

Symmetric Function Generalizations of Graph Polynomials

by

Timothy Yi-Chung Chow

A.B., Princeton University, New Jersey (1991)

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1995

©1995, Timothy Yi-Chung Chow

The author hereby grants to MIT permission to reproduce and
to distribute copies of this thesis document in whole or in part.

Signature of Author
Department of Mathematics
April 26, 1995

Certified by
Richard P. Stanley
Professor of Applied Mathematics
Thesis Supervisor

Accepted by
David Vogan
Chairman, Departmental Graduate Committee
Department of Mathematics

Symmetric Function Generalizations of Graph Polynomials

by
Timothy Yi-Chung Chow

Submitted to the Department of Mathematics
on April 26, 1995, in partial fulfillment of the
the requirements for the degree of
Doctor of Philosophy

Abstract

Motivated by certain conjectures regarding immanants of Jacobi-Trudi matrices, Stanley has recently defined and studied a symmetric function generalization X_G of the chromatic polynomial of a graph G . Independently, Chung and Graham have defined and studied a directed graph invariant called the cover polynomial. The cover polynomial is closely related to the chromatic polynomial and to the rook polynomial, so it is natural to ask if one can mimic Stanley's construction and generalize the cover polynomial to a symmetric function. The answer is yes, and the bulk of this thesis is devoted to the study of this generalization, which we call the path-cycle symmetric function. We obtain analogues of some of Stanley's theorems about X_G and we generalize some of the theory of the cover polynomial and the rook polynomial. In addition, we are led to define a symmetric function basis that seems to be a "natural" generalization of the polynomial basis

$$\left\{ \binom{x+k}{d} \right\}_{k=0,1,\dots,d}$$

and we prove a combinatorial reciprocity theorem that gives an affirmative answer to Chung and Graham's question of whether the cover polynomial of a digraph determines the cover polynomial of its complement. The reciprocity theorem also ties together several scattered results in the literature that previously seemed unrelated.

In the remainder of the thesis, we prove some miscellaneous results about Stanley's function X_G and we also sketch briefly in the introduction a possible approach to generalizing other graph polynomials (and a few other combinatorial polynomials) to symmetric functions.

Thesis supervisor: Richard P. Stanley
Title: Professor of Applied Mathematics

To my parents

Acknowledgments

I am deeply grateful to my thesis advisor, Richard Stanley. Without him, not only would this thesis not exist, but I would be entirely ignorant of this beautiful area of combinatorics (and many other areas besides) to this day. Moreover, while many graduate students have a long stream of advisor horror stories, and regard their graduate school experience as an agony in eight fits, I have no complaints at all about Richard. My experience here at M.I.T. has been remarkably pleasant, with the main source of stress being my frustration at my own ignorance and laziness.

I also wish to thank Ira Gessel and Gian-Carlo Rota, who have been consistently encouraging and unreasonably patient with my frequent silly questions and comments. I am also indebted to John Stembridge, whose SF and posets packages for Maple were of immeasurable help to me while I was doing my research.

Seeing so many other students struggling to pay for their education, I am reminded how fortunate I am to enjoy the generous financial support of the M.I.T. mathematics department (through teaching assistantships) and the National Science Foundation (through graduate fellowship stipends), who have made it possible for me to pursue my studies without having to worry about money.

To my family and my friends from Princeton, Duluth, Chinese Christian Fellowship, Graduate Christian Fellowship, and the math department—to all who have loved me, lived with me, taught me, hung out with me, laughed with me, tolerated me, encouraged me, and prayed for me—I apologize for not expressing my appreciation to you more often. Without you, the past four years would have been bleak indeed, and I would have been lucky to survive at all, let alone write a thesis and graduate from M.I.T.

Last of all, and most of all, I thank God for giving me life, filling it with joy, and creating this thing of beauty that we call mathematics.

Contents

Acknowledgments	7
Introduction	11
Chapter 1: Preliminaries	19
1. Graphs, digraphs, and symmetric functions	19
2. χ_G , X_G , $C(D)$, and Ξ_D	21
3. Connection with rook theory	23
Chapter 2: The Path-Cycle Symmetric Function	24
1. Basic facts	24
2. Reciprocity	30
3. Rook theory	38
4. The Poset Chain Conjecture	49
Chapter 3: The Chromatic Symmetric Function	53
1. G -ascents	53
2. The Poset Chain Conjecture revisited	55
3. Reconstruction	57
4. Superfication	59
5. $X_G(t)$	64
6. Counterexamples	66
Bibliography	67

Introduction

Mathematicians are notorious for sterilizing their technical papers so that there remains no taint of motivation. All traces of how the proofs were originally discovered are completely obliterated—or, at the very least, carefully disguised—so that maximum brevity and elegance are achieved. While this approach is ideal for the reader who wishes to use the paper solely as a reference (or as an object for aesthetic contemplation), the student who really wants to understand the paper is done a disservice. On the other hand, an informal presentation that gives a great deal of intuitive motivation and works through many concrete examples may be ideal pedagogically, but it is frustrating for the researchers who wish to extract only what they need for their own work.

I have adopted a compromise here. The body of this thesis is written in the customary definition-theorem-proof style, with some motivational comments scattered haphazardly. However, this introduction is written in a more informal style, and I try to show how my work fits into the larger scheme of things. I provide precise definitions here only when I feel that they are critical to the exposition; for full definitions the reader should consult Chapter 1 and the references cited therein.

The principal objects of study in this thesis may not seem natural or intrinsically interesting a priori, so let me begin by describing two previous lines of research that *are* interesting and natural. We will then see how the present thesis arose out of the unexpected convergence of those two lines of study.

The first of these pre-existing lines of research is the study of *immanants* of matrices. To understand what these are, recall that the *determinant* of a matrix

may be thought of as a sum of the form

$$\sum_{\sigma \in S_n} (-1)^\sigma \prod_{i=1}^n x_{i, \sigma(i)},$$

where S_n denotes the symmetric group on n letters. Similarly, the permanent of a matrix is a sum of the form

$$\sum_{\sigma \in S_n} \prod_{i=1}^n x_{i, \sigma(i)}.$$

Notice that the only difference between these two expressions lies in the coefficients: in the case of the determinant, the coefficients are given by the sign character of S_n , and in the case of the permanent, the coefficients are given by the trivial character of S_n . This observation suggests that perhaps expressions of the form

$$\sum_{\sigma \in S_n} \chi^\lambda(\sigma) \prod_{i=1}^n x_{i, \sigma(i)}$$

(where χ^λ is an irreducible character of S_n), called *immanants*, might be interesting objects of study. It turns out, not entirely unexpectedly, that the theory of immanants is not as rich as the theory of determinants (see [Gre] for a brief survey of work done on immanants). However, in the special case of *Jacobi-Trudi matrices*, which are matrices that arise in the theory of symmetric functions, there is a rich array of combinatorial conjectures related to immanants. While it is tempting to give a summary of these beautiful conjectures, we shall restrict ourselves to describing just one of them, due to Stanley and Stembridge [S-S], since it is the only one directly relevant to our present purposes. The interested reader is referred to [Gre] and [S-S] for the full story.

To state the Stanley-Stembridge conjecture (also known as the *Poset Chain Conjecture*), we must first define a certain invariant X_P of a poset P . A *coloring* of a poset P is a map κ that sends each element of P to a positive integer (or “color”) and whose fibers are totally ordered subsets (or “chains”) of P . Then X_P is defined by

$$X_P \stackrel{\text{def}}{=} \sum_{\kappa} \prod_{p \in P} x_{\kappa(p)},$$

where the sum is over all colorings of P . Since the colors may be permuted indiscriminately, X_P is a symmetric function, i.e., it is invariant under any permutation of the indeterminates x_i . By the well-known fundamental theorem of symmetric functions, every symmetric function has a unique expression as a polynomial in the *elementary symmetric functions* e_n defined by

$$e_n = \sum_{0 < i_1 < i_2 < \dots < i_n} x_{i_1} x_{i_2} \cdots x_{i_n}.$$

The Poset Chain Conjecture states that if P does not contain an induced subposet isomorphic to the disjoint union of a three-element chain and a one-element chain, then X_P has nonnegative coefficients when expressed as a polynomial in the elementary symmetric functions. (For the exact connection between this conjecture and the immanants of Jacobi-Trudi matrices, the reader is referred to [S-S].) The Poset Chain Conjecture is supported by considerable numerical evidence and some partial results (notably Gasharov's theorem [Ga1] that under the hypotheses of the conjecture, X_P has nonnegative coefficients when expanded in terms of Schur functions), but remains open.

The next step in the story is a simple but important observation of Stanley that the invariant X_P can be extended from a poset invariant to a graph invariant. More specifically, let $G(P)$ denote the *incomparability graph* of P , i.e., the graph whose vertex set is the set of elements of P and in which two vertices are adjacent if and only if they are *not* comparable as elements of P . Note that chains in P correspond to independent sets of $G(P)$ (i.e., subsets of vertices with no edges between them), and that colorings of P correspond to colorings of $G(P)$ in the usual sense of assignments of colors to the vertices of $G(P)$ such that adjacent vertices are never assigned the same color. Now the notion of a coloring makes sense for any graph G and not just incomparability graphs, so we may define

$$X_G \stackrel{\text{def}}{=} \sum_{\kappa} \prod_v x_{\kappa(v)},$$

where the sum is over all colorings of G and the product is over all vertices v of G .

The reason this observation is significant is that the invariant X_G so defined turns out to be a generalization of the *chromatic polynomial* $\chi_G(n)$ of G . (This therefore opens up the possibility of applying the well-established theory of graph colorings to the solution of the Poset Chain Conjecture.) More precisely, recall that $\chi_G(n)$ is by definition the number of colorings of G such that every vertex is assigned a color between 1 and n inclusive. From this we see that if we set n of the indeterminates x_i equal to 1 and the rest equal to 0, X_G becomes $\chi_G(n)$. This specialization procedure (of setting n of the indeterminates equal to 1 and the rest equal to 0) is an important one, so we introduce the notation $g(1^n)$ to represent the result of so specializing the symmetric function g . We will come back to this specialization later.

Let me now break off this account of this line of research temporarily in order to describe the other line of research that was alluded to earlier. This second line of research is the mathematical theory of *juggling*. Apart from the fact that many mathematicians are amateur jugglers, the recreational flavor of juggling gives its mathematical study great intuitive appeal. In the delightful paper [BEGW], the authors show how some intriguing and nontrivial mathematics arises from the analysis of various juggling patterns. In particular, they are led to consider a

poset invariant called the *binomial drop polynomial* [B-G] (which turns out to be equivalent to the *factorial polynomial* of [GJW1]).

In the course of studying the binomial drop polynomial, Chung and Graham [CG1] had the idea of generalizing it to an invariant of a directed graph. More precisely, they made the simple observation that a poset P may be thought of as a directed graph $D(P)$ whose vertices are the elements of P and in which $u \rightarrow v$ is an edge if and only if $u < v$ in P . They were then led to generalize the binomial drop polynomial of a poset to an invariant of an arbitrary directed graph that they call the *cover polynomial*.

Since the cover polynomial is central to this thesis, we digress momentarily to give its definition. A *path-cycle cover* of a directed graph D is a spanning subgraph of D each of whose connected components is either a directed path or a directed cycle. (An isolated vertex is considered to be a path of zero length and an isolated loop is considered to be a cycle of length one.) The cover polynomial $C(D; i, j)$ is then defined to be

$$C(D; i, j) \stackrel{\text{def}}{=} \sum_S \underline{i}^{p(S)} j^{c(S)},$$

where the sum is over all path-cycle covers S of D , $p(S)$ is the number of connected components of S that are paths, $c(S)$ is the number of connected components of S that are cycles, and the underline indicates the *lower factorial*, i.e.,

$$\underline{i}^k \stackrel{\text{def}}{=} i(i-1) \cdots (i-k+1).$$

(This rather mysterious definition is motivated by deletion-contraction considerations; the interested reader is referred to [CG2] for more explanation.)

Returning now to our main account, we remark that based on what we have said so far, the reader would probably not expect any connection between the mathematics of immanants and the mathematics of juggling, other than the fact that posets and graphs enter the picture in both cases. The surprise, however, is that for any poset P , the polynomial (in n) $X_P(1^n)$ is the factorial polynomial of P !

This unexpected connection suggests that we ought to stop for a moment and clarify the precise relationships between all the algebraic invariants mentioned so far. First, in the realm of posets, we have the symmetric function invariant X_P . This becomes, via the specialization $X_P \mapsto X_P(1^n)$, the factorial polynomial or binomial drop polynomial of P . The invariant X_P generalizes to a graph invariant X_G , and $X_G(1^n)$ is the chromatic polynomial of G . (It follows that the chromatic polynomial is a generalization of the factorial polynomial; this has actually been known for a long time [GJW2].) We also have a generalization of the factorial polynomial to a directed graph invariant: the cover polynomial. It is therefore irresistible to ask (and Chung and Graham in fact do ask this) if there is a symmetric function invariant

of a directed graph, analogous to X_G , which turns into the cover polynomial upon applying the specialization $g \mapsto g(1^n)$. This question brings us (finally!) to the main topic of this thesis.

In Chapter 1, among other things, a symmetric function generalization (the *path-cycle symmetric function*) of the cover polynomial is defined, following a suggestion of Stanley. More precisely, the path-cycle symmetric function of a directed graph D is defined by

$$\Xi_D(x; y) \stackrel{\text{def}}{=} \sum_S \tilde{m}_{\pi(S)}(x) p_{\sigma(S)}(y),$$

where the sum is as before over all path-cycle covers S of D , x and y are (countably infinite) sets of commuting independent indeterminates, $\pi(S)$ denotes the integer partition consisting of the lengths of the directed paths in S , $\sigma(S)$ denotes the integer partition consisting of the lengths of the directed cycles of S , \tilde{m} denotes the augmented monomial symmetric functions (the same as the usual monomial symmetric functions except with a constant factor), and p denotes the power sum symmetric functions. (We shall presently give a few more words of motivation for this unwieldy-looking definition.)

It is not difficult to show that the path-cycle symmetric function has the desired property of becoming the cover polynomial under the specialization $g \mapsto g(1^n)$. In Chapter 2, which forms the bulk of this thesis, the properties of the path-cycle symmetric function are investigated, and a number of striking results are obtained. In section 1, we obtain analogues of some of the basic theorems about X_G . In section 3, we obtain generalizations of some of the basic results of rook theory, such as the Möbius inversion formula of [GJW4], and the fundamental inclusion-exclusion formula for rooks. Here already we obtain our first surprise. In the course of generalizing the inclusion-exclusion formula for rooks, we are led to define a certain (vector space) basis for symmetric functions that does not seem to have been studied before but which seems to be an important object. It has some surprising connections with monomial symmetric functions and fundamental quasi-symmetric functions (for example, any symmetric function that is a nonnegative combination of fundamental quasi-symmetric functions also has nonnegative coefficients when expanded in terms of this new symmetric function basis), and it also generalizes the polynomial basis

$$\left\{ \binom{x+k}{n} \right\}_{k=0,1,\dots,n}$$

that arises in many contexts in combinatorics.

Perhaps the most surprising result, though, appears in section 2, which gives a “reciprocity” formula relating the path-cycle symmetric functions of complementary digraphs. (Two digraphs are complementary if the edges of one are precisely the

non-edges of the other.) The way this formula was discovered is instructive: I was trying to find an analogue of [St2, Corollary 2.7], which is a theorem about ωX_G . Here ω is an involution on the space of symmetric functions that arises naturally in many contexts. It is therefore natural to consider what ω does to the path-cycle symmetric function. I did not actually succeed in proving the result I was hoping for, but ended up proving the reciprocity formula instead. I was then amazed to discover afterwards that the reciprocity formula has several beautiful corollaries. For example, when specialized to the context of the cover polynomial, it provides a remarkably simple formula relating the cover polynomials of complementary digraphs:

$$C(D; i, j) = (-1)^d C(D'; -i - j, j),$$

where d is the number of vertices of D . (Chung and Graham had conjectured that some such relation might exist, but did not even have a conjectural formula. In fact, the reciprocity formula for cover polynomials is not an easy one to guess, and it seems that generalizing to the symmetric function context actually makes it easier to find the pattern—although Ira Gessel did independently arrive at the reciprocity formula for cover polynomials without resorting to symmetric functions.) Further specialization of the reciprocity formula to the context of factorial polynomials gives a formula relating factorial polynomials of complementary boards that is much simpler and more elegant than previously known formulas. Also, a different specialization of the reciprocity formula gives a result that Stanley and Stembridge have proved in the course of their study of immanants.

The remainder of the thesis consists of a miscellaneous collection of results about the path-cycle symmetric function (scattered throughout Chapter 2) and about X_G (in Chapter 3). For example, X_G is shown to be reconstructible from the list of vertex-deleted subgraphs of G , and Chung and Graham's concept of G -descents is shown to have an application to the study of X_G . Instead of giving further details about these results, however, I wish to conclude by addressing the question: where do we go from here?

We have seen that X_G and the path-cycle symmetric function are natural objects to study (if one is familiar with the relevant background), and the attractive results obtained in the course of investigating them also provide some a posteriori justification for their study. But once we have derived the main properties of the path-cycle symmetric function, what is there left to do? Is this a dead end?

The title of this thesis suggests a possible direction for further research. We know now that the chromatic polynomial and the cover polynomial have interesting symmetric function generalizations. What about other graph polynomials, or more generally, other combinatorial polynomials? Do they have natural symmetric function generalizations?

To answer this question, let us re-examine how the chromatic polynomial and the cover polynomial are generalized. In both cases, it turns out that the method of

generalization is based on interpreting the polynomial as counting the *total* number of colorings of a certain kind. The generalization is then obtained by enumerating the same set of colorings, but keeping track of the number of times *each color* is used. (This is the motivation for the unusual-looking definition of Ξ_D given above.)

The problem with this idea is that most combinatorial polynomials do not have any obvious interpretation in terms of colorings. However, it is possible to modify the above idea slightly so that it applies to a wider class of algebraic invariants. I will not go into details here, since I hope to do so elsewhere, but the essential idea is to generalize a polynomial by first interpreting it as enumerating the total number of *partitions of a set* of a certain kind, and then obtaining a generalization by enumerating the same set of partitions but keeping track of the *type* of the partition. There are already some indications that this approach ought to be fruitful; for example, Doubilet's theory of symmetric functions [Dou] is based on this idea (and this probably accounts for why his interpretations of the change-of-basis coefficients between the various standard symmetric function bases have proved to be more useful in our work than the formulas in [Mac, Table 1, p. 56] have). The theory of polynomials of *binomial type* [R-R] also generalizes nicely using this idea; the binomial type identity generalizes to the coproduct in the Hopf algebra of symmetric functions. In short, there are many promising directions for further study.

Chapter 1

Preliminaries

In this chapter, we provide some basic definitions and facts. The absence of a reference does not necessarily imply a claim to originality, since some references for “standard” concepts and facts have been omitted.

1. GRAPHS, DIGRAPHS, AND SYMMETRIC FUNCTIONS

We shall assume that reader is familiar with the basic facts about set partitions, posets, Möbius functions, permutations, and so on; a good reference is [St3].

Throughout, the unadorned term *graph* will mean a finite simple labelled undirected graph and the term *digraph* will mean a finite labelled directed graph without multiple edges but possibly with loops and bidirected edges. If G is a graph or a digraph we let $V(G)$ and $E(G)$ denote its vertex set and edge set respectively. If d is a positive integer, we use the notation $[d]$ for the set $\{1, 2, \dots, d\}$. Note that with our conventions, a digraph D with d vertices is equivalent to a subset of $[d] \times [d]$, i.e., a *board*. (Consider the edge set of D .) We call this subset the *associated board*, and conversely given a board we call the corresponding digraph on $[d]$ the *associated digraph*. Since the two representations are equivalent, we shall switch freely between them, sometimes without warning.

If G is a graph, we say that a partition σ of $V(G)$ is *connected* if for each block B of σ the subgraph induced by B is connected. The set of all connected partitions of G , partially ordered by refinement, forms a lattice L_G called the *lattice of contractions* or *bond lattice* of G .

Our notation for symmetric functions and partitions for the most part follows that of Macdonald [Mac], to which the reader is referred for any facts about symmetric functions that we do not explicitly reference. We shall always deal with symmetric functions in countably many variables. If τ is a set partition or an integer partition, we write $\ell(\tau)$ for the number of parts of τ , and $|\tau|$ for the sum of the sizes of the parts of τ . We define $\text{sgn } \tau$ by

$$\text{sgn } \tau \stackrel{\text{def}}{=} (-1)^{|\tau| - \ell(\tau)}.$$

We also define

$$r_\tau! \stackrel{\text{def}}{=} r_1! r_2! \cdots,$$

where r_i is the number of parts of τ of size i . We shall denote symmetric functions by a single letter such as g or by $g(x)$ if we wish to emphasize that the symmetric function is in the variables $x = (x_1, x_2, \dots)$. Similarly, the ring of symmetric functions will be denoted by Λ or $\Lambda(x)$. In addition to the usual symmetric functions m_λ , p_λ , e_λ , h_λ , and s_λ , we shall need the *augmented monomial symmetric functions* \tilde{m}_λ [D-K], which are defined by

$$\tilde{m}_\lambda \stackrel{\text{def}}{=} r_\lambda! m_\lambda.$$

We shall also need the *forgotten symmetric functions* f_λ , which are defined by

$$f_\lambda \stackrel{\text{def}}{=} (\text{sgn } \lambda) \omega(\tilde{m}_\lambda).$$

(Warning: this is one place where we deviate from Macdonald's conventions and follow Doubilet [Dou] instead, since [Dou] contains all the results about the forgotten symmetric functions that we shall need.) The symbol ω denotes the usual involution on symmetric functions that sends e_λ to h_λ . If $g(x)$ is a symmetric function, we shall write $g(-x)$ for the function obtained by negating each variable, and we shall write $g(1^n)$ for the polynomial in the variable n obtained by setting n variables equal to one and the rest equal to zero. We will sometimes use *set* partitions instead of *integer* partitions in subscripts; for example, if π is a set partition then the expression p_π is to be understood as an abbreviation for $p_{\text{type}(\pi)}$. We also say that a symmetric function g is *u-positive* if $\{u_\lambda\}$ is a symmetric function basis and the expansion of g in terms of this basis has nonnegative coefficients.

We shall be dealing frequently with functions in two sets of variables, i.e., elements of $\Lambda(x) \otimes \Lambda(y)$, so we fix some notation here. Let $\{x_1, x_2, x_3, \dots\}$ and $\{y_1, y_2, y_3, \dots\}$ be two sets of independent indeterminates. (Everything commutes

with everything else.) An expression like $g(x; y)$ indicates that g is invariant under any permutation of the x variables and any permutation of the y variables. If g is also invariant under permutations that mix x and y variables, then we will sometimes write $g(x, y)$ instead of $g(x; y)$. For example, $p_\lambda(x, y)$ indicates the power sum symmetric function in the union of the x and y variables. Expressions like $g(x; 0)$, $g(x; -y)$ and $g(1^i; 1^j)$ have their natural meanings. The notation $\omega_x g$ will indicate that for the purposes of applying ω , g is to be interpreted as a symmetric function in the x variables with coefficients in the y 's.

2. χ_G , X_G , $C(D)$, AND Ξ_D

We now turn to the definitions of the polynomial and symmetric function invariants that we shall be studying.

Let G be a graph. A *proper coloring* or simply a *coloring* of G is an assignment of positive integers to the vertices of G such that adjacent vertices are never assigned the same integer. The *chromatic polynomial* $\chi_G(i)$ of a graph G is defined to be the number of colorings of G using at most i colors (i.e., such that the set of assigned integers is a subset of $[i]$). To see that $\chi_G(i)$ is indeed a polynomial in i , let us define a *stable partition* of G to be a partition of $V(G)$ such that no two vertices in the same block are connected by an edge. Now observe that if we take any stable partition of G and assign integers in $[i]$ to the vertices of G in such a way that two vertices are assigned the same integer if and only if they are in the same block, then we obtain a coloring of G , and that moreover each coloring of G can be constructed in one and only one way by this procedure. Hence

$$\chi_G(i) = \sum_{\pi} i^{\ell(\pi)},$$

where the sum is over all stable partitions π of G . (We are using the notation $i^{\underline{k}} = i(i-1)\cdots(i-k+1)$ and $i^{\overline{k}} = i(i+1)\cdots(i+k-1)$.) In particular, $\chi_G(i)$ is a polynomial in i , as advertised.

Now let x_1, x_2, \dots be commuting independent indeterminates. Stanley [St2] defines the *chromatic symmetric function* X_G by

$$X_G = X_G(x) \stackrel{\text{def}}{=} \sum_{\pi} \tilde{m}_{\pi}(x),$$

where the sum is over all stable partitions of G . It is easy to see that $\tilde{m}_{\pi}(1^i) = i^{\ell(\pi)}$, whence $X_G(1^i) = \chi_G(i)$, so that X_G is a generalization of the chromatic polynomial. Stanley proves a number of results about X_G in the papers [St2][St4] (see also [Ga1][Ga2]). We shall be investigating some other properties of X_G later.

Now let D be a digraph. Following Chung and Graham [CG1], we say that a subset S of the edges of D is a *path-cycle cover* of D if no two elements of S lie in the same row or column of the associated board.¹ If we think of S as a spanning subgraph of D then we see that this condition just means that S is a (vertex-)disjoint union of directed paths and directed cycles. (Isolated vertices are thought of as directed paths with zero edges, and isolated loops are thought of as cycles of length one.) A path-cycle cover with no cycles is called a *path cover*, and a path-cycle cover with no paths is called a *cycle cover*. The *type* of a path-cycle cover S is the set partition of $V(D)$ such that each block is the set of vertices of one of these directed paths or directed cycles. We write $\pi(S)$ for the set of blocks corresponding to the directed paths and $\sigma(S)$ for the set of blocks corresponding to the directed cycles, and we say that the type of S is (π, σ) if $\pi(S) = \pi$ and $\sigma(S) = \sigma$. Chung and Graham's *cover polynomial* $C(D; i, j)$ is then defined by

$$C(D; i, j) \stackrel{\text{def}}{=} \sum_S i^{\ell(\pi(S))} j^{\ell(\sigma(S))},$$

where the sum is over all path-cycle covers $S \subset E(D)$.

In view of these definitions and the fact that $p_\sigma(1^j) = j^{\ell(\sigma)}$, the following definition (suggested by Stanley [St1]) is quite natural.

DEFINITION. Let D be a digraph, and let $x = \{x_1, x_2, \dots\}$ and $y = \{y_1, y_2, \dots\}$ be two sets of commuting independent indeterminates. The *path-cycle symmetric function* Ξ_D of D is defined by

$$\Xi_D = \Xi_D(x; y) \stackrel{\text{def}}{=} \sum_S \tilde{m}_{\pi(S)}(x) p_{\sigma(S)}(y),$$

where the sum is over all path-cycle covers $S \subset E(D)$.

The path-cycle symmetric function has not been investigated before, and the bulk of this thesis is devoted to its study.

Note that if we only care about *path* covers we can simply consider $\Xi_D(x; 0)$. In addition, if B is the board associated with D , then $\Xi_D(0; y)$ is equivalent to what Stanley and Stembridge call $Z[B]$ ([S-S, section 3]). Thus we may regard Ξ_D as a further generalization of Stanley and Stembridge's generalization of the theory of permutations with restricted position.

The following fact is immediate.

¹ Chung and Graham have a revised version [CG2] of the preprint [CG1]. We shall be referring to both versions.

PROPOSITION 1. $\Xi_D(1^i; 1^j) = C(D; i, j)$. ■

There is a close connection between X_G and Ξ_D . Given a poset P , let $G(P)$ denote its incomparability graph (in which two vertices of the poset are adjacent if and only if they are incomparable), and let $D(P)$ denote the digraph with edge set $\{(i, j) \mid i < j\}$. Chung and Graham observe that for any poset P ,

$$C(D(P); i, 0) = \chi_{G(P)}(i).$$

This connection generalizes readily to the symmetric function case.

PROPOSITION 2. *For any poset P , $\Xi_{D(P)} = X_{G(P)}$.*

Proof. Since $D(P)$ is acyclic, all path-cycle covers are in fact just path covers, so the y variables can be deleted from the definition of Ξ_D in this case. But path covers of $D(P)$ correspond to partitions of P into chains, which correspond to stable partitions of $G(P)$. Comparing the definitions of Ξ_D and X_G yields the proposition. ■

3. CONNECTION WITH ROOK THEORY

Let $B \subset [d] \times [d]$ be a board, and let the *rook number* r_k^B denote the number of ways of placing k non-taking rooks on B (i.e., the number of subsets of B such that no two squares lie in the same row or in the same column). Following Goldman, Joichi and White [GJW1], we define the *d -factorial polynomial* (or simply the *factorial polynomial*) $R(B; i)$ by

$$R(B; i) \stackrel{\text{def}}{=} \sum_k r_k^B i^{\underline{d-k}}.$$

If D is the digraph associated with B , we also write r_k^D for r_k^B and $R(D; i)$ for $R(B; i)$. (With this equivalence between boards and digraphs, the factorial polynomial is the same as Chung and Graham's *binomial drop polynomial*.) The study of the factorial polynomial and other rook polynomials is a well-established area of combinatorics (see for example [Rio, Chapters 7 and 8][GJW1][GJRW][GJW2][GJW3][GJW4]). The definition of a path-cycle cover already suggests a connection with rook theory. More precisely, we have the following proposition.

PROPOSITION 3. *For any digraph D , $R(D; i) = C(D; i, 1) = \Xi_D(1^i; 1)$.*

Proof. The first equality is demonstrated in [CG1] (or [CG2]) and the second equality follows from Proposition 1. ■

Proposition 3 (as well as, for example, [S-S, section 3] and [St2, Proposition 5.5]) suggests that some of the theory of rook polynomials might generalize to Ξ_D . This is indeed the case, as we shall see in more detail later.

Chapter 2

The Path-Cycle Symmetric Function

In this chapter, we carry out a fairly systematic investigation of the path-cycle symmetric function. For example, we try to derive as many analogues of theorems about X_G and generalizations of results from rook theory as we can. In addition, as we shall see, we uncover some unexpected results as a byproduct of this investigation.

1. BASIC FACTS

It is always good practice to begin the study of any mathematical object by computing a few examples. The chromatic symmetric functions of some graphs are computed in [St2], and Proposition 2 tells us that when these graphs are incomparability graphs, the chromatic symmetric functions are also path-cycle symmetric functions. But it would also be nice to see some examples of path-cycle symmetric functions that do not arise in this way.

In general, computing the path-cycle symmetric function is very hard (in the sense of computational complexity). Even the computation of the cover polynomial is $\#P$ -hard, since, as Chung and Graham observe, $C(D; 1, 0)$ is the number

of Hamiltonian paths, and computing this number is $\#P$ -hard (an indication of how this last fact may be proved can be found in section 7.3 of [G-J]). Furthermore, there are not many digraphs D for which Ξ_D has a simple closed form. For example, since Ξ_D is a generalization of Goldman, Joichi and White's factorial polynomial, for which there is a nice factorization theorem in the case of Ferrers shapes ([GJW1]), we might hope for a factorization theorem for Ξ_D . Unfortunately, this is not the case; the factorization theorem does not even generalize straightforwardly to the cover polynomial (although Dworkin [Dwo] has found some cases where the cover polynomial does factorize). Part of the problem is that unlike the factorial polynomial, Ξ_D (or even the cover polynomial) does not remain invariant under permutations of rows and columns. Furthermore, like X_G , the path-cycle symmetric function does not satisfy a deletion-contraction recurrence.

There are, however, some nice formulas for the path-cycle symmetric functions of directed paths and directed cycles.

PROPOSITION 4. *Let P_d be the directed path with d vertices and let C_d be the directed cycle with d vertices. Then*

$$\Xi_{P_d} = \sum_{\lambda \vdash d} \ell(\lambda)! m_\lambda(x) = \sum_{r=0}^{d-1} u_r s_{d-r, 1^r}(x)$$

and

$$\Xi_{C_d} = d \sum_{\lambda \vdash d} (\ell(\lambda) - 1)! m_\lambda(x) + p_d(y) = d \sum_{r=0}^{d-1} v_r s_{d-r, 1^r}(x) + p_d(y),$$

where u_r is the number of permutations of $r+1$ with no consecutive ascending pairs, and v_r is the number of permutations of r with no fixed points.

(Remark: the Schur function expansions were obtained with the aid of [Slo].)

Proof. Observe that every subset S of the edges of P_d is a path cover. Now $\pi(S)$ has type λ if and only if the complement of S is a set of edges whose removal from P_d breaks P_d into paths whose sizes are given by the parts of λ . So the number of such sets S is

$$\binom{\ell(\lambda)}{r_1, r_2, \dots},$$

where r_i is the number of parts of λ of size i . Hence

$$\Xi_{P_d} = \sum_{\lambda \vdash d} \binom{\ell(\lambda)}{r_1, r_2, \dots} \tilde{m}_\lambda = \sum_{\lambda \vdash d} \ell(\lambda)! m_\lambda.$$

To establish the Schur function expansion of Ξ_{P_d} , first note that

$$\sum_{r=0}^{k-1} u_r \binom{k-1}{r} = k!$$

(see [Kre] or [R-P] for a proof). Thus it suffices to show that the coefficient of m_λ in $s_{d-r,1^r}$ is

$$\binom{\ell(\lambda) - 1}{r}$$

whenever $\lambda \vdash d$, for then it will follow that

$$\sum_{r=0}^{d-1} u_r s_{d-r,1^r} = \sum_{\lambda \vdash d} m_\lambda \sum_{r=0}^{d-1} u_r \binom{\ell(\lambda) - 1}{r} = \sum_{\lambda \vdash d} \ell(\lambda)! m_\lambda,$$

since $d \geq \ell(\lambda)$ for all $\lambda \vdash d$.

The coefficient of m_λ in $s_{d-r,1^r}$ is just the number of column-strict Young tableaux of type $(d-r, 1^r)$ and content λ (see [Mac, (5.12)]). These Young tableaux are precisely those obtained by the following procedure: put a 1 in the corner of the tableau, and then choose any r distinct integers in the set $\{2, 3, \dots, \ell(\lambda)\}$ and put them in order down the left column of the tableau. Then fill in the rest of the first row with whatever integers are needed to make the content equal to λ (this can be done in exactly one way since the row must increase monotonically from left to right). Hence the number of such tableaux is

$$\binom{\ell(\lambda) - 1}{r},$$

as required.

To compute Ξ_{C_d} , observe that the set of *all* edges of C_d is a cycle cover (this accounts for the term $p_d(y)$) and that every other subset of edges is a path cover. The following procedure generates all path covers of C_d of type λ : take a set of directed paths whose sizes are given by the parts of λ , and make paths of the same length distinguishable; then arrange the paths to form a circle, and choose some vertex to be vertex 1. Now there are $(\ell(\lambda) - 1)!$ distinct circular permutations of the paths, and d ways to choose a vertex 1, but notice that this procedure generates each path cover $r_1!r_2!\cdots$ times because “in reality” directed paths of the same length are not distinguishable. This gives us the first formula for Ξ_{C_d} . To prove the other formula, all we need to establish (in view of the remarks above in the case of P_d) is that

$$\sum_{r=0}^{k-1} v_r \binom{k-1}{r} = (k-1)!,$$

but this is easy: an arbitrary permutation of $k-1$ letters can be chosen by first choosing $k-1-r$ fixed points and then choosing a permutation of the remaining r letters that has no fixed points. ■

We have found two formulas which allow the path-cycle symmetric function of a digraph to be computed from path-cycle symmetric functions of related digraphs.

One of these formulas is sufficiently interesting that we devote the entire next section to it. The other formula is a multiplicativity property that generalizes Chung and Graham’s multiplicativity formula for the cover polynomial [CG2, Corollary 2]. To prove this result, we introduce the concept of a path-cycle coloring, due to Chung and Graham (see [CG2]; the definition in [CG1] contains a minor error).

DEFINITION. A *path-cycle coloring* of a digraph D is an ordered pair (S, κ) where S is a path-cycle cover and κ is a map from $V(D)$ to the positive integers such that

1. $\kappa(v_1) = \kappa(v_2)$ if v_1 and v_2 belong to the same path or if v_1 and v_2 belong to the same cycle, and
2. $\kappa(v_1) \neq \kappa(v_2)$ if v_1 belongs to a path and v_2 belongs to a *different* path.

A *path coloring* is a path-cycle coloring with no cycles.

PROPOSITION 5. For any digraph D ,

$$\Xi_D = \sum_{(S, \kappa)} \prod_{u \text{ is in a path}} x_{\kappa(u)} \prod_{v \text{ is in a cycle}} y_{\kappa(v)},$$

where the sum is over all path-cycle colorings (S, κ) .

Proof. Regard the sum as a double sum: for each path-cycle cover S , sum over all “compatible” colorings κ , and then sum over all S . For each fixed S the paths and cycles may be colored independently so the sum over κ factors into a product of a symmetric function in x and a symmetric function in y . Clearly coloring the paths with distinct colors gives $\tilde{m}_{\pi(S)}(x)$ and coloring the cycles so that each cycle is monochromatic gives $p_{\sigma(S)}(y)$. ■

PROPOSITION 6. Suppose D is the digraph formed by joining the disjoint digraphs D_1 and D_2 with all the edges (v_1, v_2) with $v_1 \in V(D_1)$ and $v_2 \in V(D_2)$. Then $\Xi_D = \Xi_{D_1} \Xi_{D_2}$.

Proof. A path-cycle coloring of D induces path-cycle colorings of both D_1 and D_2 by restriction. Conversely, given any path-cycle coloring of D_1 and any path-cycle coloring of D_2 there exists a unique path-cycle coloring of D inducing them: if a path in D_1 has the same color as a path in D_2 , join them end to end with the appropriate edge from D_1 to D_2 . The result now follows from Proposition 5. ■

Notice that if $D_1 = D(P_1)$ and $D_2 = D(P_2)$ for some posets P_1 and P_2 , then the construction described in Proposition 6 corresponds to the *ordinal sum* $P_1 \oplus P_2$ (see [St3, Chapter 3]). Thus Proposition 6 can save us some labor in computing $\Xi_{D(P)}$ if P has a nontrivial ordinal sum decomposition. For example, it is clear from the definitions that if D has no edges at all, then $\Xi_D = \tilde{m}_{1^a} = d!e_d$. It follows that

$\Xi_{D(P)}$ is a multiple of an elementary symmetric function whenever P is an ordinal sum of antichains.

We now turn from the problem of computing Ξ_D to the problem of finding interesting facts about it. A natural thing to do is to try and find analogues of known theorems about X_G . This strategy is not always successful, but we do have the following analogues of Corollaries 2.7 and 2.11 of [St2].

PROPOSITION 7. *If D is an acyclic digraph, then $\omega_x \Xi_D$ is p -positive.*

Proof. Since D is acyclic, all path-cycle covers are path covers, and $\Xi_D = \Xi_D(x, 0)$. From Doubilet [Dou, Appendix 1] we know that for any set partition π ,

$$\tilde{m}_\pi = \sum_{\sigma \geq \pi} \mu(\pi, \sigma) p_\sigma.$$

Thus

$$\begin{aligned} \Xi_D &= \sum_S \sum_{\sigma \geq \pi(S)} \mu(\pi(S), \sigma) p_\sigma \\ &= \sum_\sigma \left(\sum_{\{S \mid \pi(S) \leq \sigma\}} \mu(\pi(S), \sigma) \right) p_\sigma, \end{aligned}$$

where S ranges over path covers. Now fix σ and let D_1, D_2, \dots, D_l be the subgraphs induced by the blocks of σ , with sizes d_1, d_2, \dots, d_l respectively. If c_i is the coefficient of p_{d_i} in Ξ_{D_i} , then we claim that the coefficient of p_σ in Ξ_D is $\prod_i c_i$. To see this, first note that choosing a path cover S of D such that $\pi(S) \leq \sigma$ is equivalent to (independently) choosing path covers for each D_i . If we let Π_n denote the lattice of partitions of $[n]$ ordered by refinement, then it is well known and easy to see that the interval $[\hat{0}, \sigma]$ in a partition lattice is isomorphic to

$$\Pi_{d_1} \times \Pi_{d_2} \times \cdots \times \Pi_{d_l}.$$

It is also well known (e.g., [St3, Prop. 3.8.2]) that the Möbius function of a product is the product of the Möbius functions. Putting these facts together readily yields our claim.

Thus to prove the theorem it suffices to prove that for an acyclic digraph with d vertices the sign of the coefficient of p_d is $(-1)^{d-1}$. For then, since any induced subgraph of an acyclic graph is acyclic, we can apply our claim above to show that the coefficient of p_σ is $\text{sgn } \sigma$.

Let D have d vertices. By specializing via Proposition 3, we see that the coefficient of p_d in Ξ_D equals the coefficient of i in $R(D; i)$. Directly from the definitions we see that this equals

$$(-1)^{d-1} \sum_{k=0}^{d-1} (-1)^k r_k^D (d-k-1)!$$

Now an acyclic digraph has at least one source and one sink, so by removing the corresponding row and column we see that B may be regarded as a subset of a $(d-1) \times (d-1)$ board. Then the above sum is a positive integer, by the inclusion-exclusion formula for rooks (see [St3, Theorem 2.3.1]). The sign of the coefficient is therefore $(-1)^{d-1}$ as desired. ■

We mention in passing that the argument used in the above proof allows [St2, Proposition 5.5] to be extended from posets to acyclic digraphs.

PROPOSITION 8. *If D is an acyclic digraph, then*

$$\frac{\partial \Xi_D}{\partial p_i} = \sum_H \left(\sum_S \mu(\pi(S), V(H)) \right) \Xi_{D \setminus H},$$

where H runs over all (vertex-)induced subgraphs of D with i vertices, S runs over all path covers of H , μ is the Möbius function of the lattice of partitions of $V(H)$, and $D \setminus H$ is the subgraph obtained by deleting $V(H)$ from D .

Proof. For the purposes of this proof only, we use the notation $S \sqsubseteq H$ to mean “ S is a path cover of H .” From the proof of Proposition 7, we have

$$\Xi_D = \sum_{\sigma} \left(\sum_{\{S \sqsubseteq D \mid \pi(S) \leq \sigma\}} \mu(\pi(S), \sigma) \right) p_{\sigma},$$


where σ ranges over all partitions of $V(D)$. Differentiating both sides with respect to p_i , we obtain

$$\frac{\partial \Xi_D}{\partial p_i} = \sum_{(H, \tau)} \left(\sum_{\{S \sqsubseteq D \mid \pi(S) \leq \tau \cup \{V(H)\}\}} \mu(\pi(S), \tau \cup \{V(H)\}) \right) p_{\tau},$$

where the outer sum ranges over all ordered pairs (H, τ) such that H is an induced subgraph of D with i vertices and τ is any partition of $V(D \setminus H)$. Since, as noted in the proof of Proposition 7, the Möbius function is multiplicative, it follows that

$$\begin{aligned} \frac{\partial \Xi_D}{\partial p_i} &= \sum_{(H, \tau)} \left(\sum_{S \sqsubseteq H} \mu(\pi(S), V(H)) \right) \left(\sum_{\{S \sqsubseteq D \setminus H \mid \pi(S) \leq \tau\}} \mu(\pi(S), \tau) \right) p_{\tau} \\ &= \sum_H \left(\sum_{S \sqsubseteq H} \mu(\pi(S), V(H)) \right) \sum_{\tau} \left(\sum_{\{S \sqsubseteq D \setminus H \mid \pi(S) \leq \tau\}} \mu(\pi(S), \tau) \right) p_{\tau} \\ &= \sum_H \left(\sum_{S \sqsubseteq H} \mu(\pi(S), V(H)) \right) \Xi_{D \setminus H}. \quad \blacksquare \end{aligned}$$

We conclude this section with two counterexamples. The digraphs



(taken from figure 2 of [Bon]) have identical lists of vertex-deleted subdigraphs but have different path-cycle symmetric functions (and even different cover polynomials), and hence the path-cycle symmetric function is not reconstructible (as opposed to the chromatic symmetric function—see section 3 of the next chapter).

The reader who is aware that every factorial polynomial is the chromatic polynomial of some graph ([GJW2]) might wonder if this generalizes to our symmetric function context. The answer is no, and the directed path on three vertices provides an example of an acyclic digraph D for which Ξ_D does not equal X_G for any graph G . The philosophical reason for this is that the proof that every factorial polynomial is a chromatic polynomial relies on the fact (mentioned previously) that the factorial polynomial of a board is independent of how the board is embedded in its $[d] \times [d]$ grid, but the same cannot be said of the path-cycle symmetric function.

2. RECIPROCITY

The *complement* D' of a digraph D is the digraph on the same vertex set whose edges are precisely those pairs (i, j) that are not edges of D . In this section we prove one of the most striking facts about the path-cycle symmetric function; namely, a combinatorial reciprocity theorem relating Ξ_D and $\Xi_{D'}$. We shall need two change-of-basis formulas, which we shall now state.

As in the proof of Proposition 7, let Π_n denote the lattice of partitions of $[n]$ (ordered by refinement). Recall that if $\pi \leq \sigma$ in Π_n and r_i is the number of blocks of σ that are composed of i blocks of π , then the Möbius function satisfies

$$|\mu(\pi, \sigma)| = \prod_i (i - 1)!^{r_i}.$$

(See [St3, Example 3.10.4] for a proof.) Also, following Doubilet [Dou], define

$$\lambda(\pi, \sigma)! \stackrel{\text{def}}{=} \prod_i i!^{r_i}.$$

We then have the following change-of-basis formulas (taken from [Dou, Appendix 1]).

PROPOSITION 9.

$$f_\pi = \sum_{\sigma \geq \pi} \lambda(\pi, \sigma)! \tilde{m}_\sigma = \sum_{\sigma \geq \pi} |\mu(\pi, \sigma)| p_\sigma. \blacksquare$$

We are now ready for the main theorem of this section.

THEOREM 1. *For any digraph D ,*

$$\Xi_D(x; y) = \sum_S \operatorname{sgn} \pi(S) f_{\pi(S)}(x, y) p_{\sigma(S)}(-y),$$

where the sum is over all path-cycle covers of the complement D' . Equivalently,

$$\Xi_D(x; y) = [\omega_x \Xi_{D'}(x; -y)]_{x \rightarrow (x, y)},$$

where $[g(x; y)]_{x \rightarrow (x, y)}$ means that, treating g as a symmetric function in the x 's with coefficients in the y 's, the set of x variables is to be replaced by the union of the x and y variables.

(Remark: “replacing the x variables with the union of the x and y variables” may be formalized as “applying the endomorphism Δ of $\Lambda(x) \otimes \Lambda(y)$ that sends $p_m(x)$ to $p_m(x) + p_m(y)$ and that leaves $p_n(y)$ fixed.”)

Proof. The equivalence of the two formulations is clear. We define a *partitioned order* of D to be a partition of $V(D)$ together with either a linear order or a cyclic order on each block. If κ is a partitioned order of D , let $\pi(\kappa)$ be the set of blocks with linear orders and let $\sigma(\kappa)$ be the set of blocks with cyclic orders. Let E_κ denote the set of ordered pairs (u, v) satisfying the following two conditions.

1. u and v are in the same block of κ and u immediately precedes v in the linear or cyclic order on the block.
2. (u, v) is not an edge of D .

Note that $E_\kappa \subset E(D')$ and that there is a natural bijection between partitioned orders κ such that $E_\kappa = \emptyset$ and path-cycle covers (given such a partitioned order, take all (u, v) satisfying condition 1 above). Now for any finite set T , the alternating sum

$$\sum_{S \subset T} (-1)^{|S|}$$

equals one if $T = \emptyset$ and is zero otherwise. Thus

$$\Xi_D = \sum_\kappa \tilde{m}_{\pi(\kappa)} p_{\sigma(\kappa)} \sum_{S \subset E_\kappa} (-1)^{|S|},$$

where the first sum is over all partitioned orders of D . We now interchange the order of summation. Observe first that all sets S that arise are path-cycle covers of D' , since S is a subset of the set of all (u, v) satisfying condition 1 above for some κ . Given a path-cycle cover S of D' , we now need to determine the set \mathcal{P} of partitioned orders of D that give rise to it. Only blocks with cyclic orders can give rise to cycles of S , so for every $\kappa \in \mathcal{P}$, $\sigma(\kappa)$ must include the blocks of $\sigma(S)$ among its own blocks. On the other hand, the blocks of $\pi(S)$ can arise either from blocks with linear orders or from blocks with cyclic orders. To determine all possibilities we must consider all ways of agglomerating the blocks of $\pi(S)$ into blocks of $\pi(\kappa)$, and then for each composite block in each such agglomeration we must consider both linear and cyclic orders. The linear or cyclic order on the composite block can be viewed as a linear or cyclic order on the *blocks of $\pi(S)$* (instead of on the vertices), because the linear or cyclic order must induce the edges of S , i.e., if (u, v) is an edge of S then u must immediately precede v in the order dictated by κ , and therefore the vertices in each block of $\pi(S)$ are constrained to be consecutive and in a fixed order. Clearly, every such linear or cyclic order on the blocks gives rise to a unique $\kappa \in \mathcal{P}$. The number of ways to impose a linear order if there are i blocks is $i!$ and the number of ways to impose a cyclic order is $(i - 1)!$. Thus we can enumerate \mathcal{P} by summing over all divisions of the blocks of $\pi(S)$ into two groups α and β (linear and cyclic) and, for each such division, summing over all ways of grouping the blocks into composite blocks, weighted by a factorial factor. More precisely we have

$$\Xi_D = \sum_S (-1)^{|S|} p_{\sigma(S)}(y) \sum_{(\alpha, \beta)} \sum_{\substack{\gamma \geq \alpha \\ \delta \geq \beta}} \lambda(\alpha, \gamma)! |\mu(\beta, \delta)| \tilde{m}_\gamma(x) p_\delta(y),$$

where the first sum is over all path-cycle covers of D' . By Proposition 9, we have

$$\begin{aligned} \Xi_D &= \sum_S (-1)^{|S|} p_{\sigma(S)}(y) \sum_{(\alpha, \beta)} f_\alpha(x) f_\beta(y) \\ &= \sum_S (-1)^{|S|} p_{\sigma(S)}(y) \sum_{(\alpha, \beta)} (\text{sgn } \alpha)(\text{sgn } \beta) \omega_x \tilde{m}_\alpha(x) \omega_y \tilde{m}_\beta(y). \end{aligned}$$

Now the blocks of α and β correspond to the paths of S , so $|\alpha| - \ell(\alpha)$ is the number of edges of S in α , and similarly for β . Thus $(\text{sgn } \alpha)(\text{sgn } \beta)$ depends only on the total number of edges of S devoted to directed paths (namely, $|\pi(S)| - \ell(\pi(S))$) and does not depend on the particular choice of α or β . We have

$$\begin{aligned} \Xi_D &= \sum_S (-1)^{|S|} p_{\sigma(S)}(y) (-1)^{|\pi(S)| - \ell(\pi(S))} \sum_{(\alpha, \beta)} \omega_x \tilde{m}_\alpha(x) \omega_y \tilde{m}_\beta(y) \\ &= \sum_S (-1)^{|\sigma(S)|} p_{\sigma(S)}(y) \omega_x \omega_y \sum_{(\alpha, \beta)} \tilde{m}_\alpha(x) \tilde{m}_\beta(y). \end{aligned}$$

A moment's thought shows that the inner sum is $\tilde{m}_{\pi(S)}(x, y)$. Now

$$\begin{aligned} \omega_x \omega_y p_n(x, y) &= \omega_x \omega_y (p_n(x) + p_n(y)) = (-1)^{n-1} p_n(x) + (-1)^{n-1} p_n(y) \\ &= [\omega_x p_n(x)]_{x \rightarrow (x, y)}, \end{aligned}$$

and since ω and $x \rightarrow (x, y)$ are both homomorphisms,

$$\omega_x \omega_y g(x, y) = [\omega_x g(x)]_{x \rightarrow (x, y)}$$

for any symmetric function g . Finally, $(-1)^{|\sigma|} p_\sigma(y) = p_\sigma(-y)$, so we obtain

$$\Xi_D = \sum_S p_{\sigma(S)}(-y) [\omega_x \tilde{m}_{\pi(S)}(x)]_{x \rightarrow (x, y)}$$

as desired. ■

We remark that the appearance of ω in Theorem 1 is what leads us to call it a combinatorial reciprocity theorem.

Theorem 1 is a rather curious result in that it is not obvious that the operation

$$\iota : g(x; y) \rightarrow [\omega_x g(x; -y)]_{x \rightarrow (x, y)}$$

is an involution. One might even wonder if ι is really an involution on *all* of $\Lambda(x) \otimes \Lambda(y)$ or if it is only an involution when restricted to some subset of $\Lambda(x) \otimes \Lambda(y)$ that includes all path-cycle symmetric functions. It is easily verified, however, that

$$\iota(\iota(p_m(x)p_n(y))) = p_m(x)p_n(y)$$

for any m and n . Furthermore, ι is the composition of three operations, each of which is an endomorphism of $\Lambda(x) \otimes \Lambda(y)$. It follows that ι is an involution on all of $\Lambda(x) \otimes \Lambda(y)$.

On the other hand, one can still ask whether there is any way of reformulating Theorem 1 in a way that makes it more obvious that the operation involved is an involution. One way of doing this is to define

$$\hat{\Xi}_D(x; y) \stackrel{\text{def}}{=} \sum_S (-2)^{\ell(\sigma(S))} \tilde{m}_{\pi(S)}(x, y) p_{\sigma(S)}(y),$$

where the sum is still over all path-cycle covers S of D . Then we have the following result.

PROPOSITION 10. $\hat{\Xi}_{D'}(x; y) = \omega_x \hat{\Xi}_D(x; -y)$.

(Remark: I arrived at the mysterious-looking definition of $\hat{\Xi}_D$ by first specializing to the cover polynomial, finding a change of variables that made the reciprocity operation an obvious involution, and then generalizing back to the symmetric function case by using plethysm.)

Proof. Let κ be the endomorphism of $\Lambda(x) \otimes \Lambda(y)$ that sends $p_n(y)$ to $-2p_n(y)$ and that leaves $p_m(x)$ fixed. Then $\hat{\Xi}_D$ is obtained from Ξ_D by applying κ first and then

Δ . In view of Theorem 1, it suffices to prove that the following operations have the same effect on all elements of $\Lambda(x) \otimes \Lambda(y)$:

1. (Left-hand side.) Apply ι , then κ , then Δ .
2. (Right-hand side.) Apply κ , then apply Δ , then negate the y variables, and then apply ω_x .

Each of these operations is an endomorphism, so it suffices to check their behavior on $p_m(x)p_n(y)$. A straightforward computation shows that in both cases the result is

$$2p_n(y)(p_m(x) - p_m(y))(-1)^{m+n},$$

which proves the desired result. ■

While Proposition 10 has the advantage of involving an operation that is clearly an involution, it has the disadvantage that the definition of $\hat{\Xi}_D$ is mysterious. In particular, the -2 has no obvious combinatorial significance. For this reason, we regard $\hat{\Xi}_D$ as an artificial contrivance and we shall continue to use Ξ_D instead.

Theorem 1 readily yields several attractive corollaries. For example, by setting all the x variables equal to zero, we immediately obtain [S-S, Theorem 3.2]. More interestingly, we can obtain an affirmative answer to the question, raised by Chung and Graham [CG1, section 8(c)], of whether $C(D; i, j)$ determines $C(D'; i, j)$.

COROLLARY 1. *If D is a digraph with d vertices, then $C(D'; i, j) = (-1)^d C(D; -i - j, j)$.*

Proof. Let g be any symmetric function that is homogeneous of degree d and let $g^* = \omega g$. We claim that $g^*(1^i)$ is obtained by changing i to $-i$ in $g(1^i)$ and then multiplying by $(-1)^d$. To see this, first consider the case where $g = p_\lambda$ for some $\lambda \vdash d$. Then $g^* = (\text{sgn } \lambda)p_\lambda$ and hence

$$g^*(1^i) = (\text{sgn } \lambda)i^{\ell(\lambda)}.$$

On the other hand $g(1^i) = i^{\ell(\lambda)}$. Changing i to $-i$ and multiplying by $(-1)^d$ amounts to multiplying by $(-1)^{d-\ell(\lambda)} = \text{sgn } \lambda$, as required. The claim then follows by linearity.

Now $\tilde{m}_\pi(1^i, 1^j) = (i+j)^{\underline{\ell(\pi)}}$. Since $(\text{sgn } \pi)f_\pi = \omega\tilde{m}_\pi$, we have

$$(\text{sgn } \pi)f_\pi(1^i, 1^j) = (-1)^{|\pi|}(-i-j)^{\underline{\ell(\pi)}}.$$

Also, as noted before, $p_\sigma(-y) = (-1)^{|\sigma|}p_\sigma(y)$. Thus, specializing Theorem 1 via Proposition 1 yields

$$\begin{aligned} C(D; i, j) &= \sum_S (-1)^{|\sigma(S)|} (-1)^{|\pi(S)|} (-i-j)^{\underline{\ell(\pi(S))}} j^{\ell(\sigma(S))} \\ &= (-1)^d C(D'; -i-j, j). \blacksquare \end{aligned}$$

Corollary 1 can be proved directly using deletion-contraction techniques, and it has also been obtained independently by Gessel [Ge1]. We omit the details.

A further specialization of Theorem 1 gives a formula for rook polynomials; we defer this to the next section, where we consider rook theory in more detail.

COROLLARY 2. *For any digraph D ,*

$$\Xi_D(x; 0) = \omega_x \Xi_{D'}(x; 0).$$

Note the similarity between this result and Stanley's reciprocity theorem [St2, Theorem 4.2]. In fact, the two reciprocity theorems overlap, because of Proposition 2, so Corollary 2 gives a new interpretation of $\omega \Xi_{D(P)} = \omega X_{G(P)}$ when P is a poset.

Corollary 2 follows immediately from Theorem 1, but we shall give two other proofs because they illustrate connections with other known results. The first proof is due to Gessel [Ge1], and it derives Corollary 2 from a result of Carlitz, Scoville and Vaughan [CSV, Theorem 7.3]. We need some preliminaries. Given a digraph D with d vertices, let

$$A_D = \{a_1, a_2, \dots, a_d\}$$

be a set of commuting independent indeterminates, and define

$$\alpha_{D,n} = \sum_{i_1, i_2, \dots, i_n} a_{i_1} a_{i_2} \cdots a_{i_n},$$

where the sum is over all i_1, i_2, \dots, i_n such that $(a_{i_j}, a_{i_{j+1}})$ is an edge of D for all $j < n$. Similarly, let

$$\alpha'_{D,n} = \sum_{i_1, i_2, \dots, i_n} a_{i_1} a_{i_2} \cdots a_{i_n},$$

where this time the sum is over all i_1, i_2, \dots, i_n such that $(a_{i_j}, a_{i_{j+1}})$ is an edge of the complement D' for all $j < n$. With this notation, the result of Carlitz, Scoville and Vaughan is (essentially) the following.

PROPOSITION 11. *For any digraph D ,*

$$\sum_n (-1)^n \alpha'_{D,n} = \left(\sum_n \alpha_{D,n} \right)^{-1}. \blacksquare$$

We can now give Gessel's proof of Corollary 2.

First proof of Corollary 2. Let $\theta_{D,y}$ be the homomorphism from the ring of symmetric functions in the variables $y = \{y_1, y_2, \dots\}$ to the ring of formal power series

in A_D that sends the complete symmetric function $h_n(y)$ to $\alpha_{D,n}$. Similarly, let $\theta'_{D,y}$ be the homomorphism that sends $h_n(y)$ to $\alpha'_{D,n}$. From [Mac, (4.2)] we have

$$\prod_{i,j} \frac{1}{1-x_i y_j} = \sum_{\lambda} h_{\lambda}(y) m_{\lambda}(x).$$

Applying $\theta_{D,y}$ gives

$$\theta_{D,y} \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \sum_{\lambda} \alpha_{D,\lambda_1} \alpha_{D,\lambda_2} \cdots m_{\lambda}(x).$$

Now $\Xi_D(x, 0)$ is just the coefficient of $a_1 a_2 \cdots a_d$ in this expression, since this coefficient counts all path covers of type π exactly $r_{\pi}!$ times, and $r_{\pi}! m_{\pi}(x) = \tilde{m}_{\pi}(x)$. Similarly, $\Xi_{D'}(x, 0)$ is the coefficient of $a_1 a_2 \cdots a_d$ in

$$\theta'_{D,y} \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \sum_{\lambda} \alpha'_{D,\lambda_1} \alpha'_{D,\lambda_2} \cdots m_{\lambda}(x).$$

Thus it suffices to prove that

$$\omega_x \theta_{D,y} \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \theta'_{D,y} \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right).$$

Now from [Mac, (4.3)] we have

$$\left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \sum_{\lambda} s_{\lambda}(x) s_{\lambda}(y),$$

so

$$\omega_x \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \sum_{\lambda} s_{\lambda'}(x) s_{\lambda}(y) = \sum_{\lambda} s_{\lambda}(x) s_{\lambda'}(y) = \omega_y \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right).$$

Thus

$$\omega_x \theta_{D,y} \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \theta_{D,y} \omega_x \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right) = \theta_{D,y} \omega_y \left(\prod_{i,j} \frac{1}{1-x_i y_j} \right).$$

So it suffices to show that $\theta_{D,y} \omega_y = \theta'_{D,y}$. From [Mac, (2.6)] we have

$$\sum_n (-1)^n e_n(y) = \left(\sum_n h_n(y) \right)^{-1},$$

so applying $\theta_{D,y}$ and using Proposition 11 yields

$$\sum_n (-1)^n \theta_{D,y}(e_n(y)) = \left(\sum_n \alpha_{D,n} \right)^{-1} = \sum_n (-1)^n \alpha'_{D,n}.$$

Equating terms of the same degree, we see that

$$\theta_{D,y} \omega_y(h_n(y)) = \theta_{D,y}(e_n(y)) = \alpha'_{D,n} = \theta'_{D,y}(h_n(y)),$$

completing the proof. ■

Our second proof of Corollary 2 is similar to Stanley's proof of the reciprocity theorem for X_G . Following Gessel [Ge2] and Stanley [St2, section 3], we define a power series in the variables $x = \{x_1, x_2, \dots\}$ to be *quasi-symmetric* if the coefficients of

$$x_{i_1}^{r_1} x_{i_2}^{r_2} \cdots x_{i_k}^{r_k} \quad \text{and} \quad x_{j_1}^{r_1} x_{j_2}^{r_2} \cdots x_{j_k}^{r_k}$$

are equal whenever $i_1 < i_2 < \cdots < i_k$ and $j_1 < j_2 < \cdots < j_k$. For any subset S of $[d-1]$ define the *fundamental* quasi-symmetric function $Q_{S,d}(x)$ by

$$Q_{S,d}(x) = \sum_{\substack{i_1 \leq \cdots \leq i_d \\ i_j < i_{j+1} \text{ if } j \in S}} x_{i_1} x_{i_2} \cdots x_{i_d}.$$

If there is no danger of confusion, we will sometimes write Q_S for $Q_{S,d}$ for brevity.

We have the following expansion of $\Xi_D(x; 0)$ in terms of fundamental quasi-symmetric functions.

PROPOSITION 12. *If D is a digraph with vertex set $[d]$, then*

$$\Xi_D(x; 0) = \sum_{\pi \in S_d} Q_{S(\pi),d}(x),$$

where S_d is the group of permutations of $[d]$ and

$$S(\pi) = \{i \in [d] \mid (\pi_i, \pi_{i+1}) \text{ is not an edge of } D\}.$$

Proof. We use the expression for Ξ_D given in Proposition 5. Given a path coloring of D , arrange the paths in increasing order of their colors, and within each path arrange the vertices in the order given by the directed path. This gives a permutation of the vertices of D , and it is easy to see that $Q_{S(\pi)}(x)$ counts precisely the path colorings that give rise to π . ■

We can now give our second proof of Corollary 2.

Second proof of Corollary 2. Without loss of generality we may assume that the vertex set of D is $[d]$. From the same argument as in Proposition 12, we see that

$$\Xi_{D'}(x; 0) = \sum_{\pi \in S_d} Q_{[d] \setminus S(\pi)}(x).$$

In view of Proposition 12, it suffices to show that the map that sends Q_S to $Q_{[d] \setminus S}$ equals ω when restricted to symmetric functions. A proof of this fact may be found in the proof of [St2, Theorem 4.2]. ■

Stanley [St2] has obtained an analogue of Proposition 12 by using the theory of acyclic orientations and P -partitions. It is natural to ask if these ideas can be applied to studying Ξ_D . Unfortunately this does not seem possible. For example, a key step in the proof of the analogue of Proposition 12 involves expressing X_G as a sum of certain poset generating functions, but in general Ξ_D has no such expression, even if D is acyclic.

3. ROOK THEORY

As we explained in the previous chapter, there is a close connection between the path-cycle symmetric function and rook theory, because Ξ_D is a generalization of the factorial polynomial of Goldman, Joichi and White. Thus every theorem about Ξ_D can be specialized to a theorem in rook theory, and we can also try to generalize every theorem in rook theory to a theorem about Ξ_D .

A good example of this relationship is Theorem 1, which can be viewed as a generalization of a result in Riordan [Rio, Chapter 7, Theorem 2] relating the rook numbers of complementary boards, a result which we now state. If B is a board, we let $B' = ([d] \times [d]) \setminus B$ denote the complementary board.

PROPOSITION 13. *Let $B \subset [d] \times [d]$ be a board. Then $R(B'; i) = (-1)^d R(B; -i - 1)$.*

Proof. Let D be the associated digraph. From Corollary 1 and Proposition 3 we have

$$R(B'; i) = C(D'; i, 1) = (-1)^d C(D; -i - 1, 1) = (-1)^d R(B; -i - 1). \blacksquare$$

Riordan's original result is

$$\sum_k r_k^B (d - k)! i^k = \sum_k (-1)^k r_k^{B'} (d - k)! i^k (i + 1)^{d - k},$$

which can be shown to be equivalent to Proposition 13. However, it seems that the formulation of Proposition 13, which is much simpler, has not appeared in the

literature before (although as Gessel [Ge1] has observed, it follows immediately from the next proposition below).

Goldman [Gol] has remarked that Proposition 13 ought to have a direct combinatorial proof, and indeed it does. Begin by observing that

$$\begin{aligned} (-1)^d R(B'; -i-1) &= (-1)^d \sum_{k=0}^d r_k^{B'} (-i-1)^{d-k} \\ &= \sum_{k=0}^d (-1)^k r_k^{B'} (i+d-k)^{d-k}. \end{aligned}$$

Now add i extra rows to $[d] \times [d]$. Then $r_k^{B'} (i+d-k)^{d-k}$ is the number of ways of first placing k rooks on B' and then placing $d-k$ more rooks anywhere (i.e., on B , B' or on the extra rows) such that no two rooks can take each other in the final configuration. By a straightforward inclusion-exclusion argument, we see that the resulting configurations in which the set S of rooks B' is nonempty cancel out of the above sum, because they are counted once for each subset of S , with alternating signs. Thus what survives is the set of placements of d nontaking rooks on the extended board such that no rook lies on B' —but it is easy to see that is this precisely what $R(B; i)$ counts.

We should remark that the idea of adding i extra rows was taken from [GJW1]. Also, Gessel [Ge1] has extended Proposition 13 to the more general setting of simplicial complexes, and the above combinatorial proof can be adapted without much difficulty to proving Gessel's more general result. We omit the details since they are tangential to our main purpose. Finally, we remark that it ought to be possible to extend the above combinatorial proof to the cover polynomial, but we have not done so.

We now turn to what is perhaps the most fundamental theorem of rook theory—the inclusion-exclusion formula (which we have already mentioned in the proof of Proposition 7)—and generalize it to the context of the path-cycle symmetric function. The inclusion-exclusion formula has many equivalent formulations; the one we shall find most convenient is following one, whose proof is given implicitly in [CG1].

PROPOSITION 14. *Let D be a digraph with d vertices, and let N_k^D denote the number of ways of placing d non-taking rooks on $[d] \times [d]$ such that exactly k rooks lie on the board associated with D . Then*

$$R(D; i) = \sum_k N_k^D \binom{i+k}{d}. \blacksquare$$

To state our generalization we need a few more definitions.

DEFINITION. For any pair of integer partitions λ and μ , define $D_{\lambda,\mu}$ to be a disjoint union of directed paths and directed cycles such that the i th directed path has λ_i vertices and the j th directed cycle has μ_j vertices. Define

$$\tilde{\Xi}_{\lambda,\mu} \stackrel{\text{def}}{=} \sum_S \frac{\tilde{m}_{\pi(S)}(x) p_{\sigma(S)}(y)}{\ell(\pi(S))!},$$

where the sum is over all path-cycle covers S of $D_{\lambda,\mu}$. For brevity we shall write D_λ for $D_{\lambda,\emptyset}$ and $\tilde{\Xi}_\lambda$ for $\tilde{\Xi}_{\lambda,\emptyset}$. We then have the following result.

THEOREM 2. Let D be a digraph with d vertices and let B be the associated board. Let $\mathcal{N}_{\lambda,\mu}^D$ be the set of placements of d non-taking rooks on $[d] \times [d]$ such that the type (π, σ) of the path-cycle cover formed by the set of edges corresponding to rooks placed on B satisfies $\text{type}(\pi) = \lambda$ and $\text{type}(\sigma) = \mu$, and let $N_{\lambda,\mu}^D = |\mathcal{N}_{\lambda,\mu}^D|$. Then

$$\Xi_D = \sum_{\lambda,\mu} N_{\lambda,\mu}^D \tilde{\Xi}_{\lambda,\mu},$$

where the sum is over all integer partitions λ and μ .

Proof. The proof is similar to the proof of [St3, Theorem 2.3.1]. Given any two integer partitions ν and η , let $\mathcal{R}_{\nu,\eta}^D$ be the set of path-cycle covers S of D satisfying $\text{type}(\pi(S)) = \nu$ and $\text{type}(\sigma(S)) = \eta$. Note that every element of $\mathcal{R}_{\nu,\eta}^D$ has $\ell(\nu)$ directed paths plus some cycles and therefore has a total of $d - \ell(\nu)$ edges.

Now fix a pair of integer partitions ν and η and consider the set of pairs (S, T) such that $S \in \mathcal{R}_{\nu,\eta}^D$ and T is an extension of S (regarded as a placement of non-taking rooks on B) to a placement of d non-taking rooks on $[d] \times [d]$. This set has $\ell(\nu)! |\mathcal{R}_{\nu,\eta}^D|$ elements, since for each $S \in \mathcal{R}_{\nu,\eta}^D$ the $\ell(\nu)$ rows and columns unoccupied by S can support $\ell(\nu)!$ placements of non-taking rooks. On the other hand, we can enumerate the set in another way, by taking each placement of d non-taking rooks on $[d] \times [d]$ and counting how many $S \in \mathcal{R}_{\nu,\eta}^D$ it extends. Now if $T \in \mathcal{N}_{\lambda,\mu}^D$ for some λ and μ , then the number of elements of $\mathcal{R}_{\nu,\eta}^D$ that it extends is just the number $n_{\lambda,\mu,\nu,\eta}$ of path-cycle covers S of $D_{\lambda,\mu}$ satisfying $\text{type}(\pi(S)) = \nu$ and $\text{type}(\sigma(S)) = \eta$. Since every placement of d non-taking rooks on $[d] \times [d]$ belongs to $\mathcal{N}_{\lambda,\mu}^D$ for some λ and μ , we have

$$\sum_{\lambda,\mu} N_{\lambda,\mu}^D n_{\lambda,\mu,\nu,\eta} = \ell(\nu)! |\mathcal{R}_{\nu,\eta}^D|.$$

Now divide both sides by $\ell(\nu)!$, multiply both sides by $\tilde{m}_\nu(x) p_\eta(y)$, and sum over all ν and η to obtain the desired result. ■

The next proposition shows that Theorem 2 does indeed generalize Proposition 14, a fact which may not be obvious at first glance.

PROPOSITION 15. *For any integer partitions λ and μ ,*

$$\tilde{\Xi}_{\lambda,\mu}(1^i; 1) = \binom{i + d - \ell(\lambda)}{d},$$

where $d = |\lambda| + |\mu|$.

Proof. Directly from the definitions we have

$$\tilde{\Xi}_{\lambda,\mu}(1^i; 1) = \sum_S \frac{i^{\ell(\pi(S))}}{\ell(\pi(S))!} = \sum_S \binom{i}{\ell(\pi(S))},$$

where the sum is over all path-cycle covers of $D_{\lambda,\mu}$. The sum can be broken up into a double sum:

$$\tilde{\Xi}_{\lambda,\mu}(1^i; 1) = \sum_k \sum_{\{S \mid \ell(\pi(S))=k\}} \binom{i}{k}.$$

But $\ell(\pi(S)) = k$ if and only if $|S| = d - k$. Since every subset of the edges of $D_{\lambda,\mu}$ is a path-cycle cover, and since the total number of edges of $D_{\lambda,\mu}$ is $d - \ell(\lambda)$, we have

$$\tilde{\Xi}_{\lambda,\mu}(1^i; 1) = \sum_k \binom{d - \ell(\lambda)}{d - k} \binom{i}{k} = \binom{i + d - \ell(\lambda)}{d}. \blacksquare$$

In view of Proposition 3, Proposition 15, and the fact that the number of edges of a path-cycle cover of type (π, σ) is $d - \ell(\pi)$, we see that Theorem 2 implies Proposition 14.

We remark that in passing that Gessel [Ge1] has obtained a generalization of Proposition 14 for the cover polynomial that does not appear to follow from our results.

It might seem that Theorem 2 is contrived, since we seem to have defined $\tilde{\Xi}_{\lambda,\mu}$ just so that Theorem 2 would come out right. In fact, however, the functions $\tilde{\Xi}_{\lambda,\mu}$ are surprisingly interesting objects in their own right. For a start, we have the following easy fact.

PROPOSITION 16. *The functions $\tilde{\Xi}_\lambda$ form a linear basis for the ring of symmetric functions over the rationals, and the functions $\tilde{\Xi}_{\lambda,\mu}$ form a linear basis for the ring of symmetric functions in two sets of variables (again over the rationals).*

Proof. Write

$$\tilde{\Xi}_\lambda = \sum_\mu c_{\lambda,\mu} \tilde{m}_\mu.$$

Then it is clear from the definition of $\tilde{\Xi}_\lambda$ that $c_{\lambda,\lambda} \neq 0$ and $c_{\lambda,\mu} \neq 0$ only if $\lambda \geq \mu$ in refinement order. Thus the matrix $(c_{\lambda,\mu})$ with respect to any linear extension of refinement order is triangular with nonzero entries on the diagonal. This proves the first assertion. To prove the second assertion, define a partial order on pairs of integer partitions by setting $(\lambda, \mu) < (\nu, \eta)$ if the multiset of parts of η can be partitioned into two multisets α and β such that $\mu = \beta$ and λ is a refinement of $\nu \cup \alpha$. The same kind of reasoning as before, with this partial order in place of refinement order and with the basis $\tilde{m}_\lambda(x) p_\mu(y)$ in place of \tilde{m}_λ , can then be applied to prove the second assertion. ■

The $\tilde{\Xi}_\lambda$ turn out to be particularly interesting, as we shall see presently. The following table expresses $\tilde{\Xi}_\lambda$ in terms of the monomial symmetric functions for some small values of λ .

$$\begin{aligned} \begin{bmatrix} \tilde{\Xi}_2 \\ \tilde{\Xi}_{11} \end{bmatrix} &= \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} m_2 \\ m_{11} \end{bmatrix} \\ \begin{bmatrix} \tilde{\Xi}_3 \\ \tilde{\Xi}_{21} \\ \tilde{\Xi}_{111} \end{bmatrix} &= \begin{bmatrix} 1 & 1 & 1 \\ 0 & \frac{1}{2} & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} m_3 \\ m_{21} \\ m_{111} \end{bmatrix} \\ \begin{bmatrix} \tilde{\Xi}_4 \\ \tilde{\Xi}_{31} \\ \tilde{\Xi}_{22} \\ \tilde{\Xi}_{211} \\ \tilde{\Xi}_{1111} \end{bmatrix} &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & \frac{1}{2} & 0 & \frac{2}{3} & 1 \\ 0 & 0 & 1 & \frac{2}{3} & 1 \\ 0 & 0 & 0 & \frac{1}{3} & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} m_4 \\ m_{31} \\ m_{22} \\ m_{211} \\ m_{1111} \end{bmatrix} \\ \begin{bmatrix} \tilde{\Xi}_5 \\ \tilde{\Xi}_{41} \\ \tilde{\Xi}_{32} \\ \tilde{\Xi}_{311} \\ \tilde{\Xi}_{221} \\ \tilde{\Xi}_{2111} \\ \tilde{\Xi}_{11111} \end{bmatrix} &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & \frac{1}{2} & 0 & \frac{2}{3} & \frac{1}{3} & \frac{3}{4} & 1 \\ 0 & 0 & \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{3}{4} & 1 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{2} & 1 \\ 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{2} & 1 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} m_5 \\ m_{41} \\ m_{32} \\ m_{311} \\ m_{221} \\ m_{2111} \\ m_{11111} \end{bmatrix} \end{aligned}$$

While a casual inspection of the above table may not reveal any interesting patterns, we have the following surprising result.

THEOREM 3. *The linear map that sends $\tilde{\Xi}_\lambda$ to $(\text{sgn } \lambda) \tilde{m}_\lambda / \ell(\lambda)!$ is an involution.*

Proof. Given any integer partitions λ and μ , let π be any set partition of type λ and define

$$c_{\lambda,\mu} \stackrel{\text{def}}{=} \sum_{\{\sigma \geq \pi \mid \text{type}(\sigma) = \mu\}} \lambda(\pi, \sigma)!$$

Note that $c_{\lambda,\mu}$ does not depend on the choice of π . We claim that

$$\Xi_{D_\mu} = \sum_{\lambda} \frac{r_\mu!}{r_\lambda!} c_{\lambda,\mu} \tilde{m}_\lambda.$$

To see this, first consider the case where $r_\mu! = r_\lambda! = 1$, i.e., the case of distinct parts. We have a disjoint union D_μ of directed paths and we want to count the number of path covers of type λ . In a path cover of D_μ , each directed path is broken up into a sequence of smaller directed paths. So the path covers can be enumerated as follows: take a set partition π of type λ and consider all ways of grouping its blocks into a partition σ of type μ and then linearly ordering the blocks of π within each block of σ . Such a configuration determines a path cover: for any block b of σ , the sequence of blocks of π in b dictates the sizes of the sequence of smaller directed paths composing the directed path in D_μ corresponding to b . It is easy to see that this correspondence is bijective, and this proves our claim in the case of distinct parts. For the general case, observe that we want equal-sized parts of π to be indistinguishable and equal-sized parts of σ to be distinguishable, so we must multiply by $r_\mu!/r_\lambda!$.

From Proposition 9 we see that the matrix $((\text{sgn } \lambda) c_{\lambda,\mu})$ is the matrix of ω (with respect to the augmented monomial symmetric function basis) and is therefore an involution. From our claim it follows that the matrix relating $(\text{sgn } \lambda) \tilde{m}_\lambda / r_\lambda!$ and $\Xi_{D_\mu} / r_\mu!$ or equivalently the matrix relating

$$\frac{(\text{sgn } \lambda) \tilde{m}_\lambda}{r_\lambda! \ell(\lambda)!} \quad \text{and} \quad \frac{\tilde{\Xi}_\mu}{r_\mu!}$$

is an involution. But then the desired result follows, since the factors of $r_\lambda!$ and $r_\mu!$ amount to conjugating by a (diagonal) matrix, and this does not change the involution property. ■

Notice the close connection between the involution of Theorem 3 and the involution ω . (In fact, it was a suggestion by Stanley that the two involutions might be equal that led to the proof of Theorem 3.) The two involutions are not the same, however—the former is not even a homomorphism—and philosophically speaking it is still unclear why they are related. For example, if we specialize to the polynomial level, ω becomes (essentially) the operation of substituting $-x$ for x , but we do not know of any such intuitive interpretation for the involution of Theorem 3.

The $\tilde{\Xi}_\lambda$ are also closely related to the fundamental quasi-symmetric functions Q_S defined in the previous section. (The knowledgeable reader may already have suspected this since the fundamental quasi-symmetric functions specialize to the same polynomial basis as the $\tilde{\Xi}_\lambda$ do.) If S is a subset of $[d-1]$, then we define the *type* of S to be the integer partition whose parts are the lengths of the subwords obtained by breaking the word $123\dots d$ after each element of S .

THEOREM 4. Let g be any symmetric function. If a_λ and b_S are constants such that

$$g = \sum_{\lambda} a_{\lambda} \tilde{\Xi}_{\lambda} \quad \text{and} \quad g = \sum_S b_S Q_S,$$

then

$$a_{\lambda} = \sum_{\{S | \text{type}(S) = \lambda\}} b_S.$$

Proof. It is not difficult to see that it suffices to prove the theorem for the case $g = \tilde{\Xi}_{\mu}$. Let d be any positive integer and let S be any subset of $[d - 1]$ and define

$$\tilde{Q}_S \stackrel{\text{def}}{=} \sum_{\substack{i_1 \leq i_2 \leq \dots \leq i_d \\ i_j < i_{j+1} \text{ iff } j \in S}} x_{i_1} x_{i_2} \cdots x_{i_d}.$$

Then

$$m_{\lambda} = \sum_{\{S | \text{type}(S) = \lambda\}} \tilde{Q}_S \quad \text{and} \quad Q_S = \sum_{T \supset S} \tilde{Q}_T.$$

By an inclusion-exclusion argument,

$$m_{\lambda} = \sum_{\{S | \text{type}(S) = \lambda\}} \sum_{T \supset S} (-1)^{|T| - |S|} Q_T.$$

Let $q_{\lambda T}$ be the coefficient of Q_T in m_{λ} . We compute

$$\sum_{\{T | \text{type}(T) = \nu\}} q_{\lambda T}.$$

Observe that there is a bijection between subsets of type λ and orderings of the parts of λ : given a subset $S \subset [d - 1]$ of type λ , take the sequence of the lengths of the subwords of the word $123 \dots d$ obtained by breaking after each element of S . Thinking of such subwords as directed paths, we see that for any fixed S of type λ , the number of subsets $T \supset S$ such that $\text{type}(T) = \nu$ is just the number of path-coverings of D_{λ} of type ν , which from the proof of Theorem 3 is

$$\frac{r_{\lambda}!}{r_{\nu}!} c_{\nu \lambda}$$

(using the notation of Theorem 3). Now there are $\ell(\lambda)!/r_{\lambda}!$ subsets S of type λ , and if $\text{type}(S) = \lambda$ and $\text{type}(T) = \nu$ then

$$(-1)^{|T| - |S|} = (\text{sgn } \nu)(\text{sgn } \lambda).$$

Putting all this together, we see that

$$\sum_{\{T | \text{type}(T) = \nu\}} q_{\lambda T} = \frac{\ell(\lambda)!}{r_{\nu}!} c_{\nu \lambda} (\text{sgn } \nu)(\text{sgn } \lambda).$$

But, again from the proof of Theorem 3,

$$\tilde{\Xi}_\mu = \sum_\lambda \frac{r_\mu!}{r_\lambda!} c_{\lambda\mu} \frac{r_\lambda!}{\ell(\lambda)!} m_\lambda.$$

Hence if $g = \tilde{\Xi}_\mu$, then

$$\begin{aligned} \sum_{\{S \mid \text{type}(S) = \nu\}} b_S &= \sum_\lambda \frac{r_\mu!}{r_\lambda!} c_{\lambda\mu} \frac{r_\lambda!}{\ell(\lambda)!} \cdot \frac{\ell(\lambda)!}{r_\nu!} c_{\nu\lambda} (\text{sgn } \nu) (\text{sgn } \lambda) \\ &= \frac{r_\mu!}{r_\nu!} \sum_\lambda (\text{sgn } \nu) c_{\nu\lambda} (\text{sgn } \lambda) c_{\lambda\mu} \\ &= \delta_{\mu\nu}, \end{aligned}$$

because $((\text{sgn } \lambda) c_{\lambda\mu})$ is the matrix of ω with respect to the augmented monomial symmetric function basis (by Proposition 9), and ω is an involution. This completes the proof. ■

COROLLARY 3. *If g is a Q -positive symmetric function, then g is $\tilde{\Xi}$ -positive. In particular, X_G and the Schur functions are $\tilde{\Xi}$ -positive.*

Proof. The first assertion is immediate from Theorem 4. The fact that X_G and the Schur functions are Q -positive is “folklore”; it is implicit in [St5], but see also [Ge2] and [St2]. ■

We caution the reader not to read more into Theorem 4 than is actually there! For example, $\tilde{\Xi}_\lambda$ is *not* Q -positive. Nor is it true that the only Q_T 's in the Q -expansion of $\tilde{\Xi}_\lambda$ with nonzero coefficients are those with $\text{type}(T) = \lambda$. Thus, while Theorem 4 allows one to translate combinatorial *interpretations of the coefficients* of the Q -expansion of a symmetric function g into combinatorial interpretations of the coefficients of the $\tilde{\Xi}$ -expansion of g , there is no guarantee that combinatorial *proofs* can be so translated. Some tricky reshuffling of combinatorial information occurs in the transition from the Q 's to the $\tilde{\Xi}$'s.

We should mention another, somewhat more philosophical, reason that the $\tilde{\Xi}_\lambda$ are interesting. The bases \tilde{m}_λ , p_λ , e_λ , h_λ , s_λ , and f_λ occur frequently “in nature.” Similarly, there are certain “natural” bases for polynomials, and moreover there is a correspondence between some of the symmetric function bases and the polynomial bases given by $g \mapsto g(1^i)$, e.g., p_λ corresponds to i^n , and the reciprocally related bases \tilde{m}_λ and f_λ correspond to $i^{\bar{n}}$ and $i^{\bar{n}}$. However, so far as we are aware, no symmetric function counterpart to the polynomial basis

$$\binom{i+n}{d}_{n=0,1,\dots,d}$$

has been proposed before. Proposition 15, together with Theorem 3 and Corollary 3, suggests that the $\tilde{\Xi}_\lambda$ may be the “right” symmetric function generalization of this polynomial basis.

If this is the case, one might hope that specializing $\tilde{\Xi}_{\lambda,\mu}$ might give rise to an interesting basis for polynomials in two variables. Unfortunately, this does not seem to be true. However, specializing $\tilde{\Xi}_{\lambda,\mu}$ to a polynomial does provide a simple proof of a theorem of Chung and Graham whose original proof is quite complicated. Following Chung and Graham, for any placement T of d non-taking rooks on $[d] \times [d]$, let $\text{drop}(T)$ be the subgraph of D with edges corresponding to the squares occupied by the rooks of T . If D is a digraph with d vertices, let $\delta_D(q, r, s)$ be the number of ordered pairs (S, T) such that S is a set of r edges of D forming precisely s disjoint cycles and T is a placement of d non-taking rooks on $[d] \times [d]$ with $S \subset \text{drop}(T)$ and $|\text{drop}(T)| = q + r$. Chung and Graham’s result ([CG1, Theorem 2] or [CG2, Theorem 3]) is then the following.

PROPOSITION 17. *For any digraph D with d vertices,*

$$C(D; i, j) = \sum_{q,r,s} \delta_D(q, r, s) \binom{i+q}{d-r} (j-1)^s.$$

Proof. We can restate the desired result as

$$\begin{aligned} C(D; i, j+1) &= \sum_{q,r,s,t} \delta_D(q, r, s) \binom{q}{d-r-t} \binom{i}{t} j^s \\ &= \sum_{q,r,s,t} \delta_D(q, r, s) \binom{q}{t-d+r+q} \binom{i}{t} j^s. \end{aligned}$$

From Theorem 2 and the definition of $\tilde{\Xi}_{\lambda,\mu}$ we have

$$C(D; i, j+1) = \sum_{\lambda,\mu,t,u} N_{\lambda,\mu}^D n_{\lambda,\mu,t,u} \binom{i}{t} (j+1)^u = \sum_{\lambda,\mu,s,t,u} N_{\lambda,\mu}^D n_{\lambda,\mu,t,u} \binom{u}{s} \binom{i}{t} j^s,$$

where $n_{\lambda,\mu,t,u}$ is the number of path-cycle covers of $D_{\lambda,\mu}$ with t paths and u cycles. Now $\binom{i}{t} j^s$ is a basis for polynomials in two variables, so equating coefficients we see that we just need to prove that for any fixed s and t ,

$$\sum_{\lambda,\mu,u} N_{\lambda,\mu}^D n_{\lambda,\mu,t,u} \binom{u}{s} = \sum_{q,r} \delta_D(q, r, s) \binom{q}{t-d+r+q}.$$

We can think of both sides as counting placements of d non-taking rooks on $[d] \times [d]$ with certain multiplicities. On the left-hand side, the number of times each such

placement T is counted equals the number of path-cycle covers of $\text{drop}(T)$ with exactly t paths plus some number of cycles of which s are distinguished. As for the right-hand side, we can rewrite it as

$$\sum_{e,r} \delta_D(e-r, r, s) \binom{e-r}{t-d+e}.$$

Then for any placement T , only one value of e (namely, $e = |\text{drop}(T)|$) involves T . Thus if we let $e = |\text{drop}(T)|$, the number of times T is counted is

$$\sum_r \left(\begin{array}{l} \text{the number of ways of choosing} \\ s \text{ cycles of } \text{drop}(T) \text{ with } r \text{ edges} \end{array} \right) \cdot \binom{e-r}{t-d+e},$$

which is just the number of ways of choosing s cycles and then deleting $t - (d - e)$ of the remaining edges (i.e., creating $t - (d - e)$ new paths). But $d - e$ is the original number of paths in $\text{drop}(T)$, so this results in a total of exactly t paths. The proposition follows. ■

Let us now return from this digression to the problem of generalizing rook theory to the context of the path-cycle symmetric function. Our next result generalizes a Möbius inversion formula for factorial polynomials due to Goldman, Joichi and White [GJW4]. For simplicity we consider only the case of acyclic digraphs, although the generalization to arbitrary digraphs is straightforward. So suppose D is an acyclic digraph with d vertices and let B be its associated board. Following an idea of Goldman, Joichi and White, extend the columns of $[d] \times [d]$ infinitely downwards, so that there are now infinitely many rows. Let \mathcal{S} be the set of all placements of d rooks such that

1. every rook lies either on B or one of the appended squares, and
2. no two rooks lie in the same column.

Given $S \in \mathcal{S}$, define $\pi(S)$ to be the partition of $[d]$ in which two numbers i and j lie in the same block if and only if the rooks in columns i and j lie in the same row. To each $S \in \mathcal{S}$ we also associate a coloring of $[d]$ as follows. Color the vertex $i \in [d]$ with color j if the rook in column i lies in the j th *appended* row. Otherwise, if the rook in column i lies in the j th *original* row, make vertex i the same color as vertex j . Since there is exactly one rook in each column, and since D is acyclic, these rules give a well-defined coloring c_S . For every set partition of $[d]$, define

$$T_\pi^D = T_\pi^B \stackrel{\text{def}}{=} \sum_{\{S \in \mathcal{S} \mid \pi(S) = \pi\}} x^S,$$

where

$$x^S \stackrel{\text{def}}{=} \prod_{i \in [d]} x_{c_S(i)}.$$

Finally define

$$T_{\geq \pi}^D = T_{\geq \pi}^B \stackrel{\text{def}}{=} \sum_{\sigma \geq \pi} T_{\sigma}^B.$$

PROPOSITION 18. *For any acyclic digraph D with d vertices,*

$$\Xi_D = \sum_{\pi \in \Pi_d} \mu(\hat{0}, \pi) T_{\geq \pi}^D.$$

Proof. By Möbius inversion, the right-hand side is just $T_{\hat{0}}^D$. The $S \in \mathcal{S}$ such that $\pi(S) = \hat{0}$ are just the placements in which no two rooks lie in the same row or column. The rooks on B then define a path cover and the rooks on the appended rows then ensure that distinct paths are assigned distinct colors. The theorem follows from Proposition 5. ■

It is not hard to show that this result specializes to [GJW4, Theorem 1(a)]. One might again object that Proposition 18 is contrived because $T_{\geq \pi}^D$ is simply a formal device to represent what one gets by Möbius inversion. This time the objection is harder to meet, because $T_{\geq \pi}^D$ is not as “nice” an object as $\tilde{\Xi}_{\lambda, \mu}$. For example, $\text{type}(\pi) = \text{type}(\sigma)$ does *not* imply $T_{\geq \pi}^D = T_{\geq \sigma}^D$. However, we do have one result that gives some more information about $T_{\geq \pi}^D$.

PROPOSITION 19. *If D is an acyclic digraph then $T_{\geq \pi}^D$ is p -positive.*

Proof. Let B be the associated board. We have

$$T_{\geq \pi}^D = \sum_{\{S \in \mathcal{S} \mid \pi(S) \geq \pi\}} x^S.$$

Collect terms that have identical placements of rooks on B . From the definitions we see that each such collection of terms corresponds to the set of colorings of $V(D)$ that are monochromatic on the connected components of the subgraph of D whose edges are those selected by the placement of rooks on B , except that the condition $\pi(S) \geq \pi$ imposes the further condition that components which contain elements of the same block of π must always be colored the same color. This gives a power sum symmetric function, so $T_{\geq \pi}^D$ is a sum of power sums. ■

Note that while Proposition 18 resembles Stanley’s formula [St2, Theorem 2.6]

$$X_G = \sum_{\pi \in L_G} \mu(\hat{0}, \pi) p_{\pi}$$

(where L_G is the lattice of contractions of G), there is a significant difference in that in Proposition 18 all the dependence on the digraph is contained in the $T_{\geq \pi}^D$ whereas

for X_G all the dependence is contained in L_G . A variant of of Proposition 18 can be obtained by considering rook placements with no two rooks in the same *row*, but this result also seems contrived and does not suggest any satisfactory analogue of L_G , so we omit the details.

More rook theory can undoubtedly be generalized to our symmetric function context, but we shall now turn to a different aspect of Ξ_D .

4. THE POSET CHAIN CONJECTURE

One of the original motivations for studying X_G and Ξ_D is a conjecture by Stanley and Stembridge [S-S, Conjecture 5.5] called the *Poset Chain Conjecture*. We restate this conjecture here for convenience. Following Stanley [St2, section 5], we write $\mathbf{a} + \mathbf{b}$ for the poset that is a disjoint union of an a -element chain and a b -element chain, and we say that a poset is $(\mathbf{a} + \mathbf{b})$ -free if it contains no induced subposet isomorphic to $\mathbf{a} + \mathbf{b}$. Then the Stanley-Stembridge conjecture is equivalent to the following.

CONJECTURE 1. *If P is a $(\mathbf{3} + \mathbf{1})$ -free poset, then $X_{G(P)}$ is e -positive.*

In view of Proposition 2, this conjecture can also be viewed as a conjecture about Ξ_D . One of the most important partial results is the following theorem of Gasharov [Ga1].

PROPOSITION 20. *If P is a $(\mathbf{3} + \mathbf{1})$ -free poset, then $X_{G(P)}$ is s -positive. ■*

The main result of this section is a slight extension of Gasharov's result that will illustrate the subtlety of Conjecture 1. To state our result we need some definitions.

DEFINITION. A loopless digraph is *weakly $(\mathbf{3} + \mathbf{1})$ -free* if, for any ordered pair (u, v) of vertices of D , either D or D' fails to have a directed path of length two from u to v .

Note that weakly $(\mathbf{3} + \mathbf{1})$ -free digraphs need not be transitively closed or even acyclic. Our nomenclature is justified by the following proposition.

PROPOSITION 21. *If P is a poset, then P is $(\mathbf{3} + \mathbf{1})$ -free if and only if $D(P)$ is weakly $(\mathbf{3} + \mathbf{1})$ -free.*

Proof. Saying that $D(P)$ is weakly $(\mathbf{3} + \mathbf{1})$ -free is equivalent to saying that if $u \rightarrow v \rightarrow w$ is a directed path of length two in $D(P)$ and x is any element such that (u, x) is *not* an edge of $D(P)$, then (x, w) is an edge of $D(P)$. Saying that P is $(\mathbf{3} + \mathbf{1})$ -free is equivalent to saying that if $u \rightarrow v \rightarrow w$ is a chain in P and x is any element such that $u \not\prec x$ in P , then $x < w$ in P . Clearly these two are equivalent. ■

DEFINITION. Let D be a digraph. A D -array is an array

$$\begin{array}{ccc} v_{1,1} & v_{1,2} & \dots \\ v_{2,1} & v_{2,2} & \dots \\ \dots & & \end{array}$$

where each $v_{i,j}$ is either undefined or an element of D and such that

1. for all $i, j \geq 1$, if $v_{i,j+1}$ is defined, then $v_{i,j}$ is defined and $(v_{i,j}, v_{i,j+1})$ is an edge of D , and
2. every element of D appears exactly once in the array.

The *shape* of a D -array is the sequence of the lengths of (the defined portion of) the rows. A D -tableau is a D -array such that

3. for all $i, j \geq 1$, if $v_{i+1,j}$ is defined, then $v_{i,j}$ is defined and $(v_{i+1,j}, v_{i,j})$ is *not* an edge of D .

Our definitions of D -array and D -tableau are motivated by Gasharov's use of Gessel-Viennot [GV2] P -arrays and P -tableaux in his proof of Proposition 20. We can now state our generalization.

THEOREM 5. *If D is a weakly $(\mathbf{3} + \mathbf{1})$ -free digraph, then the coefficient of s_λ in $\Xi_D(x, 0)$ is the number of D -tableaux of shape λ .*

Proof. The proof is almost identical to Gasharov's, and we refer to his paper for some details which we shall omit. Let S_ℓ denote the group of permutations of $[\ell]$. If $\lambda = (\lambda_1, \dots, \lambda_\ell)$ is an integer partition and $\pi \in S_\ell$, then we denote by $\pi(\lambda)$ the sequence

$$(\lambda_{\pi(j)} - \pi(j) + j)_{j=1}^\ell.$$

Define c_λ by

$$\Xi_D(x, 0) = \sum_{\lambda} c_\lambda s_\lambda(x).$$

By the same Jacobi-Trudi argument that Gasharov uses,

$$c_\lambda = \sum_{\pi \in S_\ell} (\text{sgn } \pi) \cdot (\text{coefficient of } \prod x_i^{\pi(\lambda)_i} \text{ in } \Xi_D(x, 0)),$$

where $\text{sgn } \pi$ is the sign of the permutation π . Now by Proposition 5, $\Xi_D(x, 0)$ counts path colorings of D , and path colorings of D are in bijection with D -arrays (the rows of the D -array give the directed paths and the path in row i is assigned the color i). If we let

$$A = \{(\pi, T) \mid \pi \in S_\ell \text{ and } T \text{ is a } D\text{-array of shape } \pi(\lambda)\},$$

it then follows that

$$c_\lambda = \sum_{(\pi, T) \in A} \text{sgn } \pi.$$

Now let

$$B = \{(\pi, T) \in A \mid T \text{ is not a } D\text{-tableau}\}$$

and note that if T is a D -tableau, then $\pi(\lambda)_1 \geq \pi(\lambda)_2 \geq \dots$ so that π must be the identity permutation. Thus to prove the theorem it suffices to find an involution $\varphi : B \rightarrow B$ such that if $(\sigma, T') = \varphi(\pi, T)$ then $\text{sgn } \sigma = -\text{sgn } \pi$. Gasharov's involution works without modification; for completeness we restate it here. If

$$T = \begin{array}{ccc} v_{1,1} & v_{1,2} & \dots \\ v_{2,1} & v_{2,2} & \dots \\ \dots & & \end{array}$$

then let $c = c(T)$ be the smallest positive integer such that condition 3 fails for $j = c$ and some i . Let $r = r(T)$ be the largest i with this property. Define $\sigma = \pi \circ (r, r+1)$ where $(r, r+1)$ is the permutation that interchanges r and $r+1$. Define

$$T' = \begin{array}{ccc} u_{1,1} & u_{1,2} & \dots \\ u_{2,1} & u_{2,2} & \dots \\ \dots & & \end{array}$$


by letting

- (a) $u_{i,j} = v_{i,j}$ if $i \neq r$ or $i \neq r+1$ or $(i = r \text{ and } j \leq c-1)$ or $(i = r+1 \text{ and } j \leq c)$;
- (b) $u_{r,j} = v_{r+1,j+1}$ if $j \geq c$ and $v_{r+1,j+1}$ is defined;
- (c) $u_{r+1,j} = v_{r,j-1}$ if $j \geq c+1$ and $v_{r,j-1}$ is defined.

(Other values of the array T' remain undefined.) Now row $r+1$ of T' satisfies condition 1, because if $v_{r,c}$ is defined then $(v_{r+1,c}, v_{r,c})$ is an edge of D by definition of r and c . To show that T' is a D -array it suffices to show that row r satisfies condition 1 since condition 2 is obviously satisfied. Possible trouble arises only if $c \geq 2$, but then $v_{r+1,c-1} \rightarrow v_{r+1,c} \rightarrow v_{r+1,c+1}$ is a path of length two in D and $(v_{r+1,c-1}, v_{r,c-1})$ is *not* an edge in D , so it follows from the assumption that D is weakly $(\mathbf{3} + \mathbf{1})$ -free that $(v_{r,c-1}, v_{r+1,c+1})$ is an edge of D , and condition 1 is met. Now $(v_{r+1,c}, v_{r+1,c+1})$ is an edge of D so if $u_{r,c}$ is defined $(u_{r+1,c}, u_{r,c})$ is an edge of D (since $u_{r+1,c} = v_{r+1,c}$ and $u_{r,c} = v_{r+1,c+1}$) and thus T' is not a D -tableau. It is clear that T' has shape $\sigma(\lambda)$ and that $c(T') = c(T)$ and $r(T') = r(T)$. Also, $\text{sgn } \sigma = -\text{sgn } \pi$, so φ is the desired sign-reversing involution. ■

Note that a digraph is weakly $(\mathbf{3} + \mathbf{1})$ -free if and only if its complement is weakly $(\mathbf{3} + \mathbf{1})$ -free, so Corollary 2 applied to Theorem 5 does not enlarge the class of known s -positive path-cycle symmetric functions.

It is natural to conjecture that if D is weakly $(\mathbf{3} + \mathbf{1})$ -free then $\Xi_D(x, 0)$ is e -positive, but for instance if we let D be the digraph



we find (with the aid of John Stembridge's SF package for Maple) that

$$\Xi_D(x, 0) = s_4 + 2s_{31} + s_{22} + 4s_{211} + 3s_{1111} = 3e_{31} - e_{211} + e_{1111}.$$

In fact, of the five essentially distinct weakly $(\mathbf{3} + \mathbf{1})$ -free acyclic digraphs on four vertices that are not transitively closed, only one is e -positive. So the way the property of being $(\mathbf{3} + \mathbf{1})$ -free is used in Gasharov's proof is far from enough to yield e -positivity even if the condition of acyclicity is added. This shows how delicate Conjecture 1 is.

Chapter 3

The Chromatic Symmetric Function

In this chapter we prove a number of miscellaneous results about the chromatic symmetric function X_G . The basic reference for X_G is [St2] (but see also [St4][Ga1][Ga2]).

1. G -ASCENTS

As explained in [St2], the expansion of X_G in terms of fundamental quasi-symmetric functions has an interpretation in terms of P -partitions. In the previous chapter we saw that this implies that the coefficients of the $\tilde{\mathcal{E}}$ -expansion of X_G have a combinatorial interpretation in terms of P -partitions. In this section we give another combinatorial interpretation of the coefficients that is based on the concept of a G -descent (see [CG2]; the definition in [CG1] contains a minor error) or the equivalent concept of a G -ascent.

DEFINITION. Fix a graph G with vertex set $[n]$. Given a permutation π of $[n]$ and a vertex $v \in [n]$, define the *rank* $\rho_\pi(v)$ of v to be the largest integer r for which there exists an increasing sequence of positive integers

$$i_1 < i_2 < \cdots < i_r < i_{r+1} = \pi^{-1}(v)$$

such that $\{\pi(i_j), \pi(i_{j+1})\}$ is an edge of G for all j . We say that π has a G -*ascent* at v if either

- (i) $\rho_\pi(v) < \rho_\pi(w)$, or
- (ii) $\rho_\pi(v) = \rho_\pi(w)$ and $v < w$,

where $w = \pi(\pi^{-1}(v) + 1)$. The G -*ascent type* of π is the integer partition of n whose parts are the lengths of the subwords obtained by breaking the one-line representation of π after each number at which π has a G -ascent.

For example, suppose that $n = 3$ and that G has an edge between 1 and 3 and no other edges. Let π be the permutation 132, i.e., the permutation that fixes 1 and exchanges 2 and 3. Then

$$\rho_\pi(1) = \rho_\pi(2) = 0 \quad \text{and} \quad \rho_\pi(3) = 1.$$

Furthermore, π has a unique G -ascent (at 1), and breaking 132 after the 1 gives two subwords, one with two letters and the other with one letter, so the G -ascent type of 132 is the integer partition $(2, 1)$.

THEOREM 6. *Let G be a graph with $V(G) = [n]$. If*

$$X_G = \sum_{\lambda} N_{\lambda} \tilde{\Xi}_{\lambda}$$

is the $\tilde{\Xi}$ -expansion of X_G , then N_{λ} is the number of permutations of $[n]$ with G -ascent type λ .

Proof. Given a stable partition σ of G and a linear ordering $(\sigma_1, \sigma_2, \dots, \sigma_{\ell})$ of its blocks, we will define an associated permutation of $[n]$ (represented in one-line notation) that will consist of a certain permutation of the elements of σ_1 (the exact permutation will be specified in a moment) followed by a certain permutation of the elements of σ_2 , and so on.

Before we specify how the elements within each σ_i are ordered, we first notice that if π_1 and π_2 are any two permutations of the above form, then $\rho_{\pi_1}(v) = \rho_{\pi_2}(v)$ for all v , because there are no edges of G between vertices in the same block of σ . Thus we may speak of the “rank of a vertex” without ambiguity, even before specifying the permutation.

We now single out a particular permutation of the above form by arranging the vertices within each block as follows: first arrange the vertices in decreasing order of rank, and then within each rank, arrange the vertices in numerically decreasing order. The motivation for this arrangement is that in the resulting permutation,

the only place that a G -ascent can occur is between adjacent blocks of σ . This fact will play an important role shortly.

We next claim that the number of ordered stable partitions of type μ that give rise to a fixed permutation π equals $\ell(\mu)! a_{\lambda\mu}$, where λ is the G -ascent type of π and $a_{\lambda\mu}$ is the coefficient of \tilde{m}_μ in $\tilde{\Xi}_\lambda$. To see this, let σ be the ordered partition obtained by breaking π after each number at which there is a G -ascent. A *refinement* of σ is defined to be an ordered partition obtained by breaking π at each break of σ , plus (optionally) any number of other locations. From the definition of $\tilde{\Xi}_\lambda$, we see that to prove our claim, it suffices to show that it is precisely the type μ refinements of σ that generate π . We note first that it is only such refinements that can possibly generate π , by the remark at the end of the preceding paragraph. It thus remains to show that every such refinement is in fact an ordered *stable* partition that generates π .

To see that every refinement is stable, it suffices to show that σ is stable. But if $i < j$ and $\pi(i)$ and $\pi(j)$ are adjacent, then the rank of $\pi(j)$ must exceed the rank of $\pi(i)$, so there must be at least one G -ascent between i and j , and thus $\pi(i)$ and $\pi(j)$ must be in different blocks of σ .

To see that every refinement generates π , note that the only way this could fail to happen is if the elements in some block of the refinement are not ordered properly (i.e., in decreasing order of rank, and in decreasing numerical order within each rank). However, rank is defined without reference to partitions, so when we pass from σ to one of its refinements, the vertices must be ordered properly in the refinement, provided that they were ordered properly in σ . Finally, the vertices of σ are ordered properly, because otherwise there would be a G -ascent inside a block of σ . This completes the proof of our claim.

Thus, if we let b_μ be the number of stable partitions of G of type μ , then for all μ ,

$$\ell(\mu)! b_\mu = \sum_{\lambda} N_\lambda \ell(\mu)! a_{\lambda\mu}.$$

Multiplying both sides by $\tilde{m}_\mu/\ell(\mu)!$ and summing over μ ,

$$X_G = \sum_{\mu} b_\mu \tilde{m}_\mu = \sum_{\lambda} N_\lambda \sum_{\mu} a_{\lambda\mu} \tilde{m}_\mu = \sum_{\lambda} N_\lambda \tilde{\Xi}_\lambda. \blacksquare$$

2. THE POSET CHAIN CONJECTURE REVISITED

In this section we give a new combinatorial proof of a (known) special case of the Poset Chain Conjecture. We hope that our method can be generalized to handle more cases, although so far we have not been able to do so.

Let P be a poset. Recall from [GV2] or [Ga1] that a P -array is an array

$$\begin{array}{ccc} v_{1,1} & v_{1,2} & \dots \\ v_{2,1} & v_{2,2} & \dots \\ \dots & & \end{array}$$

where each $v_{i,j}$ is either undefined or an element of P and such that

1. every element of P appears exactly once in the array, and
2. for all $i, j \geq 1$, if $v_{i,j+1}$ is defined, then $v_{i,j}$ is defined and $v_{i,j} < v_{i,j+1}$ in P .

The *shape* of a P -array is the sequences of the lengths of (the defined portion of) the rows. A P -tableau is a P -array such that

3. for all $i, j \geq 1$, if $v_{i+1,j}$ is defined, then $v_{i,j}$ is defined and $v_{i+1,j} \not< v_{i,j}$ in P .

We then have the following result of Gasharov [Ga1].

PROPOSITION 22. *If P is a $(\mathbf{3} + \mathbf{1})$ -free poset and*

$$X_{G(P)} = \sum_{\lambda} a_{\lambda} s_{\lambda}$$

is the s -expansion of $X_{G(P)}$, then a_{λ} is the number of P -tableaux of shape λ . ■

It is also well known [Mac, Table 1, p. 56] that

$$e_{\lambda} = \sum_{\mu} K_{\mu' \lambda} s_{\mu},$$

where $K_{\mu' \lambda}$ is the number of semi-standard (i.e., row-nondecreasing and column-increasing) Young tableaux of shape μ' and content λ . This points the way to a possible strategy for proving e -positivity of a $(\mathbf{3} + \mathbf{1})$ -free poset P combinatorially: try to find a partition of the set of all P -tableaux such that for each block there is an integer partition λ such that the number of P -tableaux of shape μ in that block equals $K_{\mu' \lambda}$. From the above facts, we see that the existence of such a partition implies that P is e -positive.

So far this method has not yielded new e -positivity results, but it does lead to a combinatorial proof of the known fact [S-S] that $\mathbf{3}$ -free posets (i.e., posets which do not contain an induced subposet isomorphic to a three-element chain) are e -positive.

PROPOSITION 23. *If P is a $\mathbf{3}$ -free poset, then $X_{G(P)}$ is e -positive.*

Proof. Let us define a P -diagram to be an arrangement of the elements of P into two columns, justified along the top edge, such that the height of the right-hand column

does not exceed the height of the left-hand column. (The columns are allowed to be empty.) We define the *popping* operation as follows. To pop a P -diagram, remove the bottommost element from the right-hand column and place it at the bottom of the left-hand column. (It is illegal to pop P -diagrams with empty right-hand columns.)

Assume now that P is **3**-free. Notice that this implies that every P -array has at most two columns, so that in particular every P -array is a P -diagram. If T and T' are P -tableaux, we define the relation \sim by letting $T \sim T'$ if one of them can be obtained from the other by a sequence of pops. Clearly \sim is an equivalence relation, and hence it induces a partition $\pi = (\pi_1, \pi_2, \dots)$ of the set of all P -tableaux.

For each i , let T_i be the P -tableau in π_i with the highest right-hand column. Now fix any i . We claim that any P -diagram obtained from T_i by a sequence of pops is in fact a P -tableau (and therefore lies in π_i). To see this, note that the only way a problem can arise is if an element of P that is shifted to the left-hand column in the course of a pop turns out to be less than the element that it ends up sitting underneath. However, because T_i is a P -array, any element x in the right-hand column of T_i is larger than at least one other element of P (namely, the element immediately to the left of x in T_i), and since P is **3**-avoiding, x cannot be less than any other element in P . This establishes our claim.

Let λ^i be the conjugate of the shape of T_i . In view of the remarks preceding the theorem statement, we see that all there is left to prove is that for all i , the number of P -tableaux in π_i of shape μ equals K_{μ', λ^i} . Now since λ^i has at most two parts, $K_{\mu', \lambda^i} = 1$ if $|\mu'| = |\lambda^i|$, μ' has at most two parts and $\mu'_1 \geq \lambda^i_1$, and $K_{\mu', \lambda^i} = 0$ otherwise. Then from the previous paragraph we see that there is indeed exactly one P -tableau in π_i of shape μ when $K_{\mu', \lambda^i} = 1$ and there are no P -tableaux of other shapes. ■

3. RECONSTRUCTION

In [St2], several expansions of the chromatic polynomial are generalized to expansions of X_G . One notable omission is the multiplicative expansion (see [Big, Chapter 11]), which has an important application to the graph reconstruction conjecture. It turns out, however, that the multiplicative expansion of a graph invariant even more general than X_G was derived by Tutte a long time ago. In this section we indicate the connection between Tutte's work and the theory of the chromatic symmetric function, and in particular we show how Tutte's work implies that X_G is reconstructible.

Recall that the *list of vertex-deleted subgraphs* of a graph G is the multiset of unlabelled graphs

$$\{G \setminus v \mid v \in V(G)\}.$$

The graph G is *reconstructible* if no other graph has the same list of vertex-deleted subgraphs that G does. It is a major unsolved problem in graph theory whether or not every graph with more than two vertices is reconstructible. See [Bon] for a recent survey.

If G is a graph and $S \subset E(G)$, define $G \cdot S$ to be the subgraph of G consisting of the edges of S together with their incident vertices. (In particular, $G \cdot S$ has no isolated vertices.) Now for each isomorphism class K of connected graphs with at least two vertices, let x_K be an indeterminate. Also let t be an indeterminate distinct from all the x_K . Given any graph G , define

$$J(G) \stackrel{\text{def}}{=} \sum_{S \subset E(G)} \prod_{H \subset G \cdot S} t^{|V(H)|} x_H,$$

where the product is over all connected components H of $G \cdot S$. (If this product is empty we take its value to be one.) We then have the following theorem of Tutte [Tu1, 6.6].

PROPOSITION 24. *The coefficients of all the terms of $J(G)$ are reconstructible, save possibly those terms containing x_H where $V(H) = V(G)$. ■*

We remark, for the benefit of the reader who wishes to consult Tutte's paper, that we have used $J(G)$ for what Tutte calls $J(E(G))$, that we have used x_H in place of $f(H)$, and that we are restricting our attention to the $C(G; 1a)$ case.

From Proposition 24 it is easy to derive the following result.

PROPOSITION 25. *X_G is reconstructible.*

Proof. From [St2, Theorem 2.5], we have

$$X_G = \sum_{S \subset E(G)} (-1)^{|S|} p_{\lambda(S)},$$

where $\lambda(S)$ denotes the partition of d whose parts are equal to the vertex sizes of the connected components of the spanning subgraph of G with edge set S . It suffices to show that

$$\tilde{X}_G \stackrel{\text{def}}{=} \sum_{S \subset E(G)} (-1)^{|S|} \prod_{H \subset G \cdot S} p_{|V(H)|}$$

is reconstructible; what we have done is to “set $p_1 = 1$ in X_G ,” which might appear at first to result in some loss of information, but it is easy to see that X_G can be recovered from \tilde{X}_G using the fact that X_G is homogeneous.

Now set $t = 1$ and $x_H = (-1)^{|E(H)|} p_{|V(H)|}$. Then $J(G)$ becomes \tilde{X}_G , and from Proposition 24 it follows that all the coefficients in the p -expansion of \tilde{X}_G can be

reconstructed except possibly for the coefficient of p_n , where $n = |V(G)|$. However, the coefficient of p_n in \tilde{X}_G equals the coefficient of p_n in X_G , which in turn equals the coefficient of i in the chromatic polynomial $\chi_G(i)$ (as we can see by specializing $X_G(1^i) = \chi_G(i)$). Since the chromatic polynomial is reconstructible ([Tu1, 7.5] or [Tu2]), this completes the proof. ■

4. SUPERFICATION

Recall that the *superfication* of a symmetric function g is defined by

$$g(x/y) \stackrel{\text{def}}{=} \omega_y g(x, y).$$

Stanley [St2] has briefly considered the superfication of X_G and has also asked what can be said about the two-variable polynomial $X_G(1^i/1^j)$. We investigate the latter question in this section.

Some results can be obtained trivially by specializing known theorems about X_G . For example, we can specialize [St2, Theorem 4.3] as follows. Let \mathbb{P} and $\bar{\mathbb{P}}$ be two disjoint copies of the positive integers. Denote the elements of \mathbb{P} by $1, 2, 3, \dots$ and denote the elements of $\bar{\mathbb{P}}$ by $\bar{1}, \bar{2}, \bar{3}, \dots$. Linearly order the disjoint union $\mathbb{P} \cup \bar{\mathbb{P}}$ by using the natural order on each of the sets \mathbb{P} and $\bar{\mathbb{P}}$ and, additionally, declaring every element of \mathbb{P} to be less than every element of $\bar{\mathbb{P}}$. The following proposition then follows immediately from [St2, Theorem 4.3].

PROPOSITION 26. *For any graph G , $X_G(1^i/1^j)$ is the number of pairs (\mathfrak{o}, κ) such that \mathfrak{o} is an acyclic orientation of G and*

$$\kappa : V(G) \rightarrow \{1, 2, \dots, i\} \cup \{\bar{1}, \bar{2}, \dots, \bar{j}\}$$

is a map such that (a) if $u \rightarrow v$ is an edge of \mathfrak{o} then $\kappa(u) \geq \kappa(v)$, and (b) if $u \rightarrow v$ is an edge of \mathfrak{o} and both $\kappa(u)$ and $\kappa(v)$ lie in \mathbb{P} , then $\kappa(u) > \kappa(v)$. ■

A second example is the following result, whose proof we omit since it follows a standard line of argumentation in the theory of P -partitions and is somewhat long. The reader is referred to [St3, Chapter 4] and [St5] for the relevant theory and definitions.

PROPOSITION 27. *For any graph G with n vertices,*

$$X_G(1^i/1^j) = \sum_{\mathfrak{o}} \sum_{\pi \in \mathcal{L}(\bar{\mathfrak{o}})} \sum_{k=0}^n \binom{i + D_{n-k}(\pi)}{n-k} \binom{j + A_k(\pi)}{k},$$

where the first sum is over all acyclic orientations \mathfrak{o} of G , $\bar{\mathfrak{o}}$ denotes the poset that is the transitive closure of \mathfrak{o} , \mathcal{L} denotes the Jordan-Hölder set, $D_k(\pi)$ denotes the

number of descents in the first k digits of π , and $A_k(\pi)$ denotes the number of ascents in the last k digits of π . ■

For our third and final example we need a definition:

$$\tilde{\chi}_G(m, n) \stackrel{\text{def}}{=} X_G(1^{(m-n)/2}/1^{(m+n)/2}).$$

PROPOSITION 28. *The coefficients of the polynomial $\tilde{\chi}_G(m, n)$ are nonnegative integers.*

Proof. Corollary 2.7 of [St2] states that ωX_G is p -positive, i.e., X_G is a nonnegative integer combination of $(\text{sgn } \lambda)p_\lambda$, or a polynomial in the signed power sums $(\text{sgn } d)p_d = (-1)^{d-1}p_d$ with nonnegative integer coefficients. Now

$$(-1)^{d-1}p_d(x/y) = (-1)^{d-1}(p_d(x) + (-1)^{d-1}p_d(y)) = (-1)^{d-1}p_d(x) + p_d(y),$$

so

$$(-1)^{d-1}p_d(1^i/1^j) = (-1)^{d-1}i + j.$$

If we set $m = j + i$ and $n = j - i$, then we see that $X_G(1^i/1^j)$ is a polynomial in m and n with nonnegative integer coefficients. By rewriting the equations $m = j + i$ and $n = j - i$ in the form $i = (m - n)/2$ and $j = (m + n)/2$ we see that the proposition follows. ■

Apart from such specializations of known theorems, however, it seems hard to find significant results. For example, it is easily checked that no deletion-contraction recurrence is possible. The best we have been able to do is to find a recurrence for trees which has some interesting properties. Unfortunately we have not found any applications for this recurrence, although it seems that it might be possible to use it to obtain some partial results towards the question, posed by Stanley, of whether nonisomorphic trees have distinct chromatic symmetric functions.

To state the recurrence we must first define two invariants that are closely related to $\tilde{\chi}_G$. Let T be a *rooted* tree, i.e., a tree with a distinguished vertex. Let T' be the tree obtained by adjoining an extra vertex v to T that is adjacent to the root of T (and to no other vertices). Define $\alpha_T(m, n)$ to be the number of pairs (σ, κ) such that σ is an acyclic orientation of T' and $\kappa : V(T') \rightarrow \mathbb{P} \cup \bar{\mathbb{P}}$ is a map that sends v to 1, whose image lies in

$$\{1, 2, \dots, (m - n)/2\} \cup \{\bar{1}, \bar{2}, \dots, \overline{(m + n)/2}\},$$

and that also satisfies the conditions (a) and (b) of Proposition 26. Define $\beta_T(m, n)$ in the same way except with the phrase “sends v to 1” replaced by “sends v to $\bar{1}$.” We then have the following theorem.

THEOREM 7. Let T be a rooted tree with root v , and let \mathcal{F} be the family of rooted trees that results when v is deleted from T . (The root of a tree in \mathcal{F} is the vertex that was adjacent to v before deletion.) Then

$$2\alpha_T(m, n) = (m - n - 2) \prod_{S \in \mathcal{F}} \alpha_S(m, n) + (m + n) \prod_{S \in \mathcal{F}} \beta_S(m, n)$$

and

$$2\beta_T(m, n) = (m + n + 2) \prod_{S \in \mathcal{F}} \beta_S(m, n) + (m - n) \prod_{S \in \mathcal{F}} \alpha_S(m, n).$$

Proof. Fix an arbitrary map $\kappa : V(T) \rightarrow \{1, 2, \dots, (m-n)/2\} \cup \{\bar{1}, \bar{2}, \dots, \overline{(m+n)/2}\}$. We seek acyclic orientations σ of T that are compatible with κ in the sense that (σ, κ) satisfies conditions (a) and (b) of Proposition 26. Observe first that if κ maps any two adjacent vertices to the same element of \mathbb{P} , then no acyclic orientations are compatible with κ . Otherwise, we note that conditions (a) and (b) force a particular orientation of every edge except those edges whose endvertices are mapped to the same element of $\bar{\mathbb{P}}$, in which case the conditions (a) and (b) impose no constraint on the orientation. Now since T is a tree, every orientation of its edges is an acyclic orientation, and in particular every orientation of the edges of T that meets the necessary conditions just stated is in fact an acyclic orientation compatible with κ .

With this in mind, we now consider α_T . Adjoin a vertex v to form T' , and map v to 1. We wish to extend this to a map κ of all the vertices of T' into

$$\{1, 2, \dots, (m - n)/2\} \cup \{\bar{1}, \bar{2}, \dots, \overline{(m + n)/2}\},$$

and for each such κ we want to find all compatible acyclic orientations of T' . The total number of such compatible pairs will give us $\alpha_T(m, n)$. Begin by splitting into two cases: in the first case, the root of T is mapped into \mathbb{P} , and in the second case, the root of T is mapped into $\bar{\mathbb{P}}$. In the first case, mapping the root to 1 results in adjacent vertices being mapped to the same element of \mathbb{P} , so we discard this possibility. There remain

$$\frac{m - n}{2} - 1 = \frac{m - n - 2}{2}$$

other possible colors in \mathbb{P} for the root of T . Now comes the key observation: for each such coloring of the root of T , we obtain

$$\prod_{S \in \mathcal{F}} \alpha_S(m, n)$$

corresponding compatible pairs. For once the color of the root of T is fixed, each tree in \mathcal{F} may be oriented and colored independently, and the number of ways of doing this for a particular $S \in \mathcal{F}$ is just $\alpha_S(m, n)$ —the root of T plays the role of

the adjoined vertex of S' . (There is a slight technicality here in that in the definition of α_S the adjoined vertex is required to be mapped to 1, whereas here it may be mapped to an arbitrary element of \mathbb{P} , but it is clear that there is a bijection between the two sets of configurations, with the only difference being that the orientation of the edge joining the root of S to the root of T may need to be reversed in some cases.) Furthermore, the orientation of the edge between v and the root of T is forced. Thus the total number of compatible pairs in the first case is

$$\binom{m-n-2}{2} \prod_{S \in \mathcal{F}} \alpha_S(m, n).$$

If we now consider the second case, we see that there are $(m+n)/2$ choices of colors in $\bar{\mathbb{P}}$ for the root of T and that each such choice forces the orientation of the edge between v and the root of T . The number of ways of extending each such configuration to the entire tree is

$$\prod_{S \in \mathcal{F}} \beta_S(m, n).$$

Adding up the two cases proves the first formula of our theorem.

The second formula is proved similarly; the only new twist is that instead of *discarding* the case where v and the root of T are assigned the same color, we must *count it twice*, since in this case the edge between v and the root of T may be oriented in either direction. This accounts for the term $(p+q+2)$. ■

Theorem 7 allows us to state the precise relationship between $\tilde{\chi}$ and α and β .

PROPOSITION 29. *Let T be a rooted tree. Then*

$$\tilde{\chi}_T(m, n) = \frac{(n+m)\alpha_T(m, n) + (n-m)\beta_T(m, n)}{2(n+1)}.$$

Proof. Let \mathcal{F} be the family of rooted trees resulting when the root of T is deleted. Then it is clear from the definitions of $\tilde{\chi}$, α and β

$$\tilde{\chi}_T(m, n) = \binom{m-n}{2} \prod_{S \in \mathcal{F}} \alpha_S(m, n) + \binom{m+n}{2} \prod_{S \in \mathcal{F}} \beta_S(m, n),$$

where the two terms in this sum correspond to mapping the root of T into \mathbb{P} and $\bar{\mathbb{P}}$ respectively. We can then invert the formulas in Theorem 7 to solve for the products in terms of α_T and β_T (since the determinant of

$$\begin{pmatrix} m-n-2 & m-n \\ m+n & m+n+2 \end{pmatrix}$$

is not identically zero) to obtain the desired formula. ■

Notice that the left-hand side of Proposition 29 is independent of the choice of root, even though the terms on the right-hand side are not.

If we compute a few examples by hand, we quickly notice that α and β are closely related to each other. More precisely, we say that two polynomials $\gamma_1(m, n)$ and $\gamma_2(m, n)$ form a *related pair* if the sign of the coefficient of any nonzero term $m^r n^s$ in $\gamma_1(m, n)$ equals $(-1)^{d-r}$ (where d denotes the degree of γ_1) and if γ_2 can be obtained from γ_1 by changing all the minus signs to plus signs.

PROPOSITION 30. *Let T be a rooted tree with d vertices. Then the degree of α_T equals d and α_T and β_T form a related pair.*

Proof. That the degree is d follows easily (e.g., by induction). For the second part of the proposition, we proceed by induction on d . Let \mathcal{F} be the family of rooted trees obtained by deleting the root of T . Then for all $S \in \mathcal{F}$, α_S and β_S form a related pair by the induction hypothesis, since each $S \in \mathcal{F}$ has fewer than d vertices. It follows easily that

$$\prod_{S \in \mathcal{F}} \alpha_S \quad \text{and} \quad \prod_{S \in \mathcal{F}} \beta_S$$

form a related pair, with degree $d - 1$. We may rewrite Theorem 7 in the form

$$\begin{aligned} \alpha_T &= \frac{1}{2}m \left(\prod_{S \in \mathcal{F}} \beta_S + \prod_{S \in \mathcal{F}} \alpha_S \right) + \frac{1}{2}n \left(\prod_{S \in \mathcal{F}} \beta_S - \prod_{S \in \mathcal{F}} \alpha_S \right) - \prod_{S \in \mathcal{F}} \alpha_S \\ \beta_T &= \frac{1}{2}m \left(\prod_{S \in \mathcal{F}} \beta_S + \prod_{S \in \mathcal{F}} \alpha_S \right) + \frac{1}{2}n \left(\prod_{S \in \mathcal{F}} \beta_S - \prod_{S \in \mathcal{F}} \alpha_S \right) + \prod_{S \in \mathcal{F}} \beta_S \end{aligned}$$

where we have omitted some of the m 's and n 's for brevity.

We now show that the terms in α_T have the appropriate sign. Since $\prod \alpha_S$ and $\prod \beta_S$ form a related pair, their sum contains only terms whose power of m differs from the highest power of m (namely $d - 1$) by an even integer, and moreover the surviving terms are all nonnegative. Hence the contribution to α_T from the first summand has the correct signs. For the second summand, note that $\prod \beta_S - \prod \alpha_S$ contains only terms that differ from $d - 1$ by an *odd* integer, and that all these terms are nonnegative. Multiplying by n does not change any exponents of m , so the contribution from the second summand also has the correct signs. Finally, it is clear that $-\prod \alpha_S$ also contributes the correct signs.

To conclude the proof it suffices to show that β_T can be obtained by changing all the minus signs in α_T to plus signs. As we noted before, the first two summands consist entirely of nonnegative terms, so the only thing we need to check is that

changing all minus signs to plus signs in $-\prod \alpha_S$ gives $\prod \beta_S$, but this follows from the induction hypothesis. ■

We conclude this section with the remark that if we could show that α_T is always irreducible, then it would follow that distinct rooted trees always have distinct α_T 's. For suppose we are given a polynomial and are told that it equals α_T for some rooted tree T . Then Proposition 30 lets us compute β_T , and then inverting the formulas in Theorem 7 allows us to compute $\prod \alpha_S$ and $\prod \beta_S$. By irreducibility we can then recover the α_S , and by induction we may assume that all the S 's may be reconstructed from the corresponding α_S . This would then allow us to reconstruct T . Unfortunately, proving irreducibility seems even harder than proving the original conjecture!

5. $X_G(t)$

In [St4] Stanley considers briefly some generalizations of X_G , including an invariant that he calls $X_G(t)$. Conceivably, an entire thesis could be written about $X_G(t)$, but here we present only the most basic facts, including one result that is mentioned in [St4] but whose proof is omitted there.

Let G be a graph and let κ be a map from $V(G)$ into the positive integers. Define

$$x^\kappa \stackrel{\text{def}}{=} \prod_{v \in V(G)} x_{\kappa(v)}$$

and say that an edge is *monochromatic* if κ maps its endvertices to the same integer. Stanley [St4] defines

$$X_G(t) \stackrel{\text{def}}{=} \sum_{\kappa} (1+t)^{m(\kappa)} x^\kappa,$$

where the sum is over all maps κ from $V(G)$ into the positive integers and $m(\kappa)$ denotes the number of monochromatic edges. The motivation for this definition is that if we set n of the x 's equal to one and the rest equal to zero, the resulting two-variable polynomial is equivalent to the coboundary polynomial (and therefore the Tutte polynomial). See [B-O] for more details.

THEOREM 8. *For any graph G ,*

$$X_G(t) = \sum_{S \subseteq E(G)} t^{|S|} p_{\lambda(S)}.$$

(Remark: here $\lambda(S)$ is as in [St2, Theorem 2.5] or as in the proof of Proposition 25.)

Proof. Fix $k \geq 0$ and consider the coefficient of t^k in $X_G(t)$. The maps κ that contribute to this coefficient are those with $m(\kappa) \geq k$. Each such map contributes

$$\binom{m(\kappa)}{k} x^\kappa.$$

Regard the binomial coefficient here as choosing k of the $m(\kappa)$ monochromatic edges. Then we see that we can sum these contributions in another way: first list all k -subsets $S \subset E(G)$, and then for each such S , consider all maps κ that make every edge in S monochromatic. A moment's thought reveals that this sum is precisely $p_{\lambda(S)}$, and the theorem follows. ■

COROLLARY 4. *Let G be a forest. Then the coefficient of each power sum in the power sum expansion of $X_G(t)$ is a monomial in t .*

Proof. If $S \subset E(G)$, let $G(S)$ be the subgraph of G with vertex set $V(G)$ and edge set S . Fix any integer partition λ . Since G is a forest, $G(S)$ is also a forest for any subset S , and moreover the number of edges of S equals the number of vertices of G minus the number of connected components of $G(S)$, i.e.,

$$|S| = |V(G)| - \ell(\lambda(S)).$$

This means that in the sum in Theorem 8, every term in p_λ has a coefficient of $t^{|V(G)| - \ell(\lambda)}$, independent of S . This proves the corollary. ■

THEOREM 9. *For any graph G ,*

$$X_G(t) = \sum_{\pi, \sigma \in L_G} (1+t)^{n(\pi)} \mu(\pi, \sigma) p_\sigma,$$

where L_G is the lattice of contractions of G and $n(\pi)$ is the number of edges whose endvertices lie in the same block of π .

Proof. For each $\sigma \in L_G$, define

$$X_\sigma = \sum_{\kappa} x^\kappa,$$

where the sum is over all maps such that the monochromatic edges are precisely those edges with both endvertices in the same block of σ . This is the same definition that Stanley makes in his proof of the Möbius function formula for X_G [St2, Theorem 2.6]; as Stanley observes, every map κ belongs to exactly one X_σ , so

$$p_\pi = \sum_{\sigma \geq \pi} X_\sigma$$

for all $\pi \in L_G$. By Möbius inversion,

$$X_\pi = \sum_{\sigma \geq \pi} \mu(\pi, \sigma) p_\sigma.$$

But clearly

$$X_G(t) = \sum_{\pi \in L_G} (1+t)^{n(\pi)} X_\pi,$$

and the theorem follows. ■


6. COUNTEREXAMPLES

We conclude by mentioning two counterexamples. In [St2] it is shown that X_G does not determine G . One might wonder what properties of G the chromatic symmetric function *does* determine. Does, for example, X_G determine whether or not G is planar? The answer is no, because the graphs



have the same chromatic symmetric function but the first graph is not planar (it is a subdivision of K_5) while the second one is.

Another question might be whether $X_{G(P)}$ determines the dimension of the poset P . (Recall that the dimension of a poset is the minimum number of totally ordered sets needed to express the poset as an intersection of totally ordered sets—see [Tro].) Again, the answer is no; for example, the incomparability graphs of the posets with Hasse diagrams



have the same chromatic symmetric function, but the first poset has dimension 2 while the second poset has dimension 3.

Bibliography

- [Big] N. BIGGS, “Algebraic Graph Theory,” second edition, Cambridge University Press, Cambridge, 1993.
- [Bon] J. A. BONDY, A graph reconstructor’s manual, *in* “Surveys in Combinatorics, 1991” (Guildford, 1991), London Math. Soc. Lecture Note Ser. 166, Cambridge Univ. Press, Cambridge, 1991, pp. 221–252.
- [Bre] F. BRENTI, Expansions of chromatic polynomials and log-concavity, *Trans. Amer. Math. Soc.* **332** (1992), 729–756.
- [B-O] T. BRYLAWSKI AND J. OXLEY, The Tutte polynomial and its applications, *in* “Matroid Applications” (N. White, Ed.), Encyclopedia of Mathematics and Its Applications, Vol. 40, Cambridge University Press, Cambridge, 1992, pp. 123–225.
- [BEGW] J. BUHLER, D. EISENBUD, R. GRAHAM, AND C. WRIGHT, Juggling drops and descents, *Amer. Math. Monthly* **101** (1994), 507–519.
- [B-G] J. BUHLER AND R. GRAHAM, A note on the binomial drop polynomial of a poset, *J. Combin. Theory (A)* **66** (1994), 321–326.
- [CSV] L. CARLITZ, R. SCOVILLE, AND T. VAUGHAN, Enumeration of pairs of sequences by rises, falls and levels, *Manuscripta Math.* **19** (1976), 211–243.
- [Cho] T. CHOW, The path-cycle symmetric function of a digraph, *Advances in Math.*, in press.

- [CG1] F. CHUNG AND R. GRAHAM, On digraph polynomials, preprint dated December 3, 1993.
- [CG2] F. CHUNG AND R. GRAHAM, On the cover polynomial of a digraph, *J. Combin. Theory (B)*, in press.
- [D-K] F. N. DAVID AND M. G. KENDALL, Tables of symmetric functions. I, *Biometr.* **36** (1949), 431–449.
- [Dou] P. DOUBILET, On the foundations of combinatorial theory. VII: Symmetric functions through the theory of distribution and occupancy, *Studies in Applied Math.* **51** (1972), 377–396.
- [Dwo] M. DWORKIN, personal communication.
- [GV1] E. GANSNER AND K.-P. VO, The chromatic generating function, *Lin. and Multilin. Alg.* **22** (1987), 87–93.
- [G-J] M. R. GAREY AND D. S. JOHNSON, “Computers and Intractability: A Guide to the Theory of NP-Completeness,” W. H. Freeman and Co., San Francisco, CA, 1979.
- [Ga1] V. GASHAROV, Incomparability graphs of $(\mathbf{3} + \mathbf{1})$ -free posets are s -positive, in Abstracts of the Sixth Conference on Formal Power Series and Combinatorics, May 23–27, 1994.
- [Ga2] V. GASHAROV, On Stanley’s chromatic symmetric function and stable set polynomials, preprint dated January 8, 1995.
- [Ge1] I. M. GESSEL, personal communication.
- [Ge2] I. M. GESSEL, Multipartite P -partitions and inner products of skew Schur functions, in “Combinatorics and Algebra” (C. Greene, Ed.), Contemporary Mathematics Series, Vol. 34, Amer. Math. Soc., Providence, R.I., 1984, pp. 289–301.
- [GV2] I. M. GESSEL AND G. X. VIENNOT, Determinants, paths, and plane partitions, preprint dated July 28, 1989.
- [Gol] J. R. GOLDMAN, personal communication.
- [GJW1] J. R. GOLDMAN, J. T. JOICHI, AND D. E. WHITE, Rook theory I. Rook equivalence of Ferrers boards, *Proc. Amer. Math. Soc.* **52** (1975), 485–492.

- [GJRW] J. R. GOLDMAN, J. T. JOICHI, D. L. REINER, AND D. E. WHITE, Rook theory. II: Boards of binomial type, *SIAM J. Applied Math.* **31** (1976), 618–633.
- [GJW2] J. R. GOLDMAN, J. T. JOICHI, AND D. E. WHITE, Rook theory III. Rook polynomials and the chromatic structure of graphs, *J. Combin. Theory (B)* **25** (1978), 135–142.
- [GJW3] J. R. GOLDMAN, J. T. JOICHI, AND D. E. WHITE, Rook theory—IV. Orthogonal sequences of rook polynomials, *Studies in Applied Math.* **56** (1977), 267–272.
- [GJW4] J. R. GOLDMAN, J. T. JOICHI, AND D. E. WHITE, Rook theory. V. Rook polynomials, Möbius inversion and the umbral calculus, *J. Combin. Theory (A)* **21** (1976), 230–239.
- [Gre] C. GREENE, Proof of a conjecture on immanants of the Jacobi-Trudi matrix, *Lin. Alg. Appl.* **171** (1992), 65–79.
- [Kre] G. KREWERAS, The number of more or less “regular” permutations, *Fibonacci Quart.* **18** (1980), 226–229.
- [Lin] N. LINIAL, Graph coloring and monotone functions on posets, *Disc. Math.* **58** (1986), 97–98.
- [Mac] I. G. MACDONALD, “Symmetric Functions and Hall Polynomials,” Oxford University Press, Oxford, 1979.
- [Rio] J. RIORDAN, “An Introduction to Combinatorial Analysis,” John Wiley & Sons, New York, 1958.
- [R-R] S. M. ROMAN AND G.-C. ROTA, The umbral calculus, *Advances in Math.* **27** (1978), 95–188.
- [R-P] M. RUMNEY AND E. J. F. PRIMROSE, A sequence connected with the sub-factorial sequence (Mathematical notes 3207), *Math. Gazette* **52** (1968), 381–382.
- [Slo] N. J. A. SLOANE, “The On-Line Encyclopedia of Integer Sequences.” [For information on accessing this encyclopedia, see N. J. A. SLOANE, An on-line version of the encyclopedia of integer sequences, *Electronic J. Combin.* **1** (1994), F1.]
- [St1] R. P. STANLEY, personal communication.

- [St2] R. P. STANLEY, A symmetric function generalization of the chromatic polynomial of a graph, *Advances in Math.* **111** (1995), 166–194.
- [St3] R. P. STANLEY, “Enumerative Combinatorics,” vol. 1, Wadsworth & Brooks/Cole, Pacific Grove, CA, 1986.
- [St4] R. P. STANLEY, Graph colorings and related symmetric functions: ideas and applications, preprint.
- [St5] R. P. STANLEY, Ordered structures and partitions, *Memoirs Amer. Math. Soc.* **119** (1972).
- [S-S] R. P. STANLEY AND J. STEMBRIDGE, On immanants of Jacobi-Trudi matrices and permutations with restricted position, *J. Combin. Theory (A)* **62** (1993), 261–279.
- [Tro] W. T. TROTTER, “Combinatorics and Partially Ordered Sets: Dimension Theory,” Johns Hopkins University Press, Baltimore, MD, 1992.
- [Tu1] W. T. TUTTE, All the king’s horses (a guide to reconstruction), in “Graph Theory and Related Topics” (J. A. Bondy and U. S. R. Murty, Eds.), Proceedings of the conference held in honor of Professor W. T. Tutte on the occasion of his sixtieth birthday, Waterloo, July 5–9, 1977, Academic Press, New York, NY, 1979, pp. 15–33.
- [Tu2] W. T. TUTTE, The reconstruction problem in graph theory, *Brit. Polymer J.*, September 1977, 180–183.
- [Vo] K.-P. VO, Graph colorings and acyclic orientations, *Lin. and Multilin. Alg.* **22** (1987), 161–170.