)$. However, whereas (4.15) is universally valid for all factorizations in $\mathcal{D}_\star(n_1, n_2)$, the elements of $\mathcal{D}_\star(n_1, n_2, n_3)$ fall into five disjoint categories, each of which has its own canonical form. Because of their length, these forms are listed in Appendix A. Only under the strong restrictions $n_2 = n_3 = 1$ have we been able to obtain enumerative results based on such decompositions. (We do not report further here, as these results are superseded by Theorem 4.3.13, to follow.) In all other cases, some analogue of the inclusion-exclusion engine employed in [32] must be found before further progress can be made.

4.3 Inequivalent Cycle Factorizations

In this section we apply the methods developed in the previous chapters to the enumeration of inequivalent cycle factorizations. In particular, we describe how equivalence classes of cycle factorizations can be represented by certain decorated polymaps, and further show how cacti can be pruned from these polymaps so as to simplify their enumeration.

These graphical connections are exploited to count inequivalent minimal transitive cycle factorizations of permutations π with $\ell(\pi) = 1$ or $\ell(\pi) = 2$, thereby generalizing the results outlined in §4.2.1 and §4.2.3. The case $\ell(\pi) = 3$ is far more complex, but by restricting our attention to factorizations into transpositions we are able to derive a rough form of the series $\widetilde{\Psi}_3(\mathbf{x}, u)$ defined in (4.10). When $\ell(\pi) = 1$, our methods are closely related to work done by Springer [65], so we begin the section with a brief description of his work.

4.3.1 Factorizations of Full Cycles

Goulden and Jackson were first to obtain a result concerning inequivalent cycle factorizations into factors other than transpositions. In [28], they utilize a link between the connection coefficients of $\mathbb{C}\mathfrak{S}_n$ and those of a certain symmetric function algebra to obtain a simple formula for the number $\tilde{s}(n, k)$ of inequivalent minimal transitive k-cycle factorizations of a full cycle in $\mathfrak{S}_n$. In particular, it is shown there that

$$\tilde{s}(n, k) = \frac{1}{r}\binom{(2k-1)r}{r-1}$$

in the case that $n = 1 + r(k-1)$ for some positive integer r, and $\tilde{s}(n, k) = 0$ otherwise. This is done with the aid of the Cartier-Foata monoid, which is initially used to reduce the problem to a coefficient extraction involving class sums in $\mathbb{C}\mathfrak{S}_n$. (See (4.6) for an analogous expression. In fact, the class sum $\mathsf{K}_{[1^{n-2j}\,2^j]}$ appearing there need only be replaced with $\mathsf{K}_{[1^{n-kj}\,k^j]}$ to count factorizations into k-cycles.) The extraction is transformed into a computation involving symmetric functions, through which the series $s_k(x) = \sum_{n\geq 1} \tilde{s}(n,k)x^{n-1}$ is found to satisfy the functional equation

$$s_k(x) = 1 + x^{k-1} s_k(x)^{2k-1}. \tag{4.17}$$

The formula for $\tilde{s}_k(n)$ given above then follows by Lagrange inversion.

Observe that (4.17) gives $s_2(x) = 1 + x s_2(x)^3$ when $k = 2$. This identifies $s_2(x)$ with Longyear's series $h(x)$, defined in (4.2). However, whereas (4.2) was obtained through a straightforward combinatorial decomposition, the circuitous derivation of (4.17) leaves it devoid of combinatorial meaning. The quest for a combinatorial explanation of this functional equation is left in [28] as an open problem.

Springer [65] found such an explanation by generalizing Longyear's canonical form (4.1) to cycle factorizations. If f is a minimal transitive cycle factorization of $(1\,2\,\cdots\,n)$, he shows that there is a unique $k \geq 1$, and unique $a_1, \ldots, a_k, b_1, \ldots, b_k$ with $1 < a_1 \leq b_1 < a_2 \leq b_2 < \cdots < a_k \leq b_k \leq n$, such that

$$\begin{aligned} f \sim & \big[(2\,3\,\cdots\,a_1)\big] * \big[((b_1+1)\ (b_1+2)\ \cdots\ a_2)\big] * \cdots * \big[((b_k+1)\ (b_k+2)\ \cdots\ n\,1)\big] \\ & * (1\,a_1\,a_2\,\cdots\,a_k) * \big[(a_1\ (a_1+1)\ \cdots\ b_1)\big] * \cdots * \big[(a_k\ (a_k+1)\ \cdots\ b_k)\big], \end{aligned} \tag{4.18}$$

where $\big[\pi\big]$ represents a minimal transitive cycle factorization of π. This decomposition is then used to recursively define a rooted plane tree associated with the equivalence class of f. The non-leaf vertices of these trees are all of odd degree. In fact, if f has i_k k-cycle factors, for $k \geq 2$, then the tree corresponding to $\tilde{f}$ has exactly i_k vertices of degree $2k-1$. Inequivalent k-cycle factorizations therefore correspond to trees whose non-leaf vertices are all of degree $2k-1$. The series counting such trees satisfies the functional equation (4.17), thus explaining its combinatorial significance. In particular, when $k = 2$ we have the previously mentioned bijection between inequivalent minimal factorizations of a full cycle into transpositions and plane cubic trees.

More generally, trees with a specified number of internal vertices of given degree can be counted (see [18], for example) to obtain the following formula for the number of inequivalent minimal transitive cycle factorizations of $(1\,2\,\cdots\,n)$ of cycle index $(i_2, i_3, \ldots)$:

$$\frac{\left(\sum_{k\geq 2}(2k-1)i_k\right)!}{\left(1+\sum_{k\geq 2}(2k-2)i_k\right)!\,\prod_{k\geq 2} i_k!}. \tag{4.19}$$

We shall derive this formula later (Theorem 4.3.6) by different, but closely related, methods.

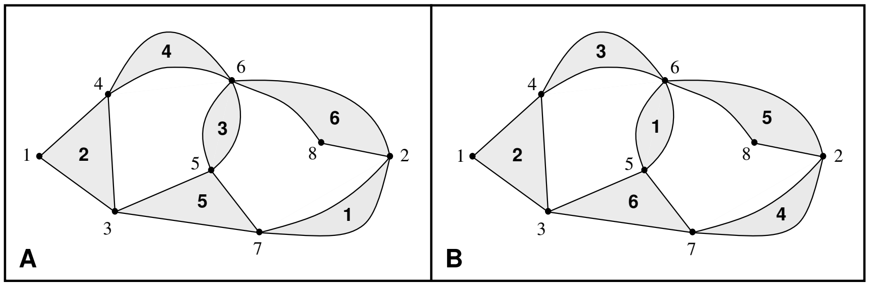

Figure 4.1: Polymaps of equivalent cycle factorizations.

4.3.2 Graphical Representation of Equivalence Classes

Commuting the factors of a cycle factorization clearly maintains the relative ordering of the factors that move any given symbol. Here a factor σ is understood to *move* the symbol i if $\sigma(i) \neq i$, or, equivalently, if i lies on the cycle σ. Thus cycle factorizations f and g of $\pi \in \mathfrak{S}_n$ are equivalent if and only if (1) they have precisely the same factors, and (2) for each $i \in [n]$, the factors that move i appear in the same order in f as they do in g.

From these comments, we see that commuting the factors of f is synonymous with relabelling the polygons of its (reduced) polymap $\mathscr{M}_f^\dagger$ in such a way that the relative order of the labels of the polygons incident with any given vertex is preserved. Thus, in particular, we have $f \sim g$ if and only if $\mathscr{M}_f^\dagger$ and $\mathscr{M}_g^\dagger$ have the same descent structure.

Consider, for example, the following equivalent cycle factorizations of $(1\,2\,3)(4\,5)(6\,7\,8)$:

$$(2\,8\,6)(3\,5\,7)(4\,6)(5\,6)(1\,4\,3)(2\,7) \sim (3\,5\,7)(2\,8\,6)(2\,7)(4\,6)(1\,4\,3)(5\,6). \tag{4.20}$$

The factors moving symbol 6 are $(5\,6)$, $(4\,6)$, and $(2\,8\,6)$, and they appear in exactly this right-to-left order in both factorizations. Let f and g, respectively, be the factorizations on the left and right of (4.20). The polymaps $\mathscr{M}_f^\dagger$ and $\mathscr{M}_g^\dagger$ are drawn in Figure 4.1. Note that the descent structure of these polymaps is identical. That is, vertex i is at a descent of a given face of $\mathscr{M}_f^\dagger$ (Figure 4.1A) if and only if it is at a descent of the corresponding face of $\mathscr{M}_g^\dagger$ (Figure 4.1B).

In this way, the equivalence classes of cycle factorizations are seen to have a natural graphical representation. The class $\tilde{f}$ containing the factorization f is represented by the polymap that results from stripping the polygon labels of $\mathscr{M}_f^\dagger$ and recording, instead, only the location of its descents. For instance, the decorated polymap corresponding to both factorizations in (4.20) is shown in Figure 4.2A.

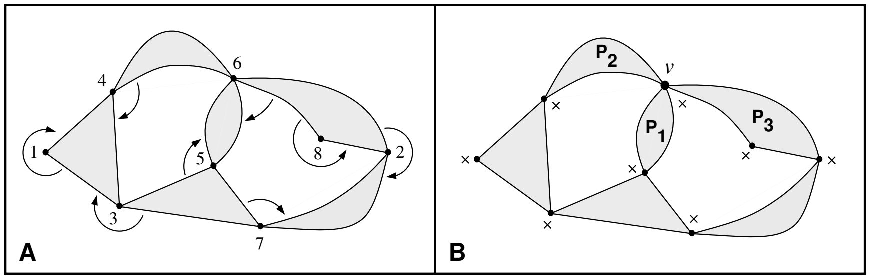

Figure 4.2: Descent-marked polymaps.

Definition 4.3.1. *A polymap is said to be* ***marked*** *if certain of its corners have been distinguished so that every vertex is at exactly one distinguished corner. The* ***rotator*** *of a vertex v in a marked polymap is the tuple $(P_1, \ldots, P_m)$ of polygons incident with v, listed in order as they are encountered along a clockwise tour about v beginning in the unique distinguished corner containing v. A* ***valid labelling*** *of a marked polymap is a polygon-labelling under which descent corners coincide with distinguished corners.*

For instance, the polymap in Figure 4.2B is marked, with its distinguished corners being indicated by small crosses. (This will be our standard convention for drawings of marked polymaps.) Note that the rotator of vertex v is (P_1, P_2, P_3). Figure 4.1 shows two valid labellings of this marked polymap.

Definition 4.3.2. *A loopless polymap $\mathscr{M}$ is said to be* ***descent-marked*** *if it is marked and admits a valid labelling. The distinguished corners of a descent-marked polymap $\mathscr{M}$ are called* ***descents****. The* ***descent set*** *of a face F of $\mathscr{M}$ is composed of all vertices at descents of F. If $\mathscr{M}$ has m_i faces containing exactly i descents, then its* ***descent partition*** *is $[1^{m_1}2^{m_2}\cdots]$.*

Observe that the descent structure of a descent-marked polymap is consistent with that induced by any of its valid labellings. From Theorem 3.2.8 we can immediately deduce that inequivalent cycle factorizations of genus g, class α, and cycle index $(i_2, i_3, \ldots)$ are in bijection with vertex-labelled, descent-marked polymaps of genus g with descent partition α and polygon index $(i_2, i_3, \ldots)$. However, as usual, we prefer to work with face-labelled maps.

Definition 4.3.3. *Let $\mathscr{M}$ be a descent-marked polymap with m labelled faces. The* ***descent class*** *of $\mathscr{M}$ is the composition $(\alpha_1, \ldots, \alpha_m)$, where α_i is the number of descents in face i, for $1 \le i \le m$. We say $\mathscr{M}$ is* ***properly labelled*** *if its vertices are also labelled in such a way that face s has descent set $\mathbb{D}_s(\alpha)$, for $1 \le s \le m$.*

The following result comes immediately from Theorem 3.3.2.

Theorem 4.3.4. *Inequivalent genus g cycle factorizations* $(\sigma_r, \ldots, \sigma_1)$ *satisfying* $\sigma_r \cdots \sigma_1 \in \mathfrak{S}(\alpha)$ *are in bijection with genus g, properly labelled, descent-marked polymaps that are of descent class* α *and contain r polygons. Moreover, under this bijection, a factorization with cycle index* $(i_2, i_3, \ldots)$ *corresponds to a polymap with polygon index* $(i_2, i_3, \ldots)$. □

For a vector $\mathbf{i} = (i_2, i_3, \ldots)$ of nonnegative integers and a composition α, let $\widetilde{M}_g(\alpha\,;\,\mathbf{i})$ denote the number of genus g, properly labelled, descent-marked polymaps of descent class α and polygon index $\mathbf{i}$. For $m \geq 1$ and $g \geq 0$, let

$$\widetilde{\Psi}_m^{(g)}(\mathbf{x}, \mathbf{p}, u) = \sum_{n \geq 1} \sum_{\mathbf{i} \geq \mathbf{0}} \sum_{\substack{\alpha \models n \\ \ell(\alpha) = m}} \widetilde{M}_g(\alpha\,;\,\mathbf{i}) \, \frac{\mathbf{x}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} \, \mathbf{p}^{\mathbf{i}} \, u^{r(\mathbf{i})},$$

where $\mathbf{x} = (x_1, \ldots, x_m)$, $\mathbf{p} = (p_2, p_3, \ldots)$, and $r(\mathbf{i}) = i_2 + i_3 + \cdots$. As usual, we write $\widetilde{\Psi}_m$ instead of $\widetilde{\Psi}_m^{(0)}$ for the genus 0 series. Notice that Theorem 4.3.4 implies $\widetilde{H}_g(\alpha\,;\,\mathbf{i}) = \widetilde{M}_g(\alpha\,;\,\mathbf{i}) \cdot \prod_i (\alpha_i - 1)!$. We therefore have the following corollary.

Corollary 4.3.5. *Let* $\alpha = (\alpha_1, \ldots, \alpha_m)$ *be a partition and let* π *be any permutation with cycle type* α. *Then, for* $g \geq 0$ *and any vector* $\mathbf{i} = (i_2, i_3, \ldots)$ *of nonnegative integers, the number of inequivalent genus g factorizations of* π *with cycle index* $\mathbf{i}$ *is given by* $\alpha_1 \cdots \alpha_m \cdot [\mathbf{x}^{\boldsymbol{\alpha}} \mathbf{p}^{\mathbf{i}} u^{i_2 + i_3 + \cdots}] \, \widetilde{\Psi}_m^{(g)}(\mathbf{x}, \mathbf{p}, u)$. □

4.3.3 Descent-Marked Cacti

Recall that a *cactus* is a planar polymap with only one face. Notice that any marked cactus is necessarily descent-marked. As a result, descent-marked cacti admit a particularly elegant recursive decomposition.

Let $\mathcal{C}$ be the set of vertex-rooted, descent-marked cacti with labelled non-root vertices. For the remainder of this section, we refer to elements of $\mathcal{C}$ simply as **cacti**. Let $w = w(x, \mathbf{p}, u)$ be the generating series for $\mathcal{C}$, with respect to labelled vertices (marked by x), polygon index (marked by $\mathbf{p}$), and total polygons (marked by u). Then we have

$$xw = x \frac{d}{dx} \widetilde{\Psi}_1(x, \mathbf{p}, u). \tag{4.21}$$

Let $C \in \mathcal{C}$, and suppose its root vertex has rotator $(P_1, \ldots, P_m)$. Detach polygons $P_2, \ldots, P_m$ from the root to form a cactus C' whose root has rotator $(P_2, \ldots, P_m)$, as shown in Figure 4.3.

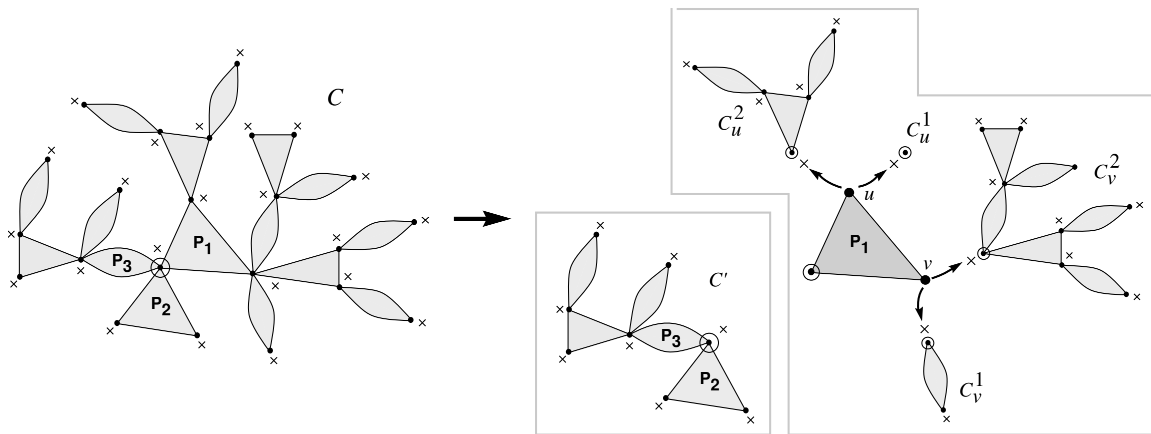

Figure 4.3: Decomposition of a descent-marked cactus.

(Vertex labels have been suppressed for clarity.) Notice that 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 $(P_v^1, \ldots, P_v^r, P_1, P_v^{r+1}, \ldots, P_v^s)$ 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 $(P_v^1, \ldots, P_v^r)$ and $(P_v^{r+1}, \ldots, P_v^s)$, respectively. See Figure 4.3 for an illustration. Thus C decomposes into a cactus C', together with k-gon P_1 and a $(k-1)$-tuple of triples (v, C_v^1, C_v^2).

It follows that $w = 1 + \sum_{k \geq 2} w \cdot up_k(xw^2)^{k-1}$, where the presence of 1 accounts for the case in which C consists of only one vertex. As in (3.14), we define $P \in \mathbb{Q}[\mathbf{p}][[z]]$ by

$$P(z) = \sum_{k \geq 2} p_k z^{k-1} \tag{4.22}$$

so that we can write

$$w = 1 + uwP(xw^2). \tag{4.23}$$

Notice that setting $u = p_k = 1$, and $p_i = 0$ for $i \neq k$, in this identity gives $w = 1 + x^{k-1}w^{2k-1}$. Thus w, under these restrictions, is identified with the series $s_k(x)$ of (4.17). That this should be the case is clear from (4.21).

Through Lagrange inversion, (4.21) and (4.23) yield the following result, which is equivalent to Springer's formula (4.19).

Theorem 4.3.6. *Let* $(i_2, i_3, \ldots)$ *be a sequence of nonnegative integers and set* $r = i_2 + i_3 + \cdots$. *Then the number of inequivalent minimal transitive cycle factorizations of* $(1\,2\,\cdots\,n)$ *with cycle index* $(i_2, i_3, \ldots)$ *is*

$$\frac{(2n+r-2)!}{(2n-1)!\,\prod_{k\geq 2} i_k!}$$

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

Proof. Set $v = w - 1$ so that (4.23) becomes

$$v = u(1+v)P(x(1+v)^2).$$

By (4.21) and Corollary 4.3.5, we wish to determine $[x^n u^r \mathbf{p}^{\mathbf{i}}]\, xw = [x^{n-1} u^r \mathbf{p}^{\mathbf{i}}]\,(1+v)$. This is accomplished through Lagrange inversion:

$$\begin{aligned}
[x^n u^r \mathbf{p}^{\mathbf{i}}]\,(1+v) &= [x^{n-1}\mathbf{p}^{\mathbf{i}}]\ \frac{1}{r}\,[\lambda^{r-1}]\,(1+\lambda)^r P(x(1+\lambda)^2)^r \\
&= \frac{1}{r}\,[\lambda^{r-1}]\,(1+\lambda)^r\ [x^{n-1}\mathbf{p}^{\mathbf{i}}]\left(\sum_{k\geq 2} p_k x^{k-1}(1+\lambda)^{2k-2}\right)^r \\
&= \frac{1}{r}\,[\lambda^{r-1}]\,(1+\lambda)^r (1+\lambda)^{2n-2}\binom{r}{i_2,\ i_3,\ i_4,\ \ldots} \\
&= \frac{(r-1)!}{\prod_{k\geq 2} i_k!}\,[\lambda^{r-1}]\,(1+\lambda)^{2n+r-2} \\
&= \frac{(r-1)!}{\prod_{k\geq 2} i_k!}\binom{2n+r-2}{r-1}.
\end{aligned}$$

□

Corollary 4.3.7. *Let* $n \geq 1$ *and* $k \geq 2$ *be such that* $n = 1 + r(k-1)$ *for some positive integer* r. *Then there are*

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

inequivalent minimal transitive k*-cycle factorizations of the full cycle* $(1\,2\,\cdots\,n)$. □

In hindsight, we remark that the decomposition of descent-marked cacti described here is essentially a high-level graphical interpretation of Springer's canonical form (4.18). First observe that an equivalence class $\tilde{f}$ of factorizations of the fixed full cycle $(1\,2\,\cdots\,n)$ corresponds with a cactus C whose root has label 1. The labels of all other vertices of C are determined from its descent structure. Let f be any member of the class $\tilde{f}$. Then the polygon P_1 in our decomposition of C corresponds with the rightmost factor of f that moves 1. If, as in (4.18), this factor is the cycle $(1\, a_1\, a_2\, \cdots\, a_k)$, then P_1 is a $(k+1)$-gon with vertices labelled $1, a_1, \ldots, a_k$ in clockwise order about its perimeter. Moreover, for $1 \leq i \leq k$, the factorizations $\big[(a_i\ (a_i+1)\ \cdots\ b_i)\big]$ and $\big[((b_{i-1}+1)\ (b_{i-1}+2)\ \cdots\ a_i)\big]$ appearing in (4.18) correspond with cacti $C^1_{a_i}$ and $C^2_{a_i}$ of our decomposition, respectively, while $\big[((b_k+1)\ \cdots\ n\ 1)\big]$ corresponds with C'. Here we have let $b_0 = 1$.

One benefit of our graphical approach to Theorem 4.3.6 is that it emphasizes "larger structure" by eliminating the need to invoke intricate "element-wise" decompositions such as (4.18).

4.3.4 Pruning Cacti

The absence of polygon labels makes pruning cacti from descent-marked polymaps a less involved process than that described by the cacti-pruning bijections of Chapter 3. The enumerative consequence of such pruning is given by Theorem 4.3.9, below, which is an analogue of Theorem 3.3.13 for descent-marked polymaps. As to be expected, it describes a relationship between $\widetilde{\Psi}^{(g)}_m$ and a certain series $\widetilde{\Gamma}^{(g)}_m$ counting smooth descent-marked polymaps. However, the series $\widetilde{\Gamma}^{(g)}_m$ introduced here is a refinement of its earlier counterparts, in that it accounts for an extra statistic, namely face degree. We define the **face degree sequence** of a polymap with m labelled faces to be the m-tuple $\mathbf{d} = (d_1, \ldots, d_m)$, where d_s is the degree of the face labelled s, for $1 \leq s \leq m$. Equivalently, d_s is the total number of corners in face s.

Definition 4.3.8. *Let $\widetilde{S}_g(\alpha\,;\ \mathbf{i}\,;\ \mathbf{d})$ denote the number of smooth, properly labelled, descent-marked polymaps of genus g and descent class α with polygon index $\mathbf{i}$ and face degree sequence $\mathbf{d}$. For $m \geq 1$, let*

$$\widetilde{\Gamma}^{(g)}_m(\mathbf{z}, \mathbf{t}, \mathbf{p}, u) = \sum_{n\geq 1} \sum_{\mathbf{i}\geq \mathbf{0}} \sum_{\substack{\alpha\models n\\ \ell(\alpha)=m}} \widetilde{S}_g(\alpha\,;\ \mathbf{i}\,;\ \mathbf{d})\, \frac{\mathbf{z}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!}\, \mathbf{p}^{\mathbf{i}}\, \mathbf{t}^{\mathbf{d}} u^{r(\mathbf{i})},$$

where $\mathbf{z} = (z_1, \ldots, z_m)$, $\mathbf{t} = (t_1, \ldots, t_m)$, $\mathbf{p} = (p_2, p_3, \ldots)$, and $r(\mathbf{i}) = i_2 + i_3 + \cdots$. We typically write $\widetilde{\Gamma}_m$ in place of $\widetilde{\Gamma}^{(g)}_m$ for the genus 0 series.

The construction of the core of a polymap (see Definition 3.3.11) must be modified slightly to account for distinguished corners in marked polymaps. Observe that removal of a leaf from a marked polymap $\mathscr{M}$ results in the amalgamation of two of its corners, c_1 and c_2, that contain the same vertex. The amalgamated corner is to be distinguished if and only if either of c_1 or c_2 isa descent. (Note that c_1 and c_2 cannot both be descents.) With this convention, the core of $\mathscr{M}$ is defined as before by the iterated removal of leaves.

Theorem 4.3.9. *Let* $g \geq 0$ *and* $m \geq 1$ *with* $(g, m) \neq (0, 1)$. *For* $1 \leq i \leq m$, *let* $w_i = w(x_i, \mathbf{p}, u)$,*where* w *is given by* (4.23)*. Then we have*

$$\widetilde{\Psi}_m^{(g)}(\mathbf{x}, \mathbf{p}, u) = \widetilde{\Gamma}_m^{(g)}(\mathbf{x} \circ \mathbf{w},\ \mathbf{w},\ \mathbf{p},\ u),$$

where $\mathbf{x} = (x_1, \ldots, x_m)$, $\mathbf{w} = (w_1, \ldots, w_m)$ *and* $\mathbf{x} \circ \mathbf{w} = (x_1 w_1, \ldots, x_m w_m)$.

Proof. Let $\mathscr{M}$ be a smooth, face-labelled, descent-marked polymap of genus g with m faces. Let c be a corner of the face of $\mathscr{M}$ labelled s, let v be the vertex at this corner, and let $(c_1, P_1, \ldots, c_k, P_k)^\circ$ be the alternating cyclic list of corners and polygons encountered along a clockwise tour about v. Assume this list is indexed so that $c = c_1$.

If c is a descent corner, then $(P_1, \ldots, P_k)$ is the rotator of v. Let C_1 and C_2 be cacti and let R_1 and R_2, respectively, be the rotators of their root vertices. Then, by identifying their roots with v, cacti C_1 and C_2 can be attached to $\mathscr{M}$ in corner c in a unique way so that the rotator of v becomes $(R_1, P_1, \cdots, P_k, R_2)$. The construction is illustrated below.

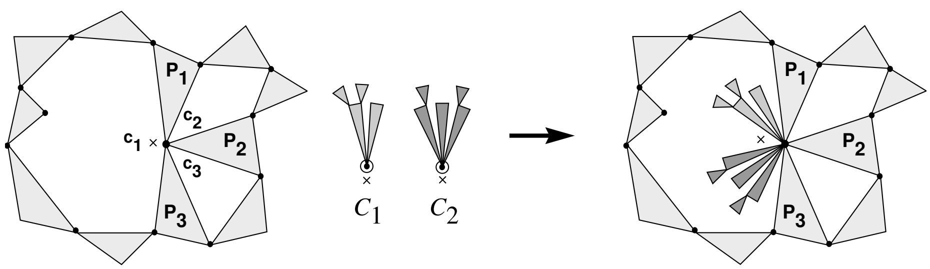

Observe that the marked polymap so formed is descent-marked, since a valid labelling is readily obtained from any valid labelling of $\mathscr{M}$.

Similarly, if c is not a descent corner, then $(P_i, \ldots, P_k, P_1, \ldots, P_{i-1})$ is the rotator of v, for some $i \neq 1$. Any cactus C whose root has rotator R can be attached to $\mathscr{M}$ in corner c so that the rotator of v becomes $(P_i, \ldots, P_k, R, P_1, \ldots, P_{i-1})$. This is shown below, in the case $i = 2$.

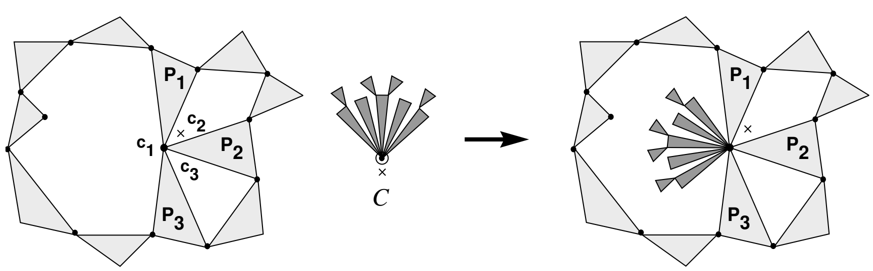

In either case, each non-root vertex of the attached cacti contributes a descent to the face of s of the newly formed map.

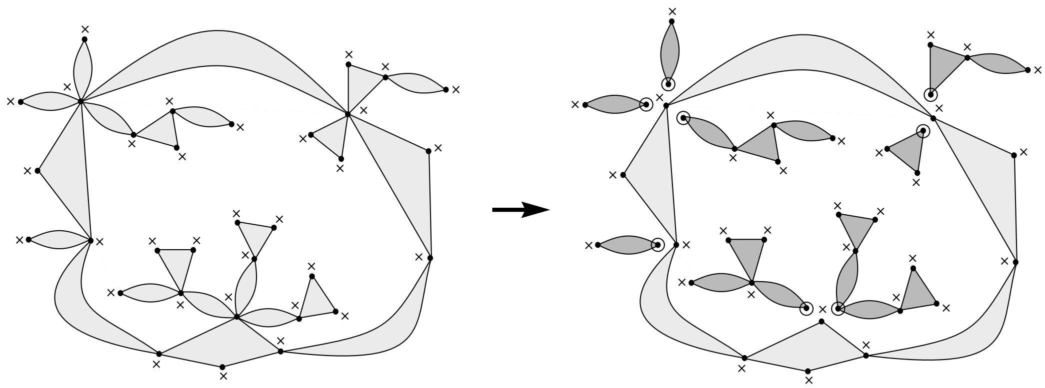

Figure 4.4: Pruning cacti from a descent-marked polymap.

Clearly any face-labelled, descent-marked polymap with core $\mathscr{M}$ can be created by carrying out this attachment process at all corners of $\mathscr{M}$. Moreover, if the vertices of $\mathscr{M}$ are labelled, then its corners are distinguishable and the process is reversible.

Observe that two cacti are to be attached at each of the α_s descent corners of face s of $\mathscr{M}$, while one cactus is attached at its remaining $d_s - \alpha_s$ corners. Thus $\alpha_s + d_s$ cacti are attached in face s altogether. Since the series $w(x, \mathbf{p}, u)$ counts cacti (with respect to the usual parameters), it follows that $\widetilde{\Psi}_m^{(g)}(\mathbf{x}, \mathbf{p}, u)$ is obtained from $\widetilde{\Gamma}_m^{(g)}(\mathbf{z}, \mathbf{t}, \mathbf{p}, u)$ through the substitutions $z_s \mapsto x_s w_s$ and $t_s \mapsto w_s$, for $1 \leq s \leq m$. □

The pruning of all cacti from a descent-marked polymap is illustrated in Figure 4.4. For the process to be reversible, vertex labels (or some other identifying mechanism) must be preserved. To avoid clutter, these are not shown in the diagram.

4.3.5 Factorizations of Class (n_1, n_2)

We now apply Theorem 4.3.9 to evaluate the series $\widetilde{\Psi}_2$, thereby generalizing (4.11) to factorizations with arbitrary cycle index. The main result comes as Corollary 4.3.12, which gives an expression for $\widetilde{\Gamma}_2$. Note the similarities between the derivation here and that of Ψ_2 given in §3.3.6. Throughout, P and w are defined as in (4.22) and (4.23).

Lemma 4.3.10. *Let $m, n \geq 1$. Up to rotational symmetry, there are*

$$\left[\frac{x^n}{n!}\frac{y^m}{m!}\right] \log\left(1 + \frac{xy}{1-(x+y)}\right)$$

distinct necklaces made of n labelled white beads and m (independently) labelled black beads.

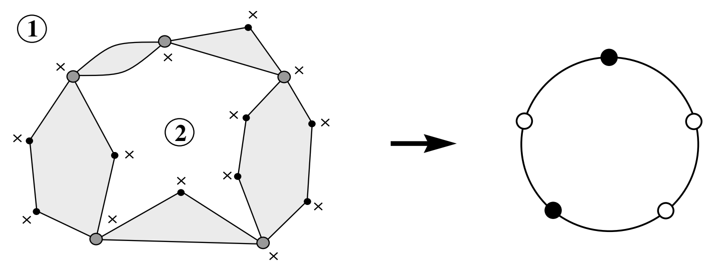

Figure 4.5: A smooth, two-face, descent-marked polymap and its associated necklace.

Proof. Any such necklace is formed by attaching the two ends of a string of labelled white and black beads. A string of this type decomposes into blocks of the form $w\,w\,\cdots\,w\,b\,b\,\cdots\,b$, where w and b represent white and black beads, respectively, and at least one bead of each colour is present. There are

$$m!n!\,[x^n y^m]\left(\frac{xy}{(1-x)(1-y)}\right)^k$$

strings consisting of k such blocks. However, by circular symmetry, exactly k of these strings form the same necklace. Thus the desired number of necklaces is

$$m!n!\,[x^n y^m]\sum_{k\geq 1}\frac{1}{k}\left(\frac{xy}{(1-x)(1-y)}\right)^k = \left[\frac{x^n}{n!}\frac{y^m}{m!}\right]\log\left(1-\frac{xy}{(1-x)(1-y)}\right)^{-1}.$$

The result follows upon rearrangement. □

Theorem 4.3.11.

$$\widetilde{\Gamma}_2(z_1,z_2,t_1,t_2,\mathbf{p},u) = \log\left(1+\frac{\delta^2 z_1 z_2}{1-\delta(z_1+z_2)}\right), \quad \textit{where}\ \ \delta = ut_1t_2\frac{P(t_1z_1)-P(t_2z_2)}{t_1z_1-t_2z_2}.$$

Proof. Let $\mathscr{M}$ be a smooth, properly-labelled, marked planar polymap with two faces. Then $\mathscr{M}$ is a closed chain of polygons, each incident with exactly two others. We say that a vertex incident with two polygons is *extremal*. Notice that $\mathscr{M}$ is descent-marked if and only if at least one extremal vertex is at a descent of each face. Suppose now that this is the case.

Let $L = (v_1,\ldots,v_r)^\circ$ be the cyclic sequence of extremal vertices encountered along the boundary walk of face 1 of $\mathscr{M}$. By regarding those v_i that are at descents of face 1 as white beads, and those at descents of face 2 as black beads, L corresponds with a necklace of the sort counted by Lemma 4.3.10. See Figure 4.5 for an illustration. Vertex (and bead) labels are not shown in the diagram, but extremal vertices are indicated in grey.

By the lemma, the generating series for such necklaces with respect to labelled white and black beads (marked by x and y, respectively) and total number of beads (marked by b) is

$$\log\left(1+\frac{bx\cdot by}{1-(bx+by)}\right). \tag{4.24}$$

Let M be the monomial in $\widetilde{\Gamma}_2(z_1,z_2,t_1,t_2,\mathbf{p},u)$ corresponding to $\mathscr{M}$. Each vertex v_i contributes the factor $z_1t_1t_2$ to M if it is at a descent of face 1, and contributes $z_2t_1t_2$ otherwise. A polygon of $\mathscr{M}$ with j_s-1 vertices incident only with face s, for $s=1,2$, further contributes the factor $up_{j_1+j_2}(t_1z_1)^{j_1-1}(t_2z_2)^{j_2-1}$ to M. Thus $\widetilde{\Gamma}_2(z_1,z_2,t_1,t_2,\mathbf{p},u)$ is obtained by performing the substitutions $x\mapsto z_1t_1t_2$, $y\mapsto z_2t_1t_2$, and

$$b\mapsto\sum_{j_1,j_2\geq 1} up_{j_1+j_2}(t_1z_1)^{j_1-1}(t_2z_2)^{j_2-1}=u\frac{P(t_1z_1)-P(t_2z_2)}{t_1z_1-t_2z_2}$$

in (4.24). The series resulting from these substitutions agrees with the claim of the theorem. □

Corollary 4.3.12. *With $w_i=w(x_i,\mathbf{p},u)$ for $i=1,2$, we have*

$$\widetilde{\Psi}_2(x_1,x_2,\mathbf{p},u)=\log\left(\frac{(x_1w_1-x_2w_2)^2}{(x_1-x_2)(x_1w_1^2-x_2w_2^2)}\right).$$

Proof. From Theorem 4.3.9 and Theorem 4.3.11 we get

$$\widetilde{\Psi}_2(x_1,x_2,\mathbf{p},u)=\widetilde{\Gamma}_2(z_1,z_2,t_1,t_2,\mathbf{p},u)\big|_{z_i=x_iw_i,\ t_i=w_i}=\log\left(1+\frac{\delta^2x_1x_2w_1w_2}{1-\delta(x_1w_1+x_2w_2)}\right),$$

where

$$\delta=uw_1w_2\frac{P(x_1w_1^2)-P(x_2w_2^2)}{x_1w_1^2-x_2w_2^2}.$$

But (4.23) gives $P(x_iw_i^2)=1-w_i^{-1}$, so we have

$$\delta=w_1w_2\frac{(1-w_1^{-1})-(1-w_2^{-1})}{x_1w_1^2-x_2w_2^2}=\frac{w_1-w_2}{x_1w_1^2-x_2w_2^2}.$$

It follows that

$$\widetilde{\Psi}_2(x_1,x_2,\mathbf{p},u)=\log\left(1+\frac{(w_1-w_2)^2x_1x_2}{(x_1w_1^2-x_2w_2^2)(x_1-x_2)}\right), \tag{4.25}$$

which can be rearranged to give the result. □

Under the restrictions $u = p_2 = 1$, and $p_i = 0$ for $i \geq 3$, the functional equation (4.23) becomes $w = 1 + xw^3$. That is, w restricts to the series h of (4.2). In this case we also have $xw^2 = 1 - w^{-1}$, so that (4.25) yields

$$\begin{aligned}\widetilde{\Psi}_2(x_1, x_2, \mathbf{p}, u)\big|_{u=p_2=1,\ 0=p_3=p_4=\cdots} &= \log\left(1 + \frac{(w_1 - w_2)^2 x_1 x_2}{((1-w_1^{-1}) - (1-w_2^{-1}))(x_1 - x_2)}\right) \\ &= \log\left(1 + x_1 x_2 w_1 w_2 \frac{w_1 - w_2}{x_1 - x_2}\right). \qquad (4.26)\end{aligned}$$

This is the series (4.11) discovered by Goulden-Jackson-Latour. Note that the current derivation eliminates their intricate inclusion-exclusion argument entirely (see §4.2.3), and suggests a more natural rôle for the logarithm in (4.11). (Namely, that $\log(1-x)^{-1}$ is the exponential generating series for cycles.) In hindsight, all the constructs of the GJL argument, including the switching algorithm, have natural graphical interpretations.

4.3.6 Factorizations of Class (n_1, n_2, n_3)

As explained in §3.3.4, the analysis of polymaps with at least three faces is complicated by the fact that a single polygon may be incident with three or more faces. This technicality has prevented us from finding a general expression for $\widetilde{\Psi}_3(\mathbf{x}, \mathbf{p}, u)$. However, the difficulty does not arise when considering polymaps that contain only 2-gons, and we have been able to derive a rough form of the restricted series $\widetilde{\Psi}_3(\mathbf{x}, \mathbf{p}, u)|_{p_3=p_4=\cdots=0}$ that counts inequivalent minimal transitive factorizations (into transpositions) of class (n_1, n_2, n_3).

Since the polymaps considered here consist solely of 2-gons, we refer to them simply as *maps* and draw them accordingly, by "flattening" all 2-gons into edges. For brevity, we write $\widetilde{\Psi}_3(\mathbf{x}, u)$ and $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ for the restrictions of $\widetilde{\Psi}_3(\mathbf{x}, \mathbf{p}, u)$ and $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, \mathbf{p}, u)$ under $p_2 = 1$, $p_3 = p_4 = \cdots = 0$.

The following notation will also be convenient:

- For a triple $\mathbf{x} = (x_1, x_2, x_3)$ and a permutation $\sigma \in \mathfrak{S}_3$, we let $\sigma(\mathbf{x}) = (x_{\sigma(1)}, x_{\sigma(2)}, x_{\sigma(3)})$.

- For $f \in \mathbb{Q}[\mathbf{t}, u][[\mathbf{z}]]$ we write $\langle f \rangle$ for the series $u(1 - uf)^{-1}$.

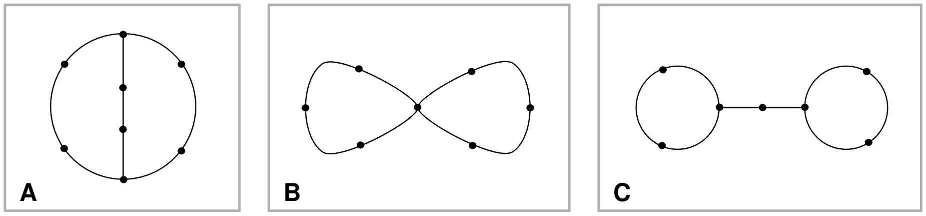

Figure 4.6: Classes of smooth, planar, three-face maps.

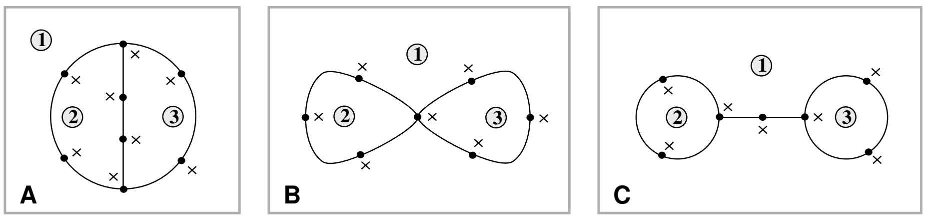

Figure 4.7: Smooth, face-labelled, marked planar maps.

Theorem 4.3.13. *Let* $\mathbf{z} = (z_1, z_2, z_3)$ *and* $\mathbf{t} = (t_1, t_2, t_3)$*. Define* $G(\mathbf{z}, \mathbf{t}, u) \in \mathbb{Q}[\mathbf{t}, u][[\mathbf{z}]]$ *by*

$$\begin{aligned} G(\mathbf{z}, \mathbf{t}, u) = \;& (V_1^{123})^2 \left(P_{13^+} P_{32^+} \langle V_1^{12}\rangle + P_{12^+} P_{23^+} \langle V_1^{13}\rangle + P_{12^+} P_{13^+} P_{23}\right) && (4.27)\\ &+ 2V_2^{123} V_3^{123} \left(P_{12} P_{23} P_{13} - \langle V_2^{12}\rangle P_{23} \langle V_3^{13}\rangle\right) && (4.28)\\ &+ 2V_1^{1123} P_{12^+} P_{13^+} + V_2^{1123} P_{21^+} P_{13^+} + V_3^{1123} P_{31^+} P_{12^+} && (4.29)\\ &+ P_{11} \left(V_2^{112} P_{21^+} + 2V_1^{112} P_{12^+}\right) \left(V_3^{113} P_{31^+} + 2V_1^{113} P_{13^+}\right), && (4.30) \end{aligned}$$

where, for $1 \le i, j, i_1, \ldots, i_m \le 3$*, we have set* $V_j^{i_1 \cdots i_m} = z_j t_{i_1} \cdots t_{i_m}$*,* $P_{ij} = \langle V_i^{ij} + V_j^{ij}\rangle$*, and* $P_{ij^+} = P_{ij} - \langle V_i^{ij}\rangle$*. Then*

$$\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u) = \sum_{\sigma \in \{\iota, (1\,2), (1\,3)\}} G(\sigma(\mathbf{z}), \sigma(\mathbf{t}), u).$$

Proof. Every smooth, planar, three-face map belongs to one of the three categories depicted in Figure 4.7. Observe that labelling the faces of any such map eliminates all non-trivial automorphisms. Thus $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ may be regarded as the (ordinary) generating series for smooth, face-labelled, descent-marked planar maps with three faces, with respect to descent class, face-degrees, and edges. Examples of such maps are shown in Figure 4.7. We shall hand-count these by category to obtain $\widetilde{\Gamma}_3$.

Category (A). Maps of this type contain exactly two vertices, u and v, of degree three. Observe that the cyclic sequences of face labels encountered on a clockwise tour about these vertices are always $(1,2,3)^{\circ}$ and $(1,3,2)^{\circ}$. We focus on three subcategories of maps defined by the following conditions: (A1) u and v are both at descents of face 1, (A2) u at a descent of face 2 while v is at a descent of face 3, and (A3) u is at a descent of face 3 while v is at a descent of face 2. These subcategories are illustrated below. Descents at low-degree vertices are not shown.

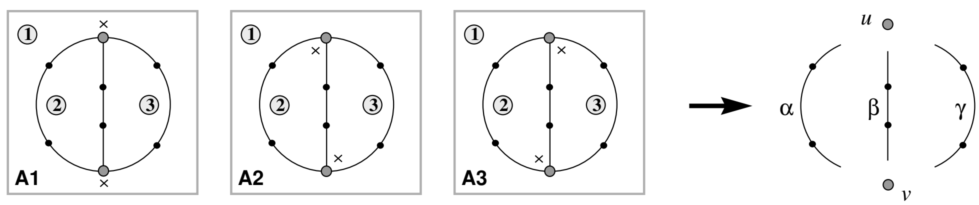

A map in any of these classes decomposes into vertices u and v together with the three paths α, β, and γ which connect them, as seen in the diagram. For each class, we count all possible descent-marked paths α, β, and γ such that every face of the map they generate contains at least one distinguished corner and at least one undistinguished corner. To do so, we exploit the observation that a smooth, three-face, planar marked map admits a valid labelling (*i.e.* is descent-marked) if and only if every face contains at least one distinguished corner and one undistinguished corner. We also make heavy use of the series $\langle V_i^{ij}\rangle$, P_{ij}, *etc.*, as defined in the statement of the theorem. Notice that these series have the following natural combinatorial interpretations in this context:

- $V_j^{i_1\cdots i_m}$ corresponds to a vertex that is at a descent of face j and is incident with m corners altogether, these belonging to faces $i_1,\ldots,i_m$.
- $\langle V_i^{ij}\rangle$ counts paths bordering faces i and j in which every vertex is at a descent of face i.
- P_{ij} counts paths bordering faces i and j.
- P_{ij^+} counts paths bordering faces i and j that have at least one vertex at a descent of face j.

Class (A1): If every vertex of α is at a descent of face 1, then some vertex of β must be at a descent of face 2 (otherwise face 2 would not contain any descents), and some vertex of γ must be at a descent of face 3 (otherwise every corner of face 1 would be a descent). The maps corresponding to this subcase are therefore counted by the series $(V_1^{123})^2\langle V_1^{12}\rangle P_{12^+}P_{23^+}$. One factor of V_1^{123} appears here for each of u and v, while α, β, and γ give rise to factors $\langle V_1^{12}\rangle$, P_{12^+}, and P_{23^+}, respectively.

If every vertex of γ is at a descent of face 3, then similar logic shows that the resulting counting series is $(V_1^{123})^2 P_{12^+} P_{23^+} \langle V_1^{13} \rangle$. Finally, if at least one vertex of α and at least one vertex of γ are at a descents of face 2, then the descents of β can be arbitrary; in fact, β can be of length 1, without any descents. The corresponding series is $(V_1^{123})^2 P_{12^+} P_{13^+} P_{23}$. The total contribution to $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ from class (A1) is therefore

$$(V_1^{123})^2 \left(P_{13^+} P_{32^+} \langle V_1^{12} \rangle + P_{12^+} P_{23^+} \langle V_1^{13} \rangle + P_{12^+} P_{13^+} P_{23} \right) \tag{4.31}$$

Classes (A2) & (A3): The analysis above could be applied here, with only minor modifications, to obtain an expression similar to (4.31). Alternatively, notice that the only way paths α, β, γ can result in a map that is not descent-marked is for α and β to contain only vertices at descents of faces 2 and 3, respectively. All other choices of α, β, γ are valid. Thus the contribution from each of classes (A1) and (A2) is

$$V_2^{123} V_3^{123} \left(P_{12} P_{23} P_{13} - \langle V_2^{12} \rangle P_{23} \langle V_3^{13} \rangle \right). \tag{4.32}$$

Summary: The total contribution to $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ from classes (A1), (A2), and (A3) is given by the sum of (4.31) and twice (4.32). Finally, observe that all other maps in category (A) are obtained uniquely by transposing either face labels 1 and 2, or 1 and 3, of the maps in these three classes. Thus (4.27) and (4.28) are accounted for.

Category (B). Maps in this category contain exactly one vertex, v of degree four. We consider four subcategories. In each of these, the cyclic sequence of face labels obtained from a tour about v is $(1, 2, 1, 3)^\circ$. The classes are characterized by which of the four corners containing v is a descent, as shown in the figure below.

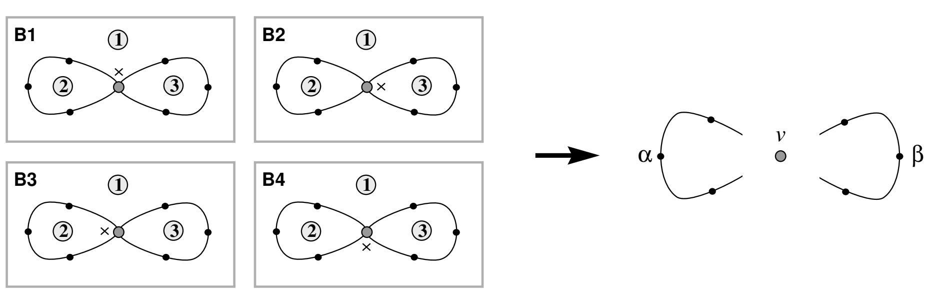

Every map in classes (B1) through (B4) decomposes into vertex v and paths α, β, as illustrated.

Applying the same analysis as in category A, we find that these classes together contribute

$$\overbrace{V_1^{1123} P_{12+} P_{13+}}^{\text{(B1)}} + \overbrace{V_3^{1123} P_{31+} P_{12+}}^{\text{(B2)}} + \overbrace{V_1^{1123} P_{12+} P_{13+}}^{\text{(B3)}} + \overbrace{V_1^{1123} P_{21+} P_{13+}}^{\text{(B4)}}$$

to $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$. In the case of (B1), for instance, path α (respectively, β) must contain at least one vertex at a descent of face 2 (respectively, face 3) to create a valid descent-marked map; otherwise, face 2 or 3 would have no descents. Again, all other maps in category (B) can be uniquely obtained from those in these four classes by transposing either face labels 1 and 2, or 1 and 3. This accounts for (4.29).

Category (C). As in category (A), these maps contain exactly two vertices, u and v, of degree three. We focus on nine subcategories. In each, the cyclic sequences of faces encountered about u and v are $(1, 2, 2)^\circ$ and $(1, 3, 3)^\circ$, respectively. The classes are distinguished by which of the three corners containing u and v are descents, as illustrated below.

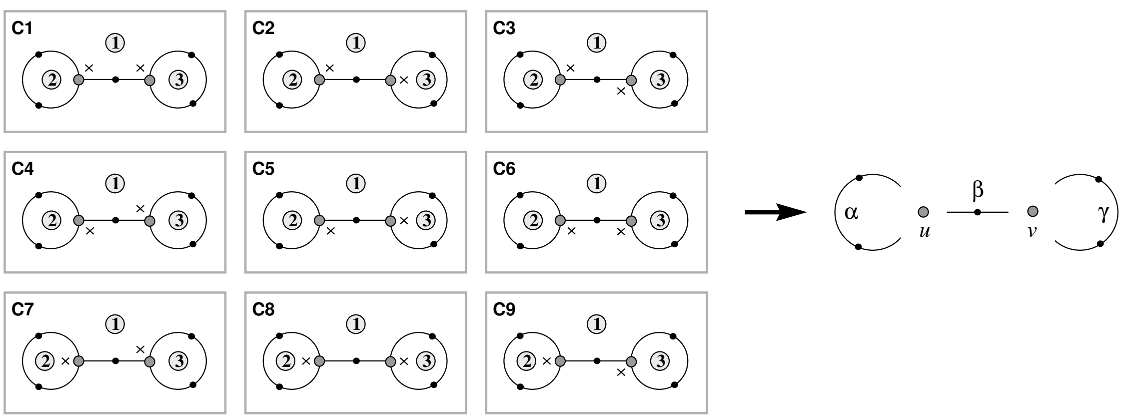

A map in any one of these classes decomposes as shown. In each case, β is an arbitrary path incident only with face 1. The corresponding counting series is P_{11}. In classes (C1) through (C6), note that α must contain at least one vertex at a descent of face 2, while in classes (C7) through (C9) it instead must contain a vertex at a descent of face 1. Note that the choice of α is always independent of that of γ. Of course, this argument is symmetric in α and γ, so we find that the total contribution to $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ of these nine classes of maps is

$$P_{11}\big(\overbrace{V_2^{112} P_{21+} + 2V_1^{112} P_{12+}}^{\alpha}\big)\big(\overbrace{V_3^{113} P_{31+} + 2V_1^{113} P_{13+}}^{\gamma}\big).$$

Once more, all other maps in category (C) can be obtained from those in these nine classes by transposing face labels 1 and 2, or 1 and 3. This accounts for (4.30), and completes the proof. □

We have been unable to combine the contributions to $\widetilde{\Gamma}_3(\mathbf{z}, \mathbf{t}, u)$ arising from the three distinct categories of smooth three-face planar maps to produce any significantly more homogeneous representation of the series than that which is given in Theorem 4.3.13. Through Theorem 4.3.9, we are therefore left with the following "rough form" of $\widetilde{\Psi}_3(\mathbf{x}, u)$.

Corollary 4.3.14. *For $i = 1, 2, 3$, let $w_i = w(x_i, u)$, where $w = w(x, u)$ is the unique series solution of $w = 1 + uxw^3$. Let $\mathbf{x} = (x_1, x_2, x_3)$ and define $F(\mathbf{x}, u) \in \mathbb{Q}[u][[\mathbf{x}]]$ by*

$$\begin{aligned} F(\mathbf{x}, u) = (X_1^{123})^2 P_{12^+}\big(P_{23^+}\langle X_1^{13}\rangle + \tfrac{1}{2}P_{13^+}P_{23}\big) + X_2^{123}X_3^{123}P_{23}\big(2P_{31^+}\langle X_2^{12}\rangle + P_{21^+}P_{31^+}\big) \\ + X_1^{1123}P_{12^+}P_{13^+} + X_2^{1123}P_{21^+}P_{13^+} \\ + \tfrac{1}{2}P_{11}\big(X_2^{112}P_{21^+} + 2X_1^{112}P_{12^+}\big)\big(X_3^{113}P_{31^+} + 2X_1^{113}P_{13^+}\big), \end{aligned}$$

where, for $1 \le i, j, i_1, \ldots, i_m \le 3$, we have $X_j^{i_1 \cdots i_m} = x_j w_j w_{i_1} \cdots w_{i_m}$, $P_{ij} = \langle X_i^{ij} + X_j^{ij}\rangle$, and $P_{ij^+} = P_{ij} - \langle X_i^{ij}\rangle$. Then

$$\widetilde{\Psi}_3(\mathbf{x}, u) = \sum_{\sigma \in \mathfrak{S}_3} F(\sigma(\mathbf{x}), u).$$

Proof. This follows directly from Theorem 4.3.13 by symmetrizing and applying Theorem 4.3.9. Note that the functional equation for w comes from restricting (4.23) with $p_3 = p_4 = \cdots = 0$. □

We have not been able to simplify this expression for $\widetilde{\Psi}_3(\mathbf{x}, u)$ in any meaningful way. The functional equation $w = 1 + uxw^3$ allows for the elimination of high powers of w, but it is unclear what general form should be targeted when using this relation for simplification. Though it represents truly minimal evidence, one might conjecture from (4.26) that $\widetilde{\Psi}_3(\mathbf{x}, u)$ can be expressed cleanly in terms of alternants involving the series w_i. Identities such as

$$x_1 x_2 w_1 w_2 \frac{w_1 - w_2}{x_1 - x_2} = X_1^{12} X_2^{12} P_{12}$$

lend support to this claim.

For now we regard Corollary 4.3.14 as a piece of raw data, and hope that it can be manipulated to uncover further structure of inequivalent factorizations. It would be more tedious than difficult to extend the methods used here to obtain a similar conglomerate expression for $\widetilde{\Psi}_m(\mathbf{x}, u)$, but there

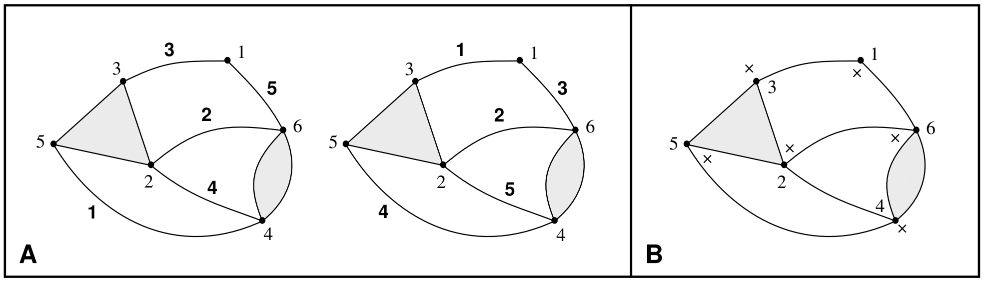

Figure 4.8: (A) Hybrid maps of equivalent β-factorizations, and (B) the vertex-labelled, descent-marked hybrid map corresponding to their common equivalence class.

does not appear to be good reason to do so until more is known about the "true" nature of the results that have already been obtained.

4.4 Inequivalent β-Factorizations

We say that two β-factorizations $(\sigma, \tau_r, \ldots, \tau_1)$ and $(\sigma', \tau_r', \ldots, \tau_1')$ are **equivalent** if the factorizations $(\tau_r, \ldots, \tau_1)$ and $(\tau_r', \ldots, \tau_1')$ are equivalent according to Definition 4.1.1. For instance we have the following equivalence amongst (3, 2, 1)-factorizations of $(1\,2)(3\,4)(5)(6)$:

$$(1)(2\,5\,3)(4\,6)\cdot(1\,6)(2\,4)(3\,1)(2\,6)(4\,5) \sim (1)(2\,5\,3)(4\,6)\cdot(2\,4)(4\,5)(1\,6)(2\,6)(3\,1). \tag{4.33}$$

The methods introduced in the previous section to count inequivalent cycle factorizations are readily altered to make them applicable to the enumeration of inequivalent β-factorizations. The nicest result that we have obtained in this way concerns the number of inequivalent minimal transitive β-factorizations of a fixed full cycle. We conclude with a brief derivation of this result.

Observe that combining the material from §3.4.2 and §4.3.2 shows equivalence classes of β-factorizations to be in correspondence with vertex-labelled *descent-marked hybrid maps* of polygon type β. The formal definition of this class of maps is the obvious hybrid map analogue of Definition 4.3.2, and will not be given here. Instead, we refer to Figure 4.8, where the hybrid maps corresponding to the β-factorizations of (4.33) are shown, along with the vertex-labelled, descent-marked hybrid map corresponding to their common equivalence class. Rotators, descents sets, *etc.*, are also defined as before. Notice that a rotator in a descent-marked hybrid map consists of either simple edges only, or simple edges and a single polygon. In the latter case, we point out that the polygon must come at the tail of the rotator. This follows from the fact that polygons of a hybrid

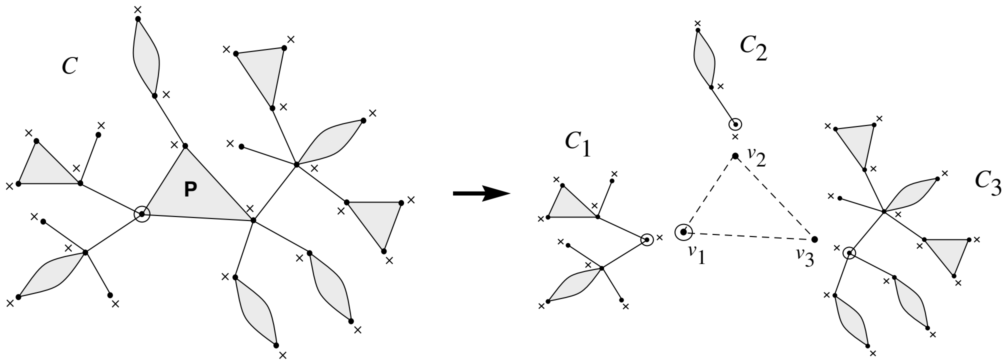

Figure 4.9: Decomposition of a rooted, descent-marked, hybrid cactus.

map are maximally labelled.

Theorem 4.4.1. *Let $\alpha, \beta \vdash n$ and $g \geq 0$. There is a bijection between inequivalent β-factorizations of genus g and class α, and vertex-labelled, descent-marked hybrid maps of genus g with descent partition α and polygon type β.* □

Theorem 4.4.1 establishes that the set of inequivalent minimal transitive β-factorizations of class (n) is in bijection with vertex-labelled descent-marked hybrid cacti on n vertices with polygon type β. Note that any such factorization has $\ell(\beta) - 1$ transposition factors, so the corresponding cacti have this number of simple edges.

Let $\mathcal{C}$ be the set of vertex-labelled, descent-marked, rooted hybrid cacti, and let $\vartheta = \vartheta(x, \mathbf{q}, u)$ be the generating series for $\mathcal{C}$, where x marks labelled vertices, u marks edges, and $\mathbf{q} = (q_1, q_2, \ldots)$ records polygon type. For brevity, we shall henceforth refer to elements of $\mathcal{C}$ simply as **cacti**. By the comments above, the number of inequivalent minimal transitive β-factorizations of $(1\,2\,\cdots\,n)$ is given by

$$[x^n q_\beta u^{\ell(\beta)-1}]\,\vartheta(x, \mathbf{q}, u). \tag{4.34}$$

Define a **simple cactus** to be a rooted, descent-marked, hybrid cactus, with labelled *non-root* vertices, whose root vertex is incident only with simple edges. Let $w = w(x, \mathbf{q}, u)$ be the generating series for simple cacti, with respect to the same statistics as above. We now develop functional equations relating ϑ and w by considering decompositions of cacti.

Let $C \in \mathcal{C}$ be a cactus whose root is incident with polygon P. If P is a k-gon, then observe that C decomposes into P and a k-tuple $((v_1, C_1), \ldots, (v_k, C_k))$, where each v_i is a vertex of P and C_i

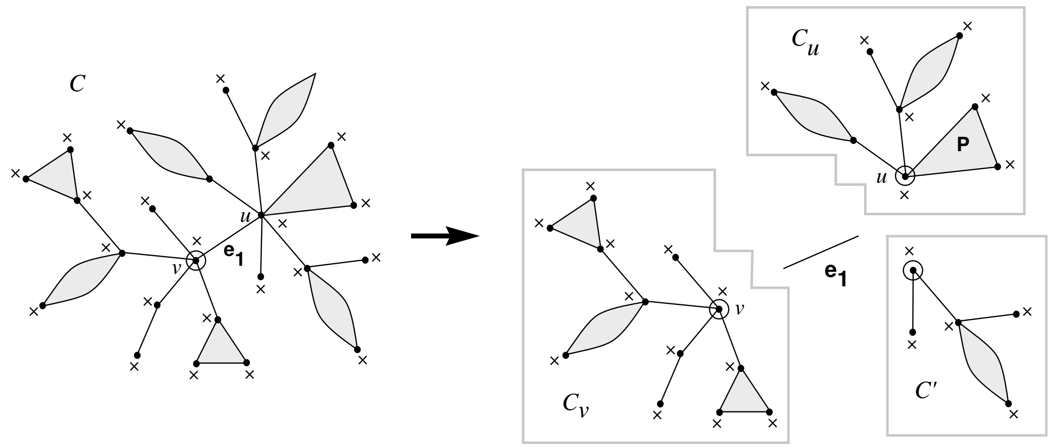

Figure 4.10: Decomposition of a simple cactus.

is a simple cactus. See Figure 4.9 for an illustration. It follows that

$$\vartheta = \sum_{k\geq 1} q_k (xw)^k = Q(xw), \tag{4.35}$$

where (as in §3.4) we have defined $Q \in \mathbb{Q}[\mathbf{q}][[z]]$ by

$$Q(z) = \sum_{k\geq 1} q_k z^k.$$

Now consider a simple cactus C whose root vertex v has rotator $(e_1, \ldots, e_m)$. Let u be one endpoint of the simple edge e_1. Then the rotator of u is either $(a_1, \ldots, a_j, e_1, b_1, \ldots, b_k, P)$ or $(a_1, \ldots, a_j, e_1, b_1, \ldots, b_k)$, where the a_i and b_i are simple edges, P is a polygon, and we allow the degenerate conditions $j = 0$ and $k = 0$ (interpreted in the obvious way). In the former case, notice that C decomposes into e_1, two simple cacti C_v, C', and a cactus C_u, where the roots of C_v, C', and C_u whose roots have rotators $(e_2, \ldots, e_m)$, $(a_1, \ldots, a_j)$, and $(b_1, \ldots, b_k, P)$. The same holds true in the latter case, except that C_u instead has rotator $(b_1, \ldots, b_k)$. See Figure 4.10.

From this decomposition there follows $w = 1 + uw^2\vartheta$, where the addition of 1 accounts for the case in which C_v consists of a single root vertex. With (4.35), we get

$$w = 1 + uw^2 Q(xw). \tag{4.36}$$

Lagrange inversion may now be applied to evaluate (4.34). This yields the following tidy result.

Theorem 4.4.2. *Let $\beta \vdash n$ and set $m = \ell(\beta)$. If $m \geq 2$, then there are*

$$\frac{n(m-2)!}{|\operatorname{Aut}(\beta)|}\binom{n+2m-3}{m-2}$$

inequivalent minimal transitive β-factorizations of the full cycle $(1\,2\,\cdots\,n)$. If $m = 1$, then the only such factorization is the trivial factorization $(1\,2\,\cdots\,n) = (1\,2\,\cdots\,n)$.

Proof. Setting $v = w - 1$ in (4.36) gives

$$v = u(1+v)^2 Q(x(1+v)),$$

so (4.35) implies $[x^n u^{m-1} q_\beta]\,\vartheta = [x^n u^{m-1} q_\beta]\,Q(x(1+v))$. Lagrange's theorem now yields

$$\begin{aligned}
[x^n u^{m-1} q_\beta]\,Q(x(1+v)) &= [x^n q_\beta]\,\frac{1}{m-1}\,[\lambda^{m-2}]\,Q'(x(1+\lambda))x\cdot(1+\lambda)^{2m-2}Q(x(1+\lambda))^{m-1} \\
&= \frac{1}{m(m-1)}\,[\lambda^{m-2}x^{n-1}]\,(1+\lambda)^{2m-3}\,[q_\beta]\,\frac{d}{dx}Q(x(1+\lambda))^m \\
&= \frac{1}{m(m-1)}\,[\lambda^{m-2}x^{n-1}]\,(1+\lambda)^{2m-3}\,\frac{d}{dx}\,[q_\beta]\left(\sum_{k\geq 1} q_k x^k (1+\lambda)^k\right)^m
\end{aligned}$$

$$= \frac{1}{m(m-1)}\,[\lambda^{m-2}x^{n-1}]\,(1+\lambda)^{n+2m-3}\,\frac{m!}{|\operatorname{Aut}(\beta)|}\,\frac{d}{dx}\,x^n,$$

and the result follows. □

A minimal transitive factorization $f = (\tau_{n-1}, \cdots, \tau_1)$ of $\pi = (1\,2\,\cdots\,n)$ into transpositions is equivalent to a minimal transitive $[1^n]$-factorization of π. Setting $\beta = [1^n]$ in Theorem 4.4.2 therefore yields Longyear's formula (4.3). Alternatively, f is associated with the (n)-factorization $(\pi, \tau_1, \ldots, \tau_{n-1})$ of the identity. Thus the equivalence class containing f corresponds with a descent-marked hybrid map with descent partition $[1^n]$, consisting of a single n-gon and $n-1$ simple edges, where the vertices of the n-gon are labelled 1 to n in clockwise order around its perimeter. Turning this polygon "inside-out" and removing the descent markings (which are superfluous) results in a non-crossing tree on the circle, as illustrated below.

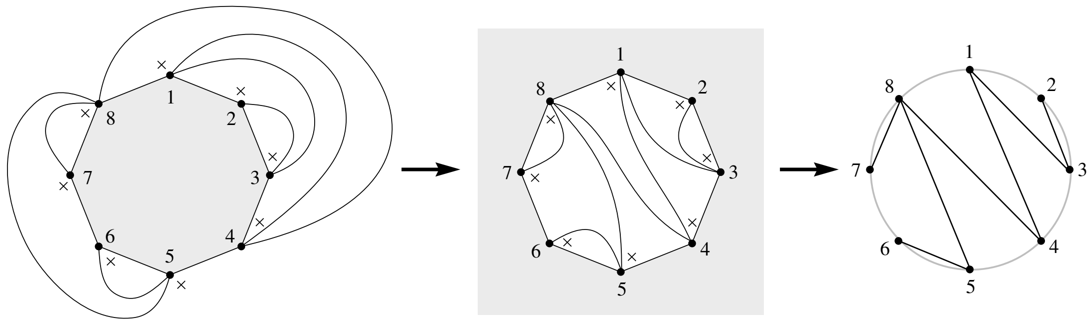

This correspondence between inequivalent minimal transitive factorizations of a fixed full cycle and non-crossing trees on the circle is essentially the same as a bijection credited to Postnikov in [68].

In closing, we mention that Theorem 4.3.9 is easily modified to describe the pruning of simple cacti from *properly labelled* descent-marked hybrid maps. This can be applied, as in §4.3.5, to determine a generating series for inequivalent minimal transitive β-factorizations of permutations composed of two cycles. The form of this series is not particularly illuminating, so we do not include it here. The usual difficulties are encountered when attempting to extend the method to the enumeration of β-factorizations of permutations with arbitrary cycle type.

Appendix A

Canonical Forms for Inequivalent Factorizations of Class (n_1, n_2, n_3)

Fix $n_1, n_2, n_3 \geq 1$, and let $\mathcal{S}^i_{n_i} = \{1^i, \ldots, n^i_i\}$ for $i = 1, 2, 3$. Let f be a minimal transitive factorization of $(1^1 \cdots n^1_1)(1^2 \cdots n^2_2)(1^3 \cdots n^3_3)$ whose left cut is $(1^1\, 1^2)$ and whose right cut is $(1^3\, a^1)$, for some a with $1 \leq a \leq n_1$. Then exactly one of the following five cases is applicable. Throughout, the symbol $\big[\pi\big]$ represents a minimal transitive factorization of π, and the notation $g * h$ indicates that the factorizations g and h are to be concatenated in the given order.

Case 1: There are unique $q,\ p_1,\ p_2,\ p_3$ with $a \leq q \leq p_1 \leq n_1$, $1 \leq p_2 \leq n_2$ and $1 \leq p_3 \leq n_3$ such that $f \sim L * (1^1\, 1^2) * C * (1^3\, a^1) * R$, where

$$\begin{aligned}
L &= \big[(1^1 (p_1+1)^1 \cdots n^1_1)\big] * \big[(1^2 (p_2+1)^2 \cdots n^2_2)\big] * \big[(1^3 (p_3+1)^3 \cdots n^3_3)\big],\\
C &= \big[(1^1 \cdots a^1 (q+1)^1 \cdots p^1_1 1^2 \cdots p^2_2)\big],\\
R &= \big[(a^1 \cdots q^1 1^3 \cdots p^3_3)\big].
\end{aligned}$$

Case 2: There are unique $q,\ p_1,\ p_2,\ p_3$ with $1 \leq p_1 < a \leq q \leq n_1$, $1 \leq p_2 \leq n_2$ and $1 \leq p_3 \leq n_3$ such that $f \sim L * (1^1\, 1^2)(1^3\, a^1) * R$, where

$$\begin{aligned}
L &= \big[(1^1 (p_1+1)^1 \cdots a^1 (q+1)^1 \cdots n^1_1)\big] * \big[(1^2 (p_2+1)^2 \cdots n^2_2)\big] * \big[(1^3 (p_3+1)^3 \cdots n^3_3)\big],\\
R &= \big[(a^1 \cdots q^1 1^3 \cdots p^3_3)\big] * \big[(1^1 \cdots p^1_1 1^2 \cdots p^2_2)\big].
\end{aligned}$$

Case 3: There are unique q, p_1, p_2, p_3 with $1 \le q < a \le p_1 \le n_1$, $1 \le p_2 \le n_2$ and $1 \le p_3 \le n_3$ such that $f \sim L * (1^1\,1^2)(1^3\,a^1) * R$, where

$$L = \big[(1^1(p_1+1)^1 \cdots n_1^1)\big] * \big[((q+1)^1 \cdots a^1)\big] * \big[(1^2(p_2+1)^2 \cdots n_2^2)\big] * \big[(1^3(p_3+1)^3 \cdots n_3^3)\big],$$
$$R = \big[(a^1 \cdots p_1^1 1^2 \cdots p_2^2 1^1 \cdots q^1 1^3 \cdots p_3^3)\big].$$

Case 4: There are unique r, p_1, p_2, p_3 with $1 \le a \le p_1 \le n_1$, $1 \le r \le p_2 \le n_2$, and $1 \le p_3 \le n_3$ such that $f \sim L * (1^1\,1^2) * C * (1^3\,a^1) * R$, where

$$L = \big[(1^1(p_1+1)^1 \cdots n_1^1)\big] * \big[(1^2(p_2+1)^2 \cdots n_2^2)\big] * \big[(1^3(p_3+1)^3 \cdots n_3^3)\big],$$
$$C = \big[(1^1 \cdots a^1(r+1)^2 \cdots p_2^2)\big],$$
$$R = \big[(a^1 \cdots p_1^1 1^2 \cdots r^2 1^3 \cdots p_3^3)\big].$$

Case 5: There are unique r, q, p_1, p_2, p_3 with $q \le p_1 < r < q < a \le n_1$, $1 \le p_2 \le n_2$ and $1 \le p_3 \le n_3$ such that $f \sim L * (1^1\,1^2)(1^3\,a^1) * R$, where

$$L = \big[(1^1(p_1+1)^1 \cdots r^1)\big] * \big[((q+1)^1 \cdots a^1)\big] * \big[(1^2(p_2+1)^2 \cdots n_2^2)\big] * \big[(1^3(p_3+1)^3 \cdots n_3^3)\big],$$
$$R = \big[(1^1 \cdots p_1^1 1^2 \cdots p_2^2)\big] * \big[(a^1 \cdots n_1^1 r^1 \cdots q^1 1^3 \cdots p_3^3)\big].$$

Appendix B

Future Work

Of the numerous questions left open in our investigations, we feel that the following three are of the greatest importance. Of these, the first is the top priority.

- A combinatorial proof of Theorem 2.6.11 remains a major goal of future research. The rational form of $\Gamma_m(\mathbf{z}, u)$ makes it an enticing object of study, yet assigning combinatorial meaning to the iterated differential operator $D = \sum_i \frac{z_i}{1-uz_i} \frac{\partial}{\partial z_i}$ seems difficult. We conjecture that the effect of this operator is to build certain trees (by repeatedly attaching paths) that could take the place of the path $\mathscr{P}$ in analogues of Theorems 2.7.11 and 2.7.14. That is, edges would be repeatedly attached from the leaves of these trees so as to form smooth maps. Schaeffer's *conjugation of trees* [62] could potentially be helpful in this context, but edge- and face-labelling complicates matters.

  In general, the possibility that Lemma 2.7.6 could be used to build smooth maps from trees in some canonical manner should be explored further. Smooth planar maps of descent class $(1, \ldots, 1)$ would provide a natural starting point for such investigations. These are counted by the *simple Hurwitz numbers* $H_0([1^n]) = n^{n-3}(2n-2)!$. While this formula is suggestive of various combinatorial interpretations, no bijective proof has been found.

  Note that the close similarities between §2.7 and §3.4.5 make it certain that a better understanding of Theorem 2.6.11 would immediately lead to further insight into the double Hurwitz problem.

- Let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be a partition of n. Let $r_\alpha = n + \ell(\alpha) - 2$ and let $G_\alpha = n \cdot H_0(\alpha) / \prod_i \alpha_i$. For $S = \{i_1, \ldots, i_k\} \subset [m]$ with $i_1 < \cdots < i_k$, let $\alpha_S = (\alpha_{i_1}, \ldots, \alpha_{i_k})$. In [72], Vakil gives the recurrence

$$G_\alpha = (r_\alpha - 1) \sum_{(\beta,\gamma)} \binom{r_\alpha - 2}{r_\beta, r_\gamma} |\beta||\gamma| G_\beta G_\gamma + \frac{1}{2} \sum_{k=1}^{m} \alpha_k \sum_{p+q=\alpha_k} \sum_{(\beta,\gamma)} \binom{r_\alpha - 1}{r_\beta, r_\gamma} G_\beta G_\gamma,$$

where the first sum is over all pairs $(\beta, \gamma) = (\alpha_S, \alpha_T)$, where (S, T) is a partition of $[m]$, and the second sum is over all pairs $(\beta, \gamma) = (\alpha'_S, \alpha'_T)$, where $\alpha' = (\alpha_1, \ldots, \widehat{\alpha_k}, \ldots, \alpha_m, p, q)$ and (S, T) is a partition of $[m + 1]$ having $p \in S, q \in T$. For example, with $\alpha = (4, 1)$, we have

$$\begin{aligned} G_{(4,1)} = 4 \left\{ \binom{3}{3} \cdot 4 \cdot 1 \cdot G_{(4)} G_{(1)} + \binom{3}{0} \cdot 1 \cdot 4 \cdot G_{(1)} G_{(4)} \right\} \\ + \frac{1}{2} \cdot 4 \cdot \left\{ \binom{4}{2} G_{(1,1)} G_{(3)} + \binom{4}{3} G_{(2,1)} G_{(2)} + \binom{4}{4} G_{(3,1)} G_{(1)} \right. \\ \left. + \binom{4}{2} G_{(3)} G_{(1,1)} + \binom{4}{1} G_{(2)} G_{(2,1)} + \binom{4}{0} G_{(1)} G_{(3,1)} \right\} \end{aligned}$$

Using results from [29], the recursion can also be simplified to

$$G_\alpha = \frac{r_\alpha (r_\alpha - 1)}{m - 1} \sum_{\{\beta,\gamma\}} \binom{r_\alpha - 2}{r_\beta, r_\gamma} |\beta||\gamma| G_\beta G_\gamma,$$

where the sum extends over all $\{\beta, \gamma\} = \{\alpha_S, \alpha_T\}$, where (S, T) is a partition of $[m]$. However, the

By Corollary 2.4.22, note that G_α is the number of planar, vertex-rooted, edge- and face-labelled maps of descent class α. Interpreting the recursions above in terms of the combinatorics of such maps would be a great step forward. No progress has yet been made along these lines.

- More work should be done to manipulate Corollary 4.3.14 into a more enlightening form. To this end, comparison with the counting series for each of the five categories of factorizations listed in Appendix A could be helpful. Further attempts should also be made at extending the switching construction of GJL [32], with graphical intuition potentially being a valuable source of insight.

Bibliography

[1] V. Arnol'd, *Critical points of functions and classification of caustics*, Uspekihi Mat. Nauk. **29** (1994), 243–244.

[2] ______, *Topological classification of complex trigonometric polynomials and the combinatorics of graphs with an equal number of vertices and edges*, Functional Anal. **39** (1996), 1–14.

[3] I. Beck, *Cycle decomposition by transpositions*, J. Combin. Theory, Ser. A **23** (1977), 198–207.

[4] C. Berge, *Principes de Combinatoire*, Dunod, Paris, 1968.

[5] E.A. Bertram and V.K. Wei, *Decomposing a permutation into two large cycles: an enumeration*, SIAM J. Algebraic and Discete Methods **1** (1980), 450–461.

[6] P. Biane, *Parking functions of types A and B*, Electronic J. Combinatorics **9** (2002), no. 7.

[7] G. Boccara, *Nombre de représentations d'une permutation comme produit de deux cycles de longueurs données*, Discrete Math. **29** (1980), 105–134.

[8] M. Bousquet-Mélou and G. Schaeffer, *Enumeration of planar constellations*, Adv. in Appl. Math. **24** (2000), 337–368.

[9] P. Cartier and D. Foata, *Problèmes combinatoires de commutation et réarrangements*, Lecture Notes in Mathematics, Vol. 85, Springer-Verlag, Berlin, 1969.

[10] A. Cayley, *A theorem on trees*, Quart. J. Math. **23** (1889), 376–378.

[11] M. Cohn and A. Lempel, *Cycle decomposition by disjoint transpositions*, J. Combin. Theory, Ser. A **13** (1972), 83–89.

[12] M. Crescimanno and W. Taylor, *Large N phases of chiral QCD_2*, Nuclear Phys. B **437** (1995), 3–24.

[13] 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. Institute Hung. **4** (1959), 63–70.

[14] S. Dulucq and J.G. Penaud, *Cordes, arbres et permutations*, Discrete Math. **117** (1993), 89–105.

[15] J. Edmonds, *Combinatorial representation for polyhedral surfaces*, Notices Amer. Math. Soc. **7** (1960), 646.

[16] J.A. Eidswick, *Short factorizations of permutations into transpositions*, Discrete Math. **73** (1989), 293–243.

[17] T. Ekedahl, S. Lando, M. Shapiro, and A. Vainshtein, *On Hurwitz numbers and Hodge integrals*, C.R. Acad. Sci. Paris Ser. I Math. **328** (1999), 1175–1180.

[18] A. Erdélyi and I.M.H. Etherington, *Some problems of non-associative combinations II*, Edinburgh Math. Notes **32** (1941), 7–12.

[19] C.L. Ezell, *Branch point structure of covering maps onto nonorientable surfaces*, Trans. Amer. Math. Soc. **243** (1978), 123–133.

[20] C. Faber, *A conjectural description of the tautological ring of the moduli space of curves*, preprint, math.AG/9711218.

[21] H.K. Farahat and F.R.S. Higman, *The centres of symmetric group rings*, Proc. Roy. Soc. London, Ser. A **250** (1950), 212–221.

[22] W. Fulton, *Algebraic Topology: A First Course*, Springer, New York, 1995.

[23] V. Goryunov and S. Lando, *On the enumeration of meromorphic functions on the line*, Proceedings of the Conference Dedicated to V.I. Arnol'd's 60th Birthday (Toronto), American Mathematical Society, 1997, pp. 209–224.

[24] I.P. Goulden, *Private communication*.

[25] ______, *A differential operator for symmetric functions and the combinatorics of multiplying transpositions*, Trans. Amer. Math. Soc. **344** (1994), 149–163.

[26] I.P. Goulden and D.M Jackson, *Combinatorial Enumeration*, Wiley Interscience, New York, 1983.

[27] I.P. Goulden and D.M. Jackson, *The combinatorial relationship between trees, cacti and certain connection coefficients for the symmetric group*, European J. Combin. **13** (1992), 357–365.

[28] ______, *Symmetric functions and Macdonald's result for top connexion coefficients in the symmetric group*, J. Algebra **166** (1994), 364–378.

[29] ______, *Transitive factorizations and holomorphic mappings of the sphere*, Proc. Amer. Math. Soc. **125** (1997), 51–60.

[30] ______, *A proof of a conjecture for the number of ramified coverings of the sphere by the torus*, J. Combin. Theory, Ser. A **88** (1999), 246–258.

[31] ______, *Transitive factorizations in the symmetric group, and combinatorial aspects of singularity theory*, Europ. J. Combinatorics **21** (2000), 1001–1016.

[32] I.P. Goulden, D.M. Jackson, and F.G. Latour, *Inequivalent transitive factorizations into transpositions*, Canad. J. Math **53** (2001), 758–779.

[33] I.P. Goulden, D.M. Jackson, and A. Vainshtein, *The number of ramified coverings of the sphere by the torus and surfaces of higher genera*, Ann. Combin. **4** (2000), 27–46.

[34] 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), 563–581.

[35] ______, *Faber's intersection number conjecture and genus 0 double Hurwitz numbers*, In preparation (2004).

[36] ______, *Towards the geometry of double Hurwitz numbers*, submitted for publication (2004), math.AG/0309440.

[37] I.P. Goulden and A. Nica, *A direct bijection for the Harer-Zagier formula*, Submitted for publication (2004).

[38] I.P. Goulden and S. Pepper, *Labelled trees and factorizations of a cycle into transpositions*, Discrete Math. **113** (1993), 263–268.

[39] I.P. Goulden and A. Yong, *Tree-like properties of cycle factorizations*, J. Combin. Theory, Ser. A **98** (2002), 106–117.

[40] A. Goupil and F. Bédard, *The poset of conjugacy classes and decompositions of products in the symmetric group*, Can. Bull. of Math. **35** (1992), 152–160.

[41] A. Goupil and G. Schaeffer, *Factoring n-cycles and counting maps of given genus*, European J. Combin. **19** (1998), 819–834.

[42] J.L. Gross and T.W. Tucker, *Topological Graph Theory*, Wiley Interscience, New York, 1987.

[43] J. Harer and D. Zagier, *The Euler characteristic of the moduli space of curves*, Invent. Math. **85** (1986), 457–486.

[44] A. Hurwitz, *Uber Riemann'sche Flächen mit gegebenen Verzweigungspunkten*, Math. Ann. **39** (1891), 1–60.

[45] ______, *Uber die anzahl der Riemann'schen Flächen mit gegebenen Verzweigungspunkten*, Math. Ann. **55** (1902), 53–66.

[46] D.M. Jackson, *Counting cycles in permutations by group characters, with an application to a topological problem*, Trans. Amer. Math. Soc. **299** (1987), 785–801.

[47] ______, *Some problems associated with products of conjugacy classes of the symmetric group*, J. Combin. Theory, Ser. A **49** (1988), 363–369.

[48] D. Kim and S. Seo, *Transitive cycle factorizations and prime parking functions*, preprint (2002).

[49] J.Q. Longyear, *A peculiar partition identity*, Discrete Math. **78** (1989), 115–118.

[50] E. Looijenga, *The complement of the bifurcation variety of a simple singularity*, Invent. Math. **23** (1973), 105–116.

[51] I.G. Macdonald, *Symmetric Functions and Hall Polynomials*, Oxford University Press (Clarendon), Oxford, 1979.

[52] A.D. Mednykh, *Nonequivalent coverings of Riemann surfaces with a prescribed ramification type*, Siber. Math. J. **25** (1984), 606–625.

[53] ______, *Branched coverings of Riemann surfaces whose branch orders coincide with the multiplicity*, Comm. in Algebra **18** (1990), 1517–1533.

[54] P. Moszkowski, *A solution to a problem of Dénes: A bijection between trees and factorizations of cyclic permutations*, European J. Combinatorics **10** (1989), 13–16.

[55] M. Noy, *Enumeration of noncrossing trees on a circle*, Discrete Math. **180** (1998), 301–313.

[56] A. Okounkov and R. Pandharipande, *Gromov-Witten theory, Hurwitz numbers, and matrix models, I*, preprint (2001), math.AG/0101147.

[57] D. Poulalhon, *Graphes et décompositions de permutations*, Mémoire de DEA, LIX, École Polytechnique, July 1997.

[58] D. Poulalhon and G. Schaeffer, *Factorizations of large cycles in the symmetric group*, Discrete Math. **254** (2002), 433–458.

[59] H. Prüfer, *Neuer Beweis eines Satzes über Permutation*, Arch. Math. Phys. **27** (1918), 142–144.

[60] A. Rattan, *Cycle factorizations and parking functions*, Submitted for publication (2004).

[61] B.E. Sagan, *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, Wadsworth, Belmont, 1991.

[62] G. Schaeffer, *Conjugaison d'arbres et cartes combinatoires aléatoires*, Phd thesis, Université Bordeaux I, 1998.

[63] J-P. Serre, *Linear Representations of Finite Groups*, Springer-Verlag, Berlin, 1974.

[64] B. Shapiro, M. Shapiro, and A. Vainshtein, *Ramified coverings of S^2 with one degenerate branching point and enumeration of edge-ordered graphs*, Adv. in Math. Sci. **34** (1997), 219–228.

[65] 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.

[66] R.P. Stanley, *Factorization of permutations into n-cycles*, Discrete Math. **37** (1981), 255–262.

[67] ———, *Enumerative Combinatorics: Volume 1*, Cambridge University Press, 1997.

[68] ———, *Enumerative Combinatorics: Volume 2*, Cambridge University Press, 1999.

[69] J. Stillwell, *Classical Topology and Combinatorial Group Theory, 2nd ed.*, Springer, New York, 1993.

[70] V. Strehl, *Minimal transitive products of transpositions: the reconstruction of a proof by A. Hurwitz*, Séminaire Lotharingien de Combinatoire **37** (1996).

[71] W.T. Tutte, *Graph Theory*, Addison-Wesley, New York, 1984.

[72] R. Vakil, *Recursions, formulas, and graph-theoretic interpretations of ramified coverings of the sphere by surfaces of genus 0 and 1*, Trans. Amer. Math. Soc. **353** (2001), 4025–4038.

[73] D.W. Walkup, *How many ways can a permutation be factored into two cycles?*, Discrete Math. **28** (1979), 315–319.

[74] H. Wilf, *generatingfunctionology*, Academic Press, San Diego, 1990.

List of Notation

$[n]$ . . . 4
$|S|$ . . . 4
$\alpha \models n$ . . . 4
$\ell(\alpha)$ . . . 4
$\alpha \vdash n$ . . . 4
$[1^{m_1} 2^{m_2} \cdots]$ . . . 4
$|\operatorname{Aut}(\alpha)|$ . . . 4
$\mathbb{Q}[[S]$ . . . 4
$\mathbf{x}^{\mathbf{i}}$ . . . 4
$[\mathbf{x}^{\mathbf{i}}]\, f(\mathbf{x})$ . . . 4
$\mathbf{i}!$ . . . 4
$\frac{df}{dx}$ . . . 4
$\int f(x)\, dx$ . . . 4
h_k . . . 6
$A(t) \circ \Delta^+(t\,;\,\mathbf{x})$ . . . 6
$(a_0, \ldots, a_n)^\circ$ . . . 7
S° . . . 7
$\mathfrak{S}_n$ . . . 7
ι . . . 7
$\operatorname{orb} \pi$ . . . 7
$\ell(\pi)$ . . . 8
$\mathscr{C}_\alpha$ . . . 8
$\mathbb{C}\mathfrak{S}_n$ . . . 8
K_α . . . 8
F_θ . . . 8
χ_α^θ . . . 9
f^θ . . . 9
$\widehat{v}$ . . . 9
$\deg(v)$ . . . 9
$((v_0, e_0), \ldots, (v_k, e_k))^\circ$ . . . 9
$\mathscr{M} \backslash e$ . . . 11
$\deg(F)$ . . . 11
$r_g(\alpha)$ . . . 20
$H_g(\alpha)$ . . . 22
$\Phi^{(g)}(z, \mathbf{p}, u)$ . . . 22
w . . . 26, 114, 132, 158, 173
$\Psi_m^{(g)}(\mathbf{x}, u)$ . . . 27, 51
w_i . . . 27, 122, 162
$\mathscr{G}_f$ . . . 33
$\mathscr{M}_f$ . . . 37, 103
MAP . . . 37
$\mathbb{D}_i(\alpha)$ . . . 49
$\mathfrak{S}(\alpha)$ . . . 49
$\Psi_m(\mathbf{x}, u)$ . . . 51
$\Gamma_m(\mathbf{z}, u)$ . . . 56, 66
$\mathscr{M}^c$ . . . 59, 119, 133
$\Gamma_m^{(g)}(\mathbf{z}, u)$ . . . 66
$\mathcal{A}_F(e)$ . . . 71
$\mathscr{M} \oplus (c, c')^e_{s,t}$ . . . 71
$\mathscr{M}^\dagger$ . . . 106
$\Psi_m^{(g)}(\mathbf{x}, \mathbf{p}, u),\ \Psi_m(\mathbf{x}, \mathbf{p}, u)$ . . . 112
$P(z)$ . . . 115

$\Psi_{m,k}(\mathbf{x})$ 117

$\Gamma_m^{(g)}(\mathbf{z}, \mathbf{p}, u)$, $\Gamma_m(\mathbf{z}, \mathbf{p}, u)$ 121

ξ 122, 135

ξ_i 122

P_i, P_i' 122

$H_g(\alpha, \beta)$ 129

$r_g(\alpha, \beta)$ 129

∞ 131

$\Theta_m^{(g)}(\mathbf{x}, \mathbf{q}, u)$, $\Theta_m(\mathbf{x}, \mathbf{q}, u)$ 131

$Q(z)$ 132

$\Lambda_m^{(g)}(\mathbf{x}, \mathbf{q}, u)$, $\Lambda_m(\mathbf{x}, \mathbf{q}, u)$ 134

$\mathfrak{F}_g(\pi)$ 139

$\mathscr{M}^\#$ 140

$\alpha_k(f)$ 141

$\mu_k(f)$ 141

$f \sim g$ 145

$\tilde{f}$ 145

$[\pi]$ 146

$f * g$ 146

$\tilde{h}_n$ 147

$h(x)$ 147

$\widetilde{H}_g(\alpha)$ 148

$\widetilde{\Phi}^{(g)}(z, \mathbf{p}, u)$ 148

$\widetilde{\Psi}_m(\mathbf{x}, u)$ 151

$s_k(x)$ 155

$\widetilde{\Psi}_m^{(g)}(\mathbf{x}, \mathbf{p}, u)$, $\widetilde{\Psi}_m(\mathbf{x}, \mathbf{p}, u)$ 158

$\widetilde{\Gamma}_m^{(g)}(\mathbf{z}, \mathbf{t}, \mathbf{p}, u)$, $\widetilde{\Gamma}_m(\mathbf{z}, \mathbf{t}, \mathbf{p}, u)$ 161

$\sigma(\mathbf{x})$ 166

$\langle f \rangle$ 166

Index

base corner 60, 119
base vertex 60, 119
β-factorization 129
boundary walk 11
branch 60, 89, 119
branched covering 28
bridge 9

cactus 104, 114, 132, 158, 173
canonical descent sets 49
carrier 35
Cartier-Foata commutation monoid 148
circle chord diagram 143
circular permutation 68
circular sequence 7
circulator 12
class 13, 97
complete symmetric function 6
component
  of a factorization 17, 97
  of a graph 9
composition 4
conjugacy class 8
connection coefficients 8
constellation 100
core 59, 119, 133, 161
corner 11, 99

cut 15
cycle factorization 110
cycle index 110
cycle type 8
cyclic list 7

descent
  of a cyclic list 7
  of a descent-marked polymap 157
  of a hybrid map 131
  of a map 39
  of a polymap 99
descent class 48, 111, 157
descent cycle 39, 100
descent partition 39, 100, 157
descent set 39, 100, 157
descent-marked polymap 157
digon 124
disjoint cycle decomposition 8
double Hurwitz numbers 129
double Hurwitz problem 129

embedding theorem 12
equivalent
  β-factorizations 172
  cycle factorizations 145
Euler-Poincaré formula 10

face . . . 10, 99
factor type . . . 97
factorization
  β-factorization . . . 129
  cycle factorization . . . 110
  generic . . . 97
  into transpositions . . . 13
  k-cycle factorization . . . 116
formal power series . . . 4
full cycle . . . 7

genus
  of a factorization . . . 20, 98
  of a map . . . 10
graph . . . 9
  of a factorization . . . 33
group algebra . . . 8

half-edges . . . 11
Hurwitz enumeration problem . . . 22
Hurwitz formula . . . 23
Hurwitz numbers . . . 22
Hurwitz series . . . 22
hybrid cactus . . . 132
hybrid map . . . 130
  descent-marked . . . 172
  properly labelled . . . 131
  smooth . . . 133
hybrid path . . . 135

increasing cyclic list . . . 7
index
  of a branch . . . 63, 120
  of a vertex . . . 89
  of an $\mathcal{N}$-branch . . . 89

isomorphism . . . 9, 11

join . . . 15

k-cycle factorization . . . 116

labelling . . . 10, 11, 157
Lagrange inversion . . . 4
leaf . . . 59, 119, 133
length
  of a composition . . . 4
  of a factorization . . . 13, 97
loop . . . 9

map . . . 10
  conventions . . . 38
  descent structure . . . 39
  hybrid . . . 130
  of a factorization . . . 37
  properly labelled . . . 49
  smooth . . . 59
  trivial . . . 88
marked polymap . . . 157
minimal factorization . . . 16
minimal transitive factorization . . . 19, 98

$\mathcal{N}$-branch . . . 89
non-crossing tree . . . 143, 175
nondecreasing cyclic list . . . 7
normal indexing . . . 61, 89

one-maps . . . 43
orbit . . . 35
ordered path . . . 73
orthogonal idempotents . . . 8

partition . . . 4

polygon . . . 99, 130
polygon index . . . 111
polygon type . . . 103, 130
polygonal path . . . 125
polymap . . . 99
  conventions . . . 99, 110
  descent structure . . . 100
  descent-marked . . . 157
  marked . . . 157
  of a factorization . . . 103
  properly labelled . . . 111
  smooth . . . 119
possible cut . . . 151
pruning . . . 43, 64, 120, 161

rank . . . 16
reduction . . . 106
rotation system . . . 12
rotator . . . 37, 99, 131, 157

simple cactus . . . 132, 173
simple edge . . . 131
skeleton . . . 11
smooth . . . 59, 119, 133
stem . . . 60, 119
submap . . . 89
support . . . 7
switching algorithm . . . 152
symmetric group . . . 7

tail . . . 73, 125, 136
transitive factorization . . . 13, 97
transitive rank . . . 19
transposition . . . 7
tree . . . 10
tree pruning bijection . . . 64
tree series . . . 26
trivial map . . . 88

valid labelling . . . 157