On the chromatic symmetric functions of trees

Michael Tang

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Ricky Ini Liu, Chair

Sara Billey

James Zhang

Program Authorized to Offer Degree:

Mathematics

University of Washington

Abstract

On the chromatic symmetric functions of trees

Michael Tang

Chair of the Supervisory Committee:

Ricky Ini Liu

Department of Mathematics

The *chromatic symmetric function* (CSF) of a graph, introduced by Richard Stanley in 1995, is a symmetric function generalization of the chromatic polynomial which enumerates proper colorings by the number of uses of each color. We attack a long-standing open question of Stanley that asks whether trees are distinguished by their CSFs. We prove new structural results about the CSFs of well-known subclasses of trees, including an involution on symmetric functions that swaps the CSFs of caterpillars in pairs, as well as a characterization of all linear relations between the CSFs of spiders. Then, we study a related graph invariant, the *generalized degree polynomial* (GDP), introduced by Crew and shown to be determined by the CSF by Aliste-Prieto et al. We present several classes of data about a tree that can be recovered from its GDP (and thus also from its CSF), such as the double-degree sequence and leaf adjacency sequence; this extends previous work of Martin, Morin, and Wagner. Then, we consider vector-valued and matrix-valued variants of the GDP for trees with distinguished vertices. By using linear relations that these GDPs satisfy, we prove new recurrence relations for the ordinary GDP and give several constructions that produce

families of ordinary trees with the same GDP. In addition, we prove in some cases that the trees produced by such constructions have different CSFs. Our work suggests a program for further work using the GDP as an intermediary.

CONTENTS

ACKNOWLEDGMENTS

First, I must thank my advisor, Ricky Liu. The work collected in this document would not exist without his guidance, both on mathematics and on comma placement. I also thank the other members of my committee: Sara Billey, James Zhang, and Heather Wilber.

I thank all the other professors, both at the University of Washington and elsewhere, who welcomed me into the mathematical community with open arms: Sara Billey, Isabella Novik, Gaku Liu, Jonah Ostroff, Natalie Naehrig, Matt Conroy, José Aliste-Prieto, Jeremy Martin, Corrine Yap, Vic Reiner, Vasu Tewari, and many more. I also thank all the extraordinarily dedicated and insightful teachers I have had the privilege of knowing throughout my life: Noah Franske, Michael Roddy, Jo Anne Link, and Andy Lawrie, among others.

I thank all my fellow graduate students in the math department at UW for their camaraderie and kindness: in particular, Tracy, Alex, Natasha, Grace, Clare, Cordelia, Yirong, Haoming, Herman, and Nelson. (Protect our unions! Pass De Morgan's laws!)

I thank John Darnielle, Rich Burlew, and Shu Takumi. Please keep doing what you do.

I thank all my teammates on NES for their wit, cleverness, and commitment to puzzling.

I thank all my friends for brightening the sunny days and sharing their umbrellas on the stormy ones. In particular, I thank Jenna and Max for always being at my side.

Finally, I thank my parents and my family for everything. I love you.

DEDICATION

For Tim, Otto, Siggy, and Pluto

Let me not think of my work only as a stepping stone to something else.

And if it is, let me become fascinated with the shape of the stone.

Ze Frank

Chapter 1

INTRODUCTION

This thesis studies a particular problem in *graph coloring*, a subfield of graph theory that is concerned with methods of assigning labels (called "colors") to the vertices or edges of a graph so that certain conditions are satisfied. By far the oldest and most famous graph coloring question is the *four-color problem*, first posed in 1852 but not settled until 1976. Over its long history, this problem pioneered the study of graph coloring and inspired a great deal of discoveries therein; graph coloring remains an extremely active area of research today.

A prominent development inspired by the four-color problem was the introduction by Birkhoff in 1912 of the *chromatic polynomial* χ_G of a graph G. For a positive integer k, the value of $\chi_G(k)$ is the number of colorings of the vertices of G with k colors (usually taken to be the integers $1, 2, \ldots, k$) which are *proper*, meaning that adjacent vertices always receive different colors. While it was not particularly useful in solving the four-color problem, the chromatic polynomial has proved fascinating in its own right. Furthermore, it directly inspired the titular object of this thesis, a symmetric function generalization of the chromatic polynomial.

A *symmetric function* is a formal power series (that is, an infinite weighted sum of monomials $x_{i_1}^{\alpha_1} \cdots x_{i_k}^{\alpha_k}$ in an infinite list of variables $x_1, x_2, \ldots$) of bounded degree with the property that it is symmetric in its variables, i.e., it is fixed under all permutations of the x_i. Examples of symmetric functions include the *power sum symmetric functions* $p_n = x_1^n + x_2^n + \cdots$ and the *elementary symmetric functions* $e_n = \sum_{i_1 < i_2 < \ldots < i_n} x_{i_1} x_{i_2} \cdots x_{i_n}$ (the sum of all products of n distinct variables).

The study of symmetric functions also has a long history because of its connection with the solutions of polynomial equations, observed as early as the 16th century [14]. In modern

terms, we view symmetric functions as having a great deal of abstract algebraic structure: in particular, the set of symmetric functions is an important example of a *Hopf algebra*, which means that it is also an algebra, a coalgebra, a ring, and a vector space. For this and many other reasons, symmetric functions are ubiquitous in mathematics (in particular, in algebra, combinatorics, and representation theory) and even appear in such far-flung realms as quantum mechanics.

Thus, this thesis is about graph coloring and symmetric functions: in particular, it is about a particular symmetric function defined using graph colorings. Let $G = (V, E)$ be a graph. The *chromatic symmetric function* $\mathbf{X}_G$, introduced by Richard Stanley in 1995 [31], is a symmetric function that enumerates proper colorings of G by the number of occurrences of each color. Formally, we have

$$\mathbf{X}_G(x_1, x_2, \dots) = \sum_{\kappa} \prod_{v \in V} x_{\kappa(v)},$$

where κ ranges over all proper colorings of G with positive integers. Thus, the coefficient of a monomial $x_1^{\alpha_1} x_2^{\alpha_2} \cdots$ in $\mathbf{X}_G$ is the number of proper colorings of G such that for each i, the color i is used exactly α_i times. This means that $\mathbf{X}_G$ provides much more information about a graph than the chromatic polynomial χ_G: in particular, the value of $\chi_G(k)$ can be recovered from $\mathbf{X}_G(x_1, x_2, \dots)$ by setting $x_1 = \ldots = x_k = 1$ and $x_i = 0$ for all $i > k$.

In the thirty years since its introduction, the chromatic symmetric function has inspired a great deal of research in algebraic combinatorics. Variants of the chromatic symmetric function for weighted graphs [12], for rooted graphs [23], and in quasisymmetric functions [30] have been developed; the latter is related to the cohomology of certain subvarieties of the flag variety called *Hessenberg varieties*. In addition, the chromatic symmetric function was shown to be a natural object associated to a graph in the sense that the mapping $G \mapsto \mathbf{X}_G$ defines a canonical morphism of *combinatorial Hopf algebras* [1]. Finally, the famous *Stanley-Stembridge conjecture* [34], recently resolved by Hikita [19], states that the chromatic symmetric functions of certain graphs called *unit interval graphs* have positive

expansions in the basis of elementary symmetric functions.

This thesis was inspired by the following question, which was posed by Stanley in the same paper [31] introducing the chromatic symmetric function:

Question 1.1. Does the chromatic symmetric function distinguish trees? In other words, must non-isomorphic trees have different chromatic symmetric functions?

To get some context for this question, note that by contrast, the chromatic polynomial is nearly useless for distinguishing trees: indeed, all trees T with n vertices have the same chromatic polynomial $\chi_T(k) = k(k-1)^n$. Meanwhile, Stanley observed that the chromatic symmetric function does not distinguish among all graphs; the smallest non-isomorphic graphs with the same chromatic symmetric function have just five vertices [31, Figure 1]. Later, Orellana and Scott showed that there are infinitely many pairs of connected graphs with exactly one cycle that have the same chromatic symmetric function [26]. However, Question 1.1 remains open: no one has discovered an example of non-isomorphic trees with the same chromatic symmetric function or found a way to recover a tree from its chromatic symmetric function.

Though Question 1.1 is unresolved, some interesting progress has been made towards answering it. One common approach is to identify or construct special classes of trees whose members are distinguished by their chromatic symmetric functions, which is a strategy used in [3, 15, 17, 22, 35]. This approach essentially resolves a restricted version of Stanley's question, and one can hope that the methods used to do so generalize beyond the particular class of trees being studied. Another common approach is to begin with the chromatic symmetric function $\mathbf{X}_T$ of an arbitrary tree T and extract information about T from the coefficients of $\mathbf{X}_T$. For instance, in 2008, Martin, Morin, and Wagner [25] proved that $\mathbf{X}_T$ determines the *degree sequence* of T (the number of vertices of T of each degree) and the *path sequence* of T (the number of paths in T of each length). As another example, Crew [11] showed that $\mathbf{X}_T$ determines the lengths of the *twigs* of T, which are the maximal paths in T consisting of a leaf and some number of degree-2 vertices. This approach shows that

trees with the same chromatic symmetric function must also agree on a wide range of more concrete statistics. One can hope that all this information recovered from the chromatic symmetric function, when taken together, will be enough to reconstruct the tree itself, or else to narrow down the search for a counterexample.

While this thesis takes inspiration from both approaches, it mainly adopts the latter. In particular, it focuses on the relationship between the chromatic symmetric function and other *graph invariants*. Indeed, the above examples of information about T recovered from $\mathbf{X}_T$ were not actually derived directly from the chromatic symmetric function; rather, the authors of [25] and [11] employed an intermediary called the *subtree polynomial*. Defined in [13], the subtree polynomial is a two-variable invariant of a tree that enumerates subtrees by their size and number of leaf edges. It was shown in [25] that the chromatic symmetric function determines the subtree polynomial: thus, any information that one can determine about a tree from its subtree polynomial is also information determined by the chromatic symmetric function. It was the subtree polynomial, whose coefficients are easier to understand and manipulate, that was used to recover the degree sequence, path sequence, and twig sequence of a tree.

In this thesis, we focus on a particular three-variable tree invariant called the *generalized degree polynomial*. Introduced by Crew in 2020 [11], the generalized degree polynomial $\mathrm{G}_T(x, y, z)$ of a tree T enumerates subsets A of the vertices of T by size as well as number of internal edges (edges of T with both endpoints in A) and boundary edges (edges of T with exactly one endpoint in A). Crew conjectured in [11] that $\mathbf{X}_T$ determines G_T, and this was later proven by Aliste-Prieto, Martin, Wagner, and Zamora [4]. We devote much of this work to examining information about a tree that can be recovered from its generalized degree polynomial, as well as pairs of trees that have the same generalized degree polynomial. In doing so, we make progress towards a better understanding of Stanley's question.

This thesis[1] comprises six chapters, the first of which you are currently reading. In

[1]We note that this thesis consists largely of joint work with Ricky Liu. In particular, much of the content of Chapter 4 also appears in the paper [21].

Chapter 2, we first introduce notations and conventions, and then present some background about symmetric functions and the CSF. In particular, we note that the CSF satisfies a *deletion–near-contraction recurrence* (Proposition 2.7).

In Chapter 3, we study the so-called *chromatic bases*, namely the *star basis* $\{\mathrm{st}_\lambda\}$ and *path basis* $\{\mathrm{pa}_\lambda\}$. These are linear bases for the algebra of symmetric functions defined using the CSFs of star and path graphs, respectively. We derive a new structural result regarding these bases (Theorem 3.6), namely, that there is an involution on symmetric functions that sends st_λ to $-\mathrm{pa}_\lambda$ for all partitions λ. Furthermore, we show that this involution swaps the CSFs of complementary pairs of *caterpillar* graphs (up to sign). Next, we use chromatic bases to study the chromatic symmetric functions of *spider* graphs. We expand the CSF of a spider graph in the path basis $\{\mathrm{pa}_\lambda\}$ (Theorem 3.12), and in doing so, we discover a connection between the chromatic symmetric functions of spiders and the *Bernstein creation operators* (Theorem 3.24). This allows us to compute the dimension of the span of $\mathbf{X}_S$ where S ranges over all spiders of a fixed order (Theorem 3.25), as well as characterize all linear relations between the $\mathbf{X}_S$ (Theorem 3.28 and Theorem 3.30).

The remaining chapters focus on the generalized degree polynomial (which we abbreviate as the GDP). In Chapter 4, we define the generalized degree polynomial of a tree and give a concise proof that it is determined by the chromatic symmetric function (Proposition 4.14) by using Hopf algebraic methods. Then, we present several kinds of information about a tree that is determined by its GDP (Theorem 4.20), organized around the notion of *degree embeddings* of trees. For example, we show that the GDP determines the *double-degree sequence* and *leaf-adjacency sequence* of a tree, which are both refinements of the degree sequence. We also consider a trio of two-variable specializations of the GDP, one of which is the subtree polynomial mentioned above; we prove a new formula relating two of these specializations (Theorem 4.3) and resolve a related conjecture of Martin, Morin, and Wagner (Theorem 4.5). Finally, we define a sequence of further generalizations of the GDP (Definition 4.30) and show that they are all determined by the CSF.

In Chapter 5, we develop the theory of the generalized degree polynomial for rooted graphs and *polarized graphs* (graphs with two distinguished vertices, a *left pole* and a *right pole*). These variants of the GDP are vector-valued and matrix-valued, respectively, and unlike the ordinary generalized degree polynomial, we can show that they are complete invariants for rooted and polarized trees (Corollary 5.5 and Theorem 5.34). Despite this, studying the GDPs of rooted and polarized trees can still tell us a great deal about the GDPs of ordinary trees. Specifically, by using linear relations that exist between polarized GDPs, as well as an operation for combining rooted and polarized trees called *concatenation* (Definition 5.19), we prove two recurrence relations that can be used to efficiently compute the ordinary GDP (Theorem 5.38 and Theorem 5.40). We also prove some surprising divisibility facts about the GDP, such as Theorem 5.47.

Finally, in Chapter 6, we apply the machinery developed in the previous chapter to give two general frameworks for constructing pairs (and larger families) of ordinary trees with the same GDP: the *reversal phenomenon* (Section 6.1) and the *switching phenomenon* (Section 6.2). We prove in some cases that these frameworks produce trees that have different CSFs (Section 6.3). In addition, by analyzing computer calculations, we show that these frameworks recover nearly all pairs of trees with order at most 23 that have the same GDP (Section 6.4). Our work suggests a two-part program to attack Stanley's question: first, characterize all pairs of ordinary trees with the same GDP, and then decide whether the trees in each pair have different chromatic symmetric functions. This thesis makes partial progress on executing both steps.

Chapter 2

PRELIMINARIES

In this chapter, we present notation, conventions, and background information about the topics studied in this thesis. Readers who are already familiar with the chromatic symmetric function may skip to Chapter 3, referencing the contents of this chapter as needed.

2.1 Conventions and basic definitions

The cardinality of a set S will be denoted $|S|$. The notation $S' \subseteq S$ means that S' is a subset of S, not necessarily proper. Given a nonnegative integer k, we write $\binom{S}{k}$ for the collection of all subsets of S with cardinality k.

We write $\mathbb{Z}$ for the set of integers and $\mathbb{Z}_{>0}$ for the set of positive integers. Given a nonnegative integer n, we write $[n]$ to mean the set $\{1, 2, \ldots, n\}$ (where $[0] = \varnothing$ is the empty set). For nonnegative integers a and b, we write $[a, b]$ to mean the set $\{a, a+1, \ldots, b\}$.

In general, the notation $[m]f$ means the coefficient of m in f. Here f is an object such as a polynomial (with m a monomial), a power series (with m a monomial), or an element of a vector space expanded in a particular basis (with m an element of that basis). The exact meaning should be clear from context.

Given objects i and j, the Kronecker delta δ_{ij} equals 1 if $i = j$ and 0 otherwise.

2.1.1 Graphs

We assume familiarity with the basic concepts of graph theory. All graphs are assumed to be simple (no parallel edges or self-loops) and finite. Graphs will generally be unlabeled and considered only up to isomorphism.

Given a graph G, we write $V(G)$ and $E(G)$ for its vertex set and edge set, respectively.

The *order* of G, denoted $|G|$, is the number of vertices of G. We denote the number of connected components of G by $c(G)$.

Given a vertex $v \in V(G)$, we denote its degree by $\deg_G(v)$, or just $\deg(v)$ if G is clear from context. A vertex v is called an *internal vertex* if $\deg(v) \geq 2$, and a *leaf vertex* if $\deg(v) = 1$. A *leaf edge* of G is an edge incident to at least one leaf of G. For a subset $A \subseteq V(G)$, we write $G|A$ to mean the induced subgraph of G with vertex set A.

2.1.2 Compositions

Here we give some definitions related to integer compositions and partitions, which we will use to discuss symmetric functions in the next section.

An *integer composition*, or *composition* for short, is a finite sequence $\alpha = (\alpha_1, \ldots, \alpha_k)$ of positive integers. Each α_i is called a *part* of α, and the *length* of α is the number of parts (in this case, k). The *size* of α, denoted $|\alpha|$, is the sum of its parts: $|\alpha| = \alpha_1 + \cdots + \alpha_k$. If $|\alpha| = n$, we say that α is a *composition of* n and write $\alpha \vDash n$. We write Comp for the set of all compositions and Comp_n for the set of all compositions of n.

There is a unique composition with no parts, which thus has length 0 and size 0; we denote this by 0. We also write positive integers interchangeably with the corresponding compositions of length 1.

When writing explicit examples of compositions, we often omit commas and parentheses. Thus, for instance, the composition $(3, 1, 3, 4)$ would be abbreviated as 3134. In addition, we use the notation a^k to mean the part a repeated k times. For instance, the composition $(1, 1, \ldots, 1, 2)$ with k occurrences of 1 would be abbreviated 1^k2.

It is well-known that the number of compositions of n is 2^{n-1} for each $n \geq 1$. This is because there is a bijection between compositions of n and subsets of $[n-1]$ given by

$$\alpha = (\alpha_1, \ldots, \alpha_k) \longmapsto \mathrm{Set}(\alpha) = \{\alpha_1,\ \alpha_1 + \alpha_2,\ \ldots,\ \alpha_1 + \cdots + \alpha_{k-1}\}.$$

Conversely, if $\mathrm{Set}(\alpha) = \{r_1 < r_2 < \cdots < r_\ell\}$, then $\alpha = (r_1, r_2 - r_1, \ldots, r_\ell - r_{\ell-1}, n - r_\ell)$.

The *complement* of a composition $\alpha \vDash n$, denoted α^c, is the composition of n such that

$$\mathrm{Set}(\alpha^c) = [n-1] \setminus \mathrm{Set}(\alpha).$$

Thus, complementation is an involution on compositions of n.

Example 2.1. Let $\alpha = 32114$, a composition of $n = 11$. Then $\mathrm{Set}(\alpha) = \{3, 5, 6, 7\}$ and so $\mathrm{Set}(\alpha^c) = \{1, 2, 4, 8, 9, 10\}$. Hence, $\alpha^c = 1124111$. $\Diamond$

The *reverse* of a composition $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$ is $\alpha^* = (\alpha_k, \ldots, \alpha_2, \alpha_1)$; we say that α is a *palindrome* (or is *palindromic*) if $\alpha = \alpha^*$.

Given two compositions $\alpha = (\alpha_1, \ldots, \alpha_k)$ and $\beta = (\beta_1, \ldots, \beta_\ell)$, their *concatenation* $\alpha\beta$ is defined by

$$\alpha\beta = (\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_\ell),$$

and their *near-concatenation* $\alpha \odot \beta$ is defined (when $\alpha, \beta \neq 0$) by

$$\alpha \odot \beta = (\alpha_1, \ldots, \alpha_{k-1}, \alpha_k + \beta_1, \beta_2, \ldots, \beta_\ell).$$

One can check that complementation, concatenation, and near-concatenation satisfy

$$(\alpha\beta)^c = \alpha^c \odot \beta^c \tag{2.1}$$

for all $\alpha, \beta \neq 0$.

Given a composition α and $k \in \mathbb{Z}_{>0}$, we write $m_k(\alpha) = |\{i : \alpha_i = k\}|$, the number of parts of α equal to k. We write $\alpha \sim \beta$ if and only if the parts of α can be rearranged to form β; equivalently, $\alpha \sim \beta$ if $m_k(\alpha) = m_k(\beta)$ for all k.

We define

$$\mathrm{aut}(\alpha) = m_1(\alpha)!\, m_2(\alpha)! \cdots, \qquad \mathrm{perm}(\alpha) = \frac{\ell(\alpha)!}{\mathrm{aut}(\alpha)},$$
$$\mathrm{prod}(\alpha) = \alpha_1 \cdots \alpha_{\ell(\alpha)}, \qquad z_\alpha = \mathrm{aut}(\alpha)\,\mathrm{prod}(\alpha).$$

The number $\mathrm{aut}(\alpha)$ can be interpreted as the number of ways to "permute" equal parts of α, i.e., the number of permutations $(w_1, \ldots, w_{\ell(\alpha)})$ of $[\ell(\alpha)]$ such that $(\alpha_{w_1}, \ldots, \alpha_{w_{\ell(\alpha)}}) = \alpha$. By the orbit-stabilizer theorem, $\mathrm{perm}(\alpha)$ is the number of compositions β with $\alpha \sim \beta$, i.e., the number of distinct permutations of the parts of α.

Example 2.2. Again, let $\alpha = 32114$. Then $\mathrm{aut}(\alpha) = 2!(1!)^3 = 2$, so $\mathrm{perm}(\alpha) = 5!/2 = 60$. We also have $\mathrm{prod}(\alpha) = 3 \cdot 2 \cdot 1 \cdot 1 \cdot 4 = 24$ and $z_\alpha = 2 \cdot 24 = 48$. ◊

2.1.3 Partitions

An (integer) *partition* is a composition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$ whose parts are in weakly decreasing order: $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_\ell$. A *partition of* n is a partition λ with size n; we write $\lambda \vdash n$ to signify this. Let Par be the set of all partitions and Par_n the set of all partitions of n. It will sometimes be convenient to consider partitions to be padded on the right by infinitely many zeros (so that a partition is an infinite, weakly decreasing sequence of nonnegative integers whose sum is finite).

For any composition α, its parts can be rearranged in weakly decreasing order to form a partition. We denote this partition by $\tilde{\alpha}$, so that $\alpha \sim \tilde{\alpha}$.

Given a partition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, its *Young diagram* is a grid of boxes with ℓ rows, such that row i from the top has λ_i boxes, all left-aligned. Then, the *conjugate* of λ is the partition λ' whose Young diagram is obtained by reflecting the Young diagram of λ over the main diagonal (top-left to bottom-right). Algebraically, we have $\lambda' = (\lambda'_1, \ldots, \lambda'_k)$ where

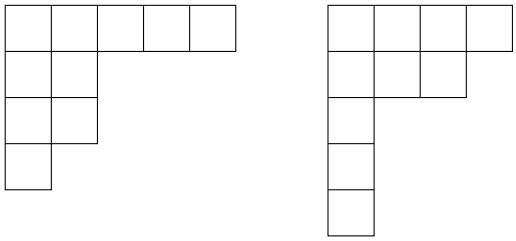

$\lambda = 5221$ $\qquad$ $\lambda' = 43111$

Figure 2.1: The Young diagrams of two conjugate partitions.

$k = \lambda_1$ and $\lambda'_j = \max\{i : \lambda_i \geq j\}$. This can be stated in an equivalent form as

$$\lambda'_j = a \qquad \text{if and only if} \qquad \lambda_a \geq j > \lambda_{a+1}.$$

Example 2.3. The partition $\lambda = 5221$ has conjugate $\lambda' = 43111$. The Young diagrams of λ and λ' are depicted in Figure 2.1. $\Diamond$

2.2 Symmetric functions

We now recall some parts of the theory of symmetric functions. Our treatment is extremely terse, covering only what we will need; a detailed account can be found in e.g. [18, 24, 33].

We work over the field of rational numbers $\mathbb{Q}$. Let $x_1, x_2, \ldots$ be an infinite list of commuting indeterminates. A *symmetric function* is a formal power series $f \in \mathbb{Q}[[x_1, x_2, \ldots]]$ of bounded degree with the property that f is invariant under all permutations of the x_i. The set of all symmetric functions forms a $\mathbb{Q}$-algebra Sym under the usual addition and multiplication of power series. In addition, Sym has a grading $\mathrm{Sym} = \bigoplus_{n \geq 0} \mathrm{Sym}_n$, where Sym_n is the set of all homogeneous symmetric functions of degree n.

The algebra Sym has several well-known bases, indexed by partitions $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, which will play important roles throughout this work. (Each basis respects the grading on Sym; the degree of a basis element equals the size of the corresponding partition.) The most obvious basis is the *monomial basis* $\{m_\lambda\}$, where m_λ is the sum of all distinct monomials

$x_{i_1}^{\alpha_1} \cdots x_{i_\ell}^{\alpha_\ell}$ such that $(\alpha_1, \ldots, \alpha_\ell) \sim \lambda$. (This is not to be confused with the notation $m_k(\alpha)$ for compositions α as defined in the previous section.) Next, for positive integers n, let

$$\begin{aligned} p_n &= m_n = x_1^n + x_2^n + \cdots, \\ e_n &= m_{1^n} = \sum_{i_1 < \ldots < i_\ell} x_{i_1} \cdots x_{i_\ell}, \\ h_n &= \sum_{\lambda \vdash n} m_n = \sum_{i_1 \leq \ldots \leq i_\ell} x_{i_1} \cdots x_{i_\ell}. \end{aligned}$$

We extend these definitions multiplicatively: given a partition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, we have $p_\lambda = p_{\lambda_1} \cdots p_{\lambda_\ell}$, and similarly for e_λ and h_λ. This yields the *power sum basis* $\{p_\lambda\}$, the *elementary basis* $\{e_\lambda\}$, and the *(complete) homogeneous basis* $\{h_\lambda\}$. One can show that these indeed form (graded) bases for Sym. By convention, we set $p_0 = e_0 = h_0 = 1$, and $p_n = e_n = h_n = 0$ for integers $n < 0$. We can also define p_α, h_α, and e_α in the same way for compositions α; thus $p_\alpha = p_{\tilde{\alpha}}$ and similarly for the other bases.

Finally, we define the celebrated *Schur basis* $\{s_\lambda\}$. Given a partition λ, a *semistandard Young tableau* of *shape* λ is a filling of the boxes in the Young diagram of λ with positive integers such that the rows weakly increase from left to right, while the columns strictly increase from top to bottom. If T is a semistandard Young tableau, we define $\mathbf{x}^T$ to be the monomial $x_1^{\alpha_1} x_2^{\alpha_2} \cdots$, where α_k is the number of occurrences of k in T. Then

$$s_\lambda = \sum_T \mathbf{x}^T$$

where T ranges over all semistandard Young tableaux of shape λ. (It is not obvious from the definition that s_λ is symmetric, but this can be shown using *Bender-Knuth involutions*.)

The Schur basis can be written in terms of the complete homogeneous basis using the

Jacobi-Trudi formula: for a partition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, we have

$$s_\lambda = \det\left[h_{\lambda_i - i + j}\right]_{1 \le i,j \le \ell} = \det \begin{bmatrix} h_{\lambda_1} & h_{\lambda_1+1} & \cdots & h_{\lambda_1+\ell-1} \\ h_{\lambda_2-1} & h_{\lambda_2} & \cdots & h_{\lambda_2+\ell-2} \\ \vdots & \vdots & \ddots & \vdots \\ h_{\lambda_\ell-\ell+1} & h_{\lambda_\ell-\ell+2} & \cdots & h_{\lambda_\ell} \end{bmatrix}. \tag{2.2}$$

The *dual Jacobi-Trudi formula* can be used to write the Schur basis in the elementary basis:

$$s_\lambda = \det\left[e_{\lambda'_i - i + j}\right]_{1 \le i,j \le \ell} = \det \begin{bmatrix} e_{\lambda'_1} & e_{\lambda'_1+1} & \cdots & e_{\lambda'_1+\ell-1} \\ e_{\lambda'_2-1} & e_{\lambda'_2} & \cdots & e_{\lambda'_2+\ell-2} \\ \vdots & \vdots & \ddots & \vdots \\ e_{\lambda'_\ell-\ell+1} & e_{\lambda'_\ell-\ell+2} & \cdots & e_{\lambda'_\ell} \end{bmatrix}. \tag{2.3}$$

2.2.1 The involution ω and the Hall inner product

We define two important structural elements of the algebra of symmetric functions. Let $\omega\colon \mathrm{Sym} \to \mathrm{Sym}$ be the (graded) algebra endomorphism defined by $\omega(e_\lambda) = h_\lambda$ for all partitions λ. Then in fact, ω satisfies $\omega(h_\lambda) = e_\lambda$ for all λ, so it is an involution on Sym. The involution ω acts on the power sum basis by

$$\omega(p_\lambda) = (-1)^{|\lambda| - \ell(\lambda)} p_\lambda$$

and on the Schur basis by

$$\omega(s_\lambda) = s_{\lambda'}.$$

Next, define a bilinear form $\langle \cdot, \cdot \rangle : \mathrm{Sym} \times \mathrm{Sym} \to \mathbb{Q}$ by

$$\langle m_\lambda, h_\mu \rangle = \delta_{\lambda\mu}$$

for partitions λ and μ, extended linearly in both arguments. It can be shown that this bilinear form is symmetric and positive definite; it is called the *Hall inner product*. In addition, the involution ω is an isometry with respect to $\langle \cdot, \cdot \rangle$, that is, $\langle \omega(f), \omega(g) \rangle = \langle f, g \rangle$ for all $f, g \in \mathrm{Sym}$.

By definition, the monomial basis $\{m_\lambda\}$ and the complete homogeneous basis $\{h_\lambda\}$ are dual with respect to the Hall inner product. In addition, we have the formulas

$$\langle p_\lambda, p_\mu \rangle = z_\lambda \delta_{\lambda\mu} \qquad \text{and} \qquad \langle s_\lambda, s_\mu \rangle = \delta_{\lambda\mu},$$

so the power sum basis is orthogonal and the Schur basis is orthonormal with respect to the Hall inner product.

2.2.2 Sym as a Hopf algebra

The algebra Sym can be given additional structure making it into a *Hopf algebra*. The particular aspect of this that we will need is the *comultiplication* map on Sym, denoted $\Delta \colon \mathrm{Sym} \to \mathrm{Sym} \otimes \mathrm{Sym}$, which is linear and multiplicative. In general, to compute $\Delta(f)$ for $f(x_1, x_2, \dots) \in \mathrm{Sym}$, one can replace the x_i with a "doubly infinite" list of variables $x_1, x_2, \dots, y_1, y_2, \dots$, evaluate f on these arguments, and then interpret the result as an element of $\mathrm{Sym} \otimes \mathrm{Sym}$ (where the x-variables correspond to the first copy of Sym and the y-variables correspond to the second). In practice, though, it suffices to know how to compute Δ on a basis of Sym, or on an algebraic generating set for Sym since Δ is multiplicative. The formulas that we will need are

$$\Delta(p_n) = 1 \otimes p_n + p_n \otimes 1 \tag{2.4}$$

and

$$\Delta(e_n) = \sum_{i+j=n} e_i \otimes e_j = e_0 \otimes e_n + e_1 \otimes e_{n-1} + \cdots + e_n \otimes e_0 \tag{2.5}$$

for $n \in \mathbb{Z}_{>0}$.

We will use *Sweedler notation* for a generic element of the image of Δ. If $f \in \mathrm{Sym}$, then there exist $k \geq 0$ and $f_{1i}, f_{2i} \in \mathrm{Sym}$ for $1 \leq i \leq k$ such that $\Delta(f) = \sum_{i=1}^{k} f_{1i} \otimes f_{2i}$. We abbreviate this to

$$\Delta(f) = \sum f_{(1)} \otimes f_{(2)}. \tag{2.6}$$

Let E be a $\mathbb{Q}$-algebra and $\mu_E \colon E \otimes E \to E$ its multiplication map. For any linear maps $\varphi, \psi \colon \mathrm{Sym} \to E$, the *convolution* of φ and ψ is the map $\varphi * \psi \colon \mathrm{Sym} \to E$ defined by

$$(\varphi * \psi)(f) = (\mu_E \circ (\varphi \otimes \psi) \circ \Delta)(f) = \sum \varphi(f_{(1)})\, \psi(f_{(2)}). \tag{2.7}$$

This is also a linear map. Moreover, it can be shown that if φ and ψ are algebra maps, then so is $\varphi * \psi$, provided that E is commutative.

2.3 The chromatic symmetric function

Finally, we define the central object of this thesis and state some key facts about it.

Let G be a graph and X a nonempty set. A *coloring* of G with colors in X is a function $\kappa \colon V(G) \to X$ that assigns an element of X (called a "color") to each vertex of G. In this thesis, we will always take X to be $\mathbb{Z}_{>0}$ or a subset thereof. A coloring κ is *proper* if adjacent vertices never get the same color: that is, if $\{v, w\} \in E(G)$, then $\kappa(v) \neq \kappa(w)$. Given a coloring κ, we define a monomial $\mathbf{x}^\kappa$ in variables $x_1, x_2, \ldots$ by

$$\mathbf{x}^\kappa = \prod_{v \in V(G)} x_{\kappa(v)} = \prod_{i \geq 1} x_i^{|\kappa^{-1}(i)|}.$$

Definition 2.4 ([31])**.** The *chromatic symmetric function* (CSF) of a graph G, denoted $\mathbf{X}_G$, is defined by

$$\mathbf{X}_G(x_1, x_2, \ldots) = \sum_{\kappa \text{ proper}} \mathbf{x}^\kappa,$$

where κ ranges over all proper colorings of G with colors in $\mathbb{Z}_{>0}$.

Thus, the coefficient of any monomial $x_1^{\alpha_1}x_2^{\alpha_2}\cdots$ in $\mathbf{X}_G$ is the number of proper colorings of G such that each color i is used exactly α_i times. This is invariant under permutations of the numbers α_i, so $\mathbf{X}_G$ is indeed a symmetric function. Furthermore, the definition implies that $\mathbf{X}_G$ is homogeneous of degree $n = |G|$. We also note that $\mathbf{X}_{G_1 \sqcup G_2} = \mathbf{X}_{G_1}\mathbf{X}_{G_2}$ for graphs G_1 and G_2, where $G_1 \sqcup G_2$ is the disjoint union.

The chromatic symmetric function $\mathbf{X}_G$ can be used to recover the classical *chromatic polynomial* χ_G. Recall that χ_G is a polynomial of degree n such that, for each positive integer k, $\chi_G(k)$ is the number of proper colorings of G with colors in $[k]$. For any k, we can take $x_1 = \cdots = x_k = 1$ and $x_i = 0$ for $i > k$ to get

$$\mathbf{X}_G(1^k) = \mathbf{X}_G(\underbrace{1, 1, \ldots, 1}_{k}, 0, 0, \ldots) = \chi_G(k).$$

(Indeed, under this substitution, the term corresponding to a proper coloring κ is sent to 1 if it only uses colors from $[k]$, otherwise 0.) Therefore χ_G can be recovered from $\mathbf{X}_G$.

A common way to study the chromatic symmetric function of a graph is to examine its expansion in one of the many bases of Sym (defined in the previous section). For two particular bases, the monomial basis and the power sum basis, the coefficients that one obtains have combinatorial significance.

First, we consider the monomial basis expansion of $\mathbf{X}_G$. A *stable partition* of a graph G is an (unordered) set partition of $V(G)$ into blocks that are independent sets of G. The *type* of a stable partition is the integer partition of n whose parts are the sizes of the blocks of the stable partition. Given $\lambda \vdash n$, let $a_\lambda(G)$ be the number of stable partitions of G of type λ. Then, from the definition of $\mathbf{X}_G$, we have the formula

$$\mathbf{X}_G = \sum_{\lambda \vdash n} \mathrm{aut}(\lambda) a_\lambda(G) m_\lambda,$$

which expresses $\mathbf{X}_G$ in the monomial basis. It follows that knowing $\mathbf{X}_G$ is equivalent to

knowing all the numbers $\alpha_\lambda(G)$.

Next, in the case of a forest F, we consider the power sum basis expansion of $\mathbf{X}_F$. A *connected partition* of F is a set partition of $V(F)$ into blocks, each of which is the vertex set of a connected subgraph of F. We write $\mathcal{C} \vdash F$ if $\mathcal{C}$ is a connected partition of F. Let $b_\lambda(F)$ be the number of connected partitions of F of type λ (where the type is defined in the same way as for stable partitions). Then we have the following result of Stanley.

Proposition 2.5 ([31, Corollary 2.8])**.** If F is a forest of order n, then

$$\mathbf{X}_F = \sum_{\lambda \vdash n} (-1)^{n-\ell(\lambda)} b_\lambda(F) p_\lambda = \sum_{\mathcal{C} \vdash F} \prod_{A \in \mathcal{C}} (-1)^{|\mathcal{A}|-1} p_{|\mathcal{A}|}.$$

Thus, for a forest F, knowing $\mathbf{X}_F$ is equivalent to knowing all of the numbers $b_\lambda(F)$. The power sum expansion of the CSF will turn out to be more useful for our purposes than the monomial expansion.

Let us state one more nice fact about the CSF.

Proposition 2.6. For any graph G, we have

$$\Delta(\mathbf{X}_G) = \sum_{A \sqcup B = V(G)} \mathbf{X}_{G|A} \otimes \mathbf{X}_{G|B}, \tag{2.8}$$

where A and B ranges over all ways to partition $V(G)$ into two sets.

This fact can be proven directly by using the monomial or power sum expansions of $\mathbf{X}_G$, or by using the theory of *combinatorial Hopf algebras* [1]. In particular, this means that if E is an algebra and $f, g \colon \mathrm{Sym} \to E$ are linear maps, then their convolution satisfies

$$(f * g)(\mathbf{X}_G) = \sum_{A \sqcup B = V(G)} f(\mathbf{X}_{G|A}) \cdot g(\mathbf{X}_{G|B}). \tag{2.9}$$

2.3.1 The deletion–near-contraction recurrence

We now describe a recurrence relation, first presented in [5], that can be used to compute the chromatic symmetric function. To do so, it is helpful to recall the classical *deletion-contraction* relation satisfied by the chromatic polynomial.

Let G be a graph and $e = \{u, v\}$ be an edge of G. The *deletion* of e from G is the graph $G - e$ with vertex set $V(G)$ and edge set $E(G) \setminus \{e\}$. The *contraction* of e in G is the graph G/e obtained from $G - e$ by identifying u and v (and then replacing any parallel edges by a single edge to obtain a simple graph). Then, we have the well-known relation

$$\chi_G(k) = \chi_{G-e}(k) - \chi_{G/e}(k),$$

which can be used to recursively compute the chromatic polynomial of any graph. Many natural graph invariants satisfy a deletion-contraction–type relation; the most general such invariant is the *Tutte polynomial*.

Unfortunately, the deletion-contraction relation does not have a direct analog for the chromatic symmetric function; one reason for this is that G/e has one fewer vertex than G and $G-e$, and so $\mathbf{X}_{G/e}$ does not have the same degree (in Sym) as $\mathbf{X}_G$ and $\mathbf{X}_{G-e}$. However, the CSF does satisfy a similar relation, called the *deletion–near-contraction relation*. If e is an edge of G, the *near-contraction* of e in G is the graph $G \odot e$ obtained from G/e by adding a new leaf adjacent to the identified vertex. (Note that if e is a leaf edge, then $G \odot e$ is isomorphic to G.)

Proposition 2.7 ([5, Proposition 3.1])**.** Let G be a graph and let e be an (non-leaf) edge of G. Then we have

$$\mathbf{X}_G = \mathbf{X}_{G-e} + \mathbf{X}_{G\odot e} - p_1 \mathbf{X}_{G/e}. \tag{2.10}$$

This proposition was proved in [5] by using a weighted variant of the CSF (introduced in [12]), but for completeness, we give a self-contained proof here.

Proof. Given any coloring of a graph, an edge is *monochromatic* if its endpoints are assigned the same color. Thus, a coloring is proper if and only if it has no monochromatic edges.

Let $e = \{u, v\}$. Let A be the set of colorings $\kappa\colon V(G) \to \mathbb{Z}_{>0}$ of G with e as the only monochromatic edge. Then the proper colorings of $G - e$ are exactly the proper colorings of G together with the colorings in A. It follows that

$$\mathbf{X}_{G-e} - \mathbf{X}_G = \sum_{\kappa \in A} \mathbf{x}^\kappa.$$

In $G \odot e$, let u be the identified vertex and let w be the new leaf attached to it. Note that $p_1 \mathbf{X}_{G/e} = \mathbf{X}_{(G/e) \sqcup \{w\}}$ since p_1 is the CSF of a single-vertex graph; furthermore, $G \odot e$ can be obtained from $(G/e) \sqcup \{w\}$ by deleting the edge $\{u, w\}$. Then, by the same logic as before, we have

$$p_1 \mathbf{X}_{G/e} - \mathbf{X}_{G \odot e} = \sum_{\kappa' \in B} \mathbf{x}^{\kappa'},$$

where B is the set of colorings κ' of $G \odot e$ with $\{u, w\}$ as the only monochromatic edge. But there is a bijection between the colorings $\kappa \in A$ and $\kappa' \in B$, where $\kappa(v) = \kappa'(w)$ and κ and κ' agree on all other vertices. This implies that $\mathbf{X}_{G-e} - \mathbf{X}_G = p_1 \mathbf{X}_{G/e} - \mathbf{X}_{G \odot e}$, and the result follows. □

By recursively applying the deletion–near-contraction relation, one can compute $\mathbf{X}_G$ for any graph G. This procedure always terminates because $G - e$ and G/e have fewer edges than G, while $G \odot e$ has the same number of edges but one fewer non-leaf edge than G. The base cases of the procedure are therefore the graphs with no non-leaf edges. Each component of such a graph is a *star graph*, whose chromatic symmetric functions we will see in the next chapter.

Chapter 3

CHROMATIC BASES AND SPECIAL CLASSES OF TREES

In this chapter, we discuss two bases for symmetric functions defined using the CSF, the *star* and *path* bases, as well as two special classes of trees, *caterpillars* and *spiders*. The star basis $\{\mathrm{st}_\lambda\}$ and path basis $\{\mathrm{pa}_\lambda\}$ were introduced in [10] and studied in more detail in [2], where it was shown that for each basis, there is an involution on Sym swapping it with the power sum basis $\{p_\lambda\}$. (This fact has come to be known as *chromatic reciprocity*.) We present a related and seemingly overlooked result, namely, that there is an involution of Sym that swaps st_λ and $-\mathrm{pa}_\lambda$ for all λ. We also show that, up to sign, this involution swaps the CSFs of pairs of "complementary" caterpillars.

Next, we use chromatic bases to study the chromatic symmetric functions of a special class of trees called *spiders*. We give a formula for the expansion of the CSF of a spider in the path basis, which specializes to a formula expressing the star basis in terms of the path basis and vice versa. Then, we consider the subspace of Sym_n spanned by the CSFs of all spiders of order n. We compute the dimension of this space and give a basis for it; along the way, we discover a surprising connection between spiders and the *Bernstein creation operators*. Finally, we characterize all linear relations between the CSFs of spiders.

3.1 Stars, paths, and caterpillars

We begin by defining the star and path bases for Sym.

Let n be a positive integer. The *star graph* of order n, which we denote St_n, is the n-vertex tree in which one vertex is adjacent to all others. The *path graph* of order n, which we denote Pa_n, is the n-vertex tree whose vertices can be labeled $v_1, \ldots, v_n$ such that its edges are $\{v_i, v_{i+1}\}$ for $1 \leq i \leq n-1$.

Definition 3.1. Let $\lambda = (\lambda_1, \ldots, \lambda_k)$ be a partition. The *star symmetric function* st_λ is defined by

$$\mathrm{st}_\lambda = \mathrm{st}_{\lambda_1} \cdots \mathrm{st}_{\lambda_k}$$

where $\mathrm{st}_n = \mathbf{X}_{\mathrm{St}_n}$ is the CSF of the star graph of order n. Similarly, the *path symmetric function* pa_λ is defined by

$$\mathrm{pa}_\lambda = \mathrm{pa}_{\lambda_1} \cdots \mathrm{pa}_{\lambda_k}$$

where $\mathrm{pa}_n = \mathbf{X}_{\mathrm{Pa}_n}$. By convention, $\mathrm{st}_0 = \mathrm{pa}_0 = 1$.

(Like with the other multiplicative bases of Sym, we define $\mathrm{st}_\alpha = \mathrm{st}_{\tilde{\alpha}}$ and $\mathrm{pa}_\alpha = \mathrm{pa}_{\tilde{\alpha}}$ for compositions α.) Using Proposition 2.5, we can derive formulas expressing st_n and pa_n in the power sum basis. These are well-known; see e.g. [10, Theorem 8].

Proposition 3.2. For a positive integer n, we have

$$\mathrm{st}_n = \sum_{k=1}^{n} (-1)^{k-1} \binom{n-1}{k-1} p_{(k,1^{n-k})} \tag{3.1}$$

and

$$\mathrm{pa}_n = \sum_{\alpha \vDash n} (-1)^{n-\ell(\alpha)} p_\alpha = \sum_{\lambda \vdash n} (-1)^{n-\ell(\lambda)} \operatorname{perm}(\lambda)\, p_\lambda. \tag{3.2}$$

The star and path bases were introduced by Cho and van Willigenburg in [10], where they showed that $\{\mathrm{st}_\lambda\}$ and $\{\mathrm{pa}_\lambda\}$ are indeed (graded) bases for Sym. Equivalently, the symmetric functions $\{\mathrm{st}_n\}_{n\geq 1}$ are algebraically independent (over $\mathbb{Q}$) and thus freely generate Sym as a polynomial ring; the same is true for $\{\mathrm{pa}_n\}_{n\geq 1}$. In fact, they showed that for *any* family of connected graphs $\{G_n\}_{n\geq 1}$ where $|G_n| = n$, the chromatic symmetric functions $\{\mathbf{X}_{G_n}\}_{n\geq 1}$ have the same property [10, Theorem 5]. However, among all possible families $\{G_n\}$, the stars and paths are special because the corresponding bases exhibit a phenomenon known as *chromatic reciprocity*:

Theorem 3.3 ([2, Theorem 6.3])**.** Let $\{u_\lambda\}$ be either the star or path basis. Then the unique linear map $\varphi\colon \mathrm{Sym} \to \mathrm{Sym}$ satisfying $\varphi(p_\lambda) = u_\lambda$ for all λ is also the unique linear map satisfying $\varphi(u_\lambda) = p_\lambda$ for all λ. In particular, φ is an involution.

(As shown in [2], the key property of the families $\{\mathrm{St}_n\}$ and $\{\mathrm{Pa}_n\}$ that gives rise to chromatic reciprocity is that every connected partition of a star or path consists only of stars and paths, respectively.) By applying the involutions in Theorem 3.3 to (3.1) and (3.2), we get formulas writing the power sum basis in terms of the star and path bases, respectively:

$$p_n = \sum_{k=1}^{n} (-1)^{k-1} \binom{n-1}{k-1} \mathrm{st}_{(k,1^{n-k})} \tag{3.3}$$

and

$$p_n = \sum_{\alpha \vDash n} (-1)^{n-\ell(\alpha)} \mathrm{pa}_\alpha = \sum_{\lambda \vdash n} (-1)^{n-\ell(\lambda)} \mathrm{perm}(\lambda)\, \mathrm{pa}_\lambda\,. \tag{3.4}$$

We will now show that, in fact, there is a third involution on Sym that (nearly) swaps the star and path bases. Also, in the next section, we will give a formula expressing the star basis in terms of the path basis and vice versa (Corollary 3.14).

Our involution behaves nicely not only on stars and paths, but on a larger class of trees called caterpillars. A *caterpillar* is a tree whose internal vertices form a path. That path is called the *spine* of the caterpillar, and the leaf edges are called the *legs*. Caterpillars can be specified by giving the number of legs attached to each vertex along the spine, or alternatively, by giving the degree of each spine vertex. We record this information in a composition as follows:

Definition 3.4. Given a composition $\alpha = (\alpha_1, \ldots, \alpha_k) \neq 0$, let $C(\alpha)$ be the caterpillar whose spine vertices $v_1, \ldots, v_k$ have degrees $\deg(v_i) = \alpha_i + 1$. Equivalently, $C(\alpha)$ is the caterpillar such that v_1 has α_1 legs attached, v_k has α_k legs attached, and v_i has $\alpha_i - 1$ legs attached for $2 \leq i \leq k-1$. Also, we define $C(0)$ to be the tree with two vertices.

Figure 3.1: The caterpillars $C(12132)$, at left, and $C(13212)$, at right.

For all α, the caterpillar $C(\alpha)$ has $|\alpha|+2$ vertices.[1] In addition, the caterpillars $C(\alpha)$ and $C(\alpha^*)$ are always isomorphic; otherwise, each caterpillar (of order at least 2) corresponds to a unique composition up to reversal. Finally, note that stars and paths are both special cases of caterpillars: $\mathrm{St}_n = C(n-2)$ and $\mathrm{Pa}_n = C(1^{n-2})$.

Example 3.5. Figure 3.1 shows two caterpillars of order 11, which correspond respectively to the compositions 12132 and 13212. ◊

We can now state our main result:

Theorem 3.6. Define a $\mathbb{Q}$-algebra map $\psi\colon \mathrm{Sym} \to \mathrm{Sym}$ by $\psi(\mathrm{st}_n) = -\mathrm{pa}_n$ for all positive integers n. Then for all compositions $\alpha \neq 0$, we have

$$\psi(\mathbf{X}_{C(\alpha)}) = -\mathbf{X}_{C(\alpha^c)}.$$

In particular, taking $\alpha = 1^{n-2}$ for $n \geq 3$, we get $\psi(\mathrm{pa}_n) = -\mathrm{st}_n$, so ψ is an involution.

Proof. In this proof, we will abbreviate $\mathbf{X}_{C(\alpha)}$ to $\mathbf{X}_\alpha$ to improve readability.

We use induction on $\ell(\alpha)$. If $\ell(\alpha) = 1$, then $\mathbf{X}_\alpha = \mathrm{st}_n$ and $\mathbf{X}_{\alpha^c} = \mathrm{pa}_n$ where $n = \alpha_1 + 2$, so the statement is true by the definition of ψ.

Now, fix a composition $\alpha = (\alpha_1, \ldots, \alpha_k)$ where $k \geq 2$. We apply the deletion–near-contraction recurrence (Proposition 2.7) to $C(\alpha)$ at the edge between the first and second

[1] One way to see this is to count edges: $C(\alpha)$ has $\alpha_1 + (\alpha_2 - 1) + \cdots + (\alpha_{k-1} - 1) + \alpha_k = |\alpha| - k + 2$ legs and $k - 1$ spine edges, so $|\alpha| + 1$ edges total, which means $|\alpha| + 2$ vertices.

spine vertices. Letting $\beta = (\alpha_3, \ldots, \alpha_k)$, this yields

$$\mathbf{X}_\alpha = \mathbf{X}_{\alpha_1 - 1}\mathbf{X}_{(\alpha_2 - 1, \beta)} + \mathbf{X}_{(\alpha_1 + \alpha_2, \beta)} - p_1 \mathbf{X}_{(\alpha_1 + \alpha_2 - 1, \beta)}.$$

Applying the algebra map ψ to both sides and using the induction hypothesis, we get

$$\begin{aligned}\psi(\mathbf{X}_\alpha) &= \psi(\mathbf{X}_{\alpha_1 - 1})\psi(\mathbf{X}_{(\alpha_2 - 1, \beta)}) + \psi(\mathbf{X}_{(\alpha_1 + \alpha_2, \beta)}) - \psi(p_1)\psi(\mathbf{X}_{(\alpha_1 + \alpha_2 - 1, \beta)}) \\ &= \mathbf{X}_{(\alpha_1 - 1)^c}\mathbf{X}_{(\alpha_2 - 1, \beta)^c} - \mathbf{X}_{(\alpha_1 + \alpha_2, \beta)^c} - p_1 \mathbf{X}_{(\alpha_1 + \alpha_2 - 1, \beta)^c}.\end{aligned}$$

(Note that $p_1 = \mathrm{st}_1 = \mathrm{pa}_1$, so $\psi(p_1) = -p_1$.) Using (2.1) on the right-hand side gives

$$\psi(\mathbf{X}_\alpha) = \mathbf{X}_{1^{\alpha_1 - 1}}\mathbf{X}_{1^{\alpha_2 - 1} \odot \beta^c} - \mathbf{X}_{1^{\alpha_1 + \alpha_2} \odot \beta^c} - p_1 \mathbf{X}_{1^{\alpha_1 + \alpha_2 - 1} \odot \beta^c}. \tag{3.5}$$

We now apply the deletion–near-contraction recurrence to the caterpillar $C(1^{\alpha_1 + \alpha_2} \odot \beta^c)$ at the edge between the α_1th and $(\alpha_1 + 1)$th spine vertices. This gives

$$\mathbf{X}_{1^{\alpha_1 + \alpha_2} \odot \beta^c} = \mathbf{X}_{1^{\alpha_1 - 1}}\mathbf{X}_{1^{\alpha_2 - 1} \odot \beta^c} + \mathbf{X}_{1^{\alpha_1 - 1} 2 1^{\alpha_2 - 1} \odot \beta^c} - p_1 \mathbf{X}_{1^{\alpha_1 + \alpha_2 - 1} \odot \beta^c}. \tag{3.6}$$

Finally, we substitute (3.6) into the right-hand side of (3.5), which gives simply

$$\psi(\mathbf{X}_\alpha) = -\mathbf{X}_{1^{\alpha_1 - 1} 2 1^{\alpha_2 - 1} \odot \beta^c}.$$

We have $\alpha^c = (\alpha_1, \alpha_2)^c \odot \beta^c = 1^{\alpha_1 - 1} 2 1^{\alpha_2 - 1} \odot \beta^c$, therefore $\psi(\mathbf{X}_\alpha) = -\mathbf{X}_{\alpha^c}$, completing the induction and the proof. □

By multiplicativity, $\psi(\mathrm{st}_\lambda) = (-1)^{\ell(\lambda)} \mathrm{pa}_\lambda$ and $\psi(\mathrm{pa}_\lambda) = (-1)^{\ell(\lambda)} \mathrm{st}_\lambda$ for all partitions λ.

Remark 3.7. Recall that there is an algebra involution ι of Sym that sends $f \in \mathrm{Sym}_n$ to $(-1)^n f$, that is, ι negates odd-degree (homogeneous) elements and fixes even-degree elements. Furthermore, ι commutes with every graded endomorphism of Sym, so it follows that $\psi' = \iota \circ \psi$ is also an involution of Sym. This is an alternative version of our star-path

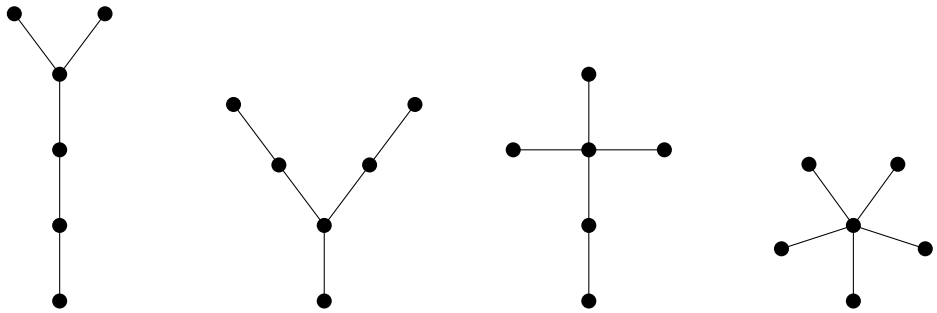

Sp(311) Sp(221) Sp(2111) Sp(11111)

Figure 3.2: The four spiders of order 6.

involution; it satisfies $\psi'(\mathrm{st}_\lambda) = (-1)^{|\lambda|-\ell(\lambda)}\,\mathrm{pa}_\lambda$ and $\psi'(\mathrm{pa}_\lambda) = (-1)^{|\lambda|-\ell(\lambda)}\,\mathrm{st}_\lambda$.

3.2 Spiders in the path basis

We now consider a different subclass of trees, namely spiders. A *spider* is a tree with exactly one vertex of degree at least 3. That vertex is called the *torso* of the spider, and the paths from the torso to the leaf vertices are called the *legs* of the spider.

Given a partition $\mu \vdash n-1$, we can define a graph $\mathrm{Sp}(\mu)$ of order n by starting with a single vertex and attaching a path of length μ_i for each i. Then the spiders of order n are exactly the graphs $\mathrm{Sp}(\mu)$ for $\ell(\mu) \geq 3$. If instead $\ell(\mu) \leq 2$, then $\mathrm{Sp}(\mu) = \mathrm{Pa}_n$, the path of order n (which can be viewed as a degenerate spider[2]). We write

$$\mathrm{Sp}_n = \{\mathrm{Sp}(\mu) : \mu \vdash n-1,\ \ell(\mu) \geq 3\}$$

for the set of all spiders of order n, and let $\mathrm{Sp}'_n = \mathrm{Sp}_n \cup \{\mathrm{Pa}_n\}$.

Example 3.8. There are four spiders with 6 vertices, corresponding to the four partitions of 5 with length at least 3. These are shown in Figure 3.2. ◊

[2]We make this distinction primarily because the term "spider" as it is typically defined (and as we have defined it here) excludes path graphs.

It was proved in [25] that spiders can be recovered from their chromatic symmetric functions, therefore settling Question 1.1 for the class of spiders. (In fact, a much stronger result is true: only the *two-part portion* of the chromatic symmetric function, which consists of all coefficients of p_λ with $\ell(\lambda) = 2$, is needed to distinguish spiders [25, Theorem 9].) Our results in the rest of this chapter will provide further insight into the CSFs of spiders.

We will need the following definitions. A *weak composition* is a finite sequence $\gamma = (\gamma_1, \ldots, \gamma_\ell)$ of *nonnegative* integers. Given weak compositions γ and δ of the same length, we write $\gamma \leq \delta$ if $\gamma_i \leq \delta_i$ for all i. We analogously define $\gamma < \delta$, $\gamma > \delta$, etc.

We begin with a preliminary formula for the chromatic symmetric function of a spider in the path basis.

Proposition 3.9. Given a partition $\mu = (\mu_1, \ldots, \mu_\ell) \vdash n - 1$, we have

$$\mathbf{X}_{\mathrm{Sp}(\mu)} = \sum_{0 \leq \gamma \leq \mu} \sum_{\alpha \vDash n - |\gamma|} (-1)^{\ell(\alpha)-1} \, \mathrm{pa}_{\alpha\gamma}, \tag{3.7}$$

where γ is a weak composition with $\gamma \leq \mu$ and α is a usual ("strong") composition.

Proof. Given a connected partition $\mathcal{P}$ of $\mathrm{Sp}(\mu)$, let A be the part containing the torso together with a (possibly empty) prefix of the vertices on each leg. Also, let γ_i be the number of vertices on the ith leg that are not in A. Then γ is a weak composition with $\gamma \leq \mu$ and A has size $n - |\gamma|$. Furthermore, $\mathcal{P}$ consists of A together with connected partitions of the paths Pa_{γ_i} for all i. Summing over all possible γ, it follows from Proposition 2.5 that

$$\mathbf{X}_{\mathrm{Sp}(\mu)} = \sum_{\mathcal{P} \vdash \mathrm{Sp}(\mu)} \prod_{B \in \mathcal{P}} (-1)^{|B|-1} p_{|B|} = \sum_{0 \leq \gamma \leq \mu} (-1)^{n-1-|\gamma|} p_{n-|\gamma|} \, \mathrm{pa}_{\gamma_1} \cdots \mathrm{pa}_{\gamma_\ell} \, .$$

Using (3.4), we can write this entirely in terms of the path basis:

$$\begin{aligned}\mathbf{X}_{\mathrm{Sp}(\mu)} &= \sum_{0\le\gamma\le\mu} (-1)^{n-1-|\gamma|}\Bigg(\sum_{\alpha\vDash n-|\gamma|} (-1)^{n-|\gamma|-\ell(\alpha)}\,\mathrm{pa}_\alpha\Bigg)\,\mathrm{pa}_\gamma \\ &= \sum_{0\le\gamma\le\mu}\sum_{\alpha\vDash n-|\gamma|} (-1)^{\ell(\alpha)-1}\,\mathrm{pa}_{\alpha\gamma}, \end{aligned} \tag{3.8}$$

as desired. □

Our goal is now to obtain a cancellation-free form of this formula. We use a generating function argument. For $m \ge 0$, define a polynomial $f_m(t)$ with coefficients in Sym by

$$f_m(t) = 1 + \mathrm{pa}_1\, t + \mathrm{pa}_2\, t^2 + \cdots + \mathrm{pa}_m\, t^m.$$

We extend this definition multiplicatively, so that $f_\gamma(t) = f_{\gamma_1\cdots\gamma_\ell}(t) = f_{\gamma_1}(t)\cdots f_{\gamma_\ell}(t)$ for a weak composition $\gamma = (\gamma_1,\ldots,\gamma_\ell)$ (and $f_0(t) = 1$). Also, define the formal power series $f(t) = \sum_{i\ge0} \mathrm{pa}_i\, t^i$.

Proposition 3.10. We have

$$\frac{f_\mu(t)}{f(t)} = \sum_{0\le\gamma\le\mu}\sum_{\alpha\in\mathrm{Comp}} (-1)^{\ell(\alpha)}\,\mathrm{pa}_{\alpha\gamma}\,t^{|\alpha|+|\gamma|}. \tag{3.9}$$

Thus, for $\mu \vdash n-1$,

$$\mathbf{X}_{\mathrm{Sp}(\mu)} = -[t^n]\frac{f_\mu(t)}{f(t)}. \tag{3.10}$$

Proof. Let $\mu = (\mu_1,\ldots,\mu_\ell)$. We first compute

$$f_\mu(t) = f_{\mu_1}(t)\cdots f_{\mu_\ell}(t) = \left(\sum_{k=0}^{\mu_1}\mathrm{pa}_k\,t^k\right)\cdots\left(\sum_{k=0}^{\mu_\ell}\mathrm{pa}_k\,t^k\right) = \sum_{0\le\gamma\le\mu}\mathrm{pa}_\gamma\,t^{|\gamma|}.$$

Next, we write

$$\frac{1}{f(t)} = \frac{1}{1-(-\sum_{k\geq 1} \mathrm{pa}_k\, t^k)} = \sum_{\ell \geq 0} \Big(-\sum_{k\geq 1} \mathrm{pa}_k\, t^k \Big)^{\ell} = \sum_{\alpha \in \mathrm{Comp}} (-1)^{\ell(\alpha)}\, \mathrm{pa}_\alpha\, t^{|\alpha|}.$$

Therefore,

$$\begin{aligned} \frac{f_\mu(t)}{f(t)} &= \Big(\sum_{0 \leq \gamma \leq \mu} \mathrm{pa}_\gamma\, t^{|\gamma|} \Big) \Big(\sum_{\alpha \in \mathrm{Comp}} (-1)^{\ell(\alpha)}\, \mathrm{pa}_\alpha\, t^{|\alpha|} \Big) \\ &= \sum_{0 \leq \gamma \leq \mu} \sum_{\alpha \in \mathrm{Comp}} (-1)^{\ell(\alpha)}\, \mathrm{pa}_{\alpha\gamma}\, t^{|\alpha|+|\gamma|}, \end{aligned}$$

as claimed. Comparing this with (3.7) yields the desired formula for $\mathbf{X}_{\mathrm{Sp}(\mu)}$. □

It will be helpful to also consider other coefficients of the quotient $f_\mu(t)/f(t)$. Given $\mu \vdash n-1$ and an integer d, we define a symmetric function $\mathbf{X}_{\mu,d}$ by

$$\mathbf{X}_{\mu,d} = -[t^{n+d}] \frac{f_\mu(t)}{f(t)} = \sum_{0 \leq \gamma \leq \mu} \sum_{\alpha \vDash n+d-|\gamma|} (-1)^{\ell(\alpha)-1}\, \mathrm{pa}_{\alpha\gamma}\,. \tag{3.11}$$

Thus, $\mathbf{X}_{\mu,0} = \mathbf{X}_{\mathrm{Sp}(\mu)}$.[3] In this section, we will find formulas for the path expansions of all $\mathbf{X}_{\mu,d}$, as well as the dimension of the span of $\mathbf{X}_{\mu,d}$ for $\mu \vdash n-1$ (given any n and d).

We now rewrite the quotient $f_\mu(t)/f(t)$ as follows. (The next result is a general fact about generating functions and does not actually involve the CSF.)

Lemma 3.11. For $m \geq 0$, define $\tilde{f}_m(t) = f(t) - f_m(t) = \sum_{i=m+1}^{\infty} \mathrm{pa}_i\, t^i$, and extend multiplicatively to $\tilde{f}_\gamma(t)$ for weak compositions γ. Then for any $\mu = (\mu_1, \ldots, \mu_\ell)$, we have

$$-\frac{f_\mu(t)}{f(t)} = \sum_{j=1}^{\ell} (-1)^j \tilde{f}_{\mu_1 \ldots \mu_{j-1}}(t) f_{\mu_{j+1} \ldots \mu_\ell}(t) + (-1)^{\ell-1} \frac{\tilde{f}_\mu(t)}{f(t)}. \tag{3.12}$$

[3] For $d \geq 0$, it can be shown that $\mathbf{X}_{\mu,d}$ is (up to sign) an example of a *weighted* chromatic symmetric function: specifically, it is the CSF of the spider $\mathrm{Sp}(\mu)$ when its torso is given weight $1+d$ and all other vertices are given weight 1. More information about weighted CSFs can be found in [12].

Proof. We have

$$f_\mu(t) = f_{\mu_1}(t)\cdots f_{\mu_\ell}(t) = \prod_{i=1}^{\ell}\left(f(t) - \tilde{f}_{\mu_i}(t)\right).$$

We can index the terms in this product by subsets $I \subseteq [\ell]$ consisting of the indices i for which we choose $f(t)$ instead of $-\tilde{f}_{\mu_i}(t)$. This gives

$$f_\mu(t) = \sum_{I\subseteq[\ell]} f(t)^{|I|} \prod_{i\in[\ell]\setminus I}\left(-\tilde{f}_{\mu_i}(t)\right).$$

When $I = \varnothing$, the corresponding term of this sum is $(-1)^\ell \tilde{f}_\mu(t)$. When instead I is nonempty, suppose we fix the value of $j = \min(I)$. Then j is always an element of I, while all of $[j-1]$ lies outside I. Let $I = \{j\}\cup J$ for $J \subseteq [j+1,\ell]$. Then, the sum of the corresponding terms for this value of j is

$$\begin{aligned}
&f(t)\cdot(-1)^{j-1}\tilde{f}_{\mu_1\cdots\mu_{j-1}}(t) \sum_{J\subseteq[j+1,\ell]} f(t)^{|J|} \prod_{i\in[j+1,\ell]\setminus J}\left(-\tilde{f}_{\mu_i}(t)\right)\\
&= f(t)\cdot(-1)^{j-1}\tilde{f}_{\mu_1\cdots\mu_{j-1}}(t) \prod_{i=j+1}^{\ell}\left(f(t) - \tilde{f}_{\mu_i}(t)\right)\\
&= f(t)\cdot(-1)^{j-1}\tilde{f}_{\mu_1\cdots\mu_{j-1}}(t) f_{\mu_{j+1}\cdots\mu_\ell}(t).
\end{aligned}$$

To finish, we sum over all j (and remember to include the term $I = \varnothing$) to obtain

$$f_\mu(t) = f(t)\sum_{j=1}^{\ell}(-1)^{j-1}\tilde{f}_{\mu_1\cdots\mu_{j-1}}(t) f_{\mu_{j+1}\cdots\mu_\ell}(t) + (-1)^\ell \tilde{f}_\mu(t).$$

Dividing by $-f(t)$ yields the desired equation (3.12). □

By combining (3.10) and (3.12), we obtain for $\mu \vdash n-1$ that

$$\mathbf{X}_{\mu,d} = \sum_{j=1}^{\ell}(-1)^j\cdot[t^{n+d}]\tilde{f}_{\mu_1\ldots\mu_{j-1}}(t) f_{\mu_{j+1}\ldots\mu_\ell}(t) + (-1)^{\ell-1}\cdot[t^{n+d}]\frac{\tilde{f}_\mu(t)}{f(t)}, \tag{3.13}$$

which we now use to give a cancellation-free formula for the path expansion of $\mathbf{X}_{\mu,d}$. The

key fact is that for any $\lambda \vdash n+d$, at most one of the $\ell+1$ terms on the right-hand side of (3.13) can contain pa_λ. (Also, there is no cancellation within any individual term because the fs and $\tilde{f}$s have all nonnegative coefficients.) Let us see why.

In the summation of (3.13), when each product $\tilde{f}_{\mu_1 \cdots \mu_{j-1}}(t) f_{\mu_{j+1} \cdots \mu_\ell}(t)$ is expanded, every pa_λ that appears (in any coefficient) has $\ell(\lambda) \leq \ell - 1$. Meanwhile, in the last term, every pa_λ that appears in a coefficient of $\tilde{f}_\mu(t)/f(t)$ has $\ell(\lambda) \geq \ell$. Thus we only need to consider the terms of the summation. Suppose that $\lambda = (\lambda_1, \ldots, \lambda_k) \vdash n+d$ with $k \leq \ell - 1$. If pa_λ appears in $[t^{n+d}]\tilde{f}_{\mu_1 \cdots \mu_{j-1}}(t) f_{\mu_{j+1} \cdots \mu_\ell}(t)$, then λ must consist of $j-1$ parts that are respectively greater than $\mu_1, \ldots, \mu_{j-1}$, and at most $\ell - j$ parts that are respectively less than or equal to (some of) $\mu_{j+1}, \ldots, \mu_\ell$. Since λ and μ are partitions, it must be the case that

$$\lambda_i > \mu_i \qquad \text{for} \qquad i = 1, \ldots, j-1$$

and

$$\lambda_i \leq \mu_{i+1} \qquad \text{for} \qquad i = j, \ldots, k.$$

This can only happen for at most one value of j: specifically, j must be the least index with[4] $\lambda_j \leq \mu_j$ (although the latter condition still might not hold, in which case $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} = 0$).

This discussion, after some careful analysis, leads to a formula for the coefficient of any pa_λ in the path expansion of $\mathbf{X}_{\mu,d}$. For a composition α, recall that $\mathrm{perm}(\alpha)$ is the number of distinct permutations of (the parts of) α. Let $\mathrm{perm}_>(\alpha;\beta)$ be the number of distinct permutations of α whose parts are respectively greater than those of β; define $\mathrm{perm}_\leq(\alpha;\beta)$ analogously.

Theorem 3.12. Let $\lambda = (\lambda_1, \ldots, \lambda_k) \vdash n+d$ and $\mu = (\mu_1, \ldots, \mu_\ell) \vdash n-1$ be partitions.

(i) If $k \geq \ell$, then

$$[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} = (-1)^{k-1} \mathrm{perm}_>(\lambda; \mu 0^{k-\ell}).$$

[4] If $\lambda_i > \mu_i$ for all $i = 1, \ldots, k$, then we take $j = k+1$.

(ii) If $k < \ell$, let j be the least index such that $\lambda_j \leq \mu_j$, or $j = k+1$ if none exists. Then

$$[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} = (-1)^j \operatorname{perm}_>(\lambda_1 \dots \lambda_{j-1}; \mu_1 \dots \mu_{j-1}) \operatorname{perm}_\leq(\lambda_j \dots \lambda_k 0^{\ell-k-1}; \mu_{j+1} \dots \mu_\ell).$$

Proof. If $k \geq \ell$, then pa_λ can only appear in the term $(-1)^{\ell-1}[t^{n+d}]\tilde{f}_\mu(t)/f(t)$ of (3.13). We have

$$\frac{\tilde{f}_\mu(t)}{f(t)} = \left(\sum_{i>\mu_1} \mathrm{pa}_i\, t^i\right) \cdots \left(\sum_{i>\mu_\ell} \mathrm{pa}_i\, t^i\right) \sum_{\alpha \in \mathrm{Comp}} (-1)^{\ell(\alpha)}\, \mathrm{pa}_\alpha\, t^{|\alpha|}.$$

To obtain $\mathrm{pa}_\lambda\, t^{n+d}$ from this product, we must choose a permutation of λ so that the first ℓ parts are respectively greater than $\mu_1, \dots, \mu_\ell$, with no restrictions on the last $k-\ell$ parts. This is exactly the same as a permutation of λ whose parts are respectively greater than those of $\mu 0^{k-\ell}$. Each permutation contributes $(-1)^{\ell-1} \cdot (-1)^{k-\ell} = (-1)^{k-1}$ to the coefficient of pa_λ, so (i) is proved.

If $k < \ell$, then define j as in the theorem statement. As stated in our previous discussion, pa_λ can only appear in the term $(-1)^j[t^{n+d}]\tilde{f}_{\mu_1 \cdots \mu_{j-1}}(t) f_{\mu_{j+1} \cdots \mu_\ell}(t)$ of (3.13). We have

$$\begin{aligned}&\tilde{f}_{\mu_1 \cdots \mu_{j-1}}(t) f_{\mu_{j+1} \cdots \mu_\ell}(t) \\ &= \left(\sum_{i>\mu_1} \mathrm{pa}_i\, t^i\right) \cdots \left(\sum_{i>\mu_{j-1}} \mathrm{pa}_i\, t^i\right) \left(\sum_{i \leq \mu_{j+1}} \mathrm{pa}_i\, t^i\right) \cdots \left(\sum_{i \leq \mu_\ell} \mathrm{pa}_i\, t^i\right).\end{aligned}$$

To obtain $\mathrm{pa}_\lambda\, t^{n+d}$ from this product, we must choose one pa_i coefficient from each term, comprising (in some order) the coefficients $\mathrm{pa}_{\lambda_1}, \dots, \mathrm{pa}_{\lambda_k}$ as well as $\ell - k - 1$ occurrences of pa_0. The coefficients chosen from the first $j-1$ terms must be $\mathrm{pa}_{\lambda_1}, \dots, \mathrm{pa}_{\lambda_{j-1}}$ in some order, since $\lambda_j \leq \mu_j \leq \mu_{j-1}$. Thus, the coefficients chosen from the last terms must be $\mathrm{pa}_{\lambda_j}, \dots, \mathrm{pa}_{\lambda_k}$ as well as $\ell - k - 1$ occurrences of pa_0. This proves (ii). □

As a corollary, we can combine the cases of Theorem 3.12 to give a concise criterion for whether pa_λ appears in the path expansion of $\mathbf{X}_{\mu,d}$.

Corollary 3.13. Let $\lambda \vdash n+d$ and $\mu \vdash n-1$. Considering λ and μ to be padded

with infinitely many zeros, let j be the least positive integer such that $\lambda_j \leq \mu_j$. Then $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} \neq 0$ if and only if $\lambda_i \leq \mu_{i+1}$ for all $i \geq j$.

Proof. Let $k = \ell(\lambda)$ and $\ell = \ell(\mu)$. If $k < \ell$, then Theorem 3.12(ii) immediately proves the desired statement. So suppose that $k \geq \ell$. If $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} \neq 0$, then Theorem 3.12(i) shows that $\lambda_i > \mu_i$ for $i = 1, \ldots, \ell$. Then $j = k + 1$, so $\lambda_i = \mu_{i+1} = 0$ for all $i \geq j$, and the condition is satisfied. Conversely, suppose that $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} = 0$, so that $\lambda_i < \mu_i$ for some $i \in [\ell]$. Then we must have $j \leq i$ by minimality, so $j \leq \ell$. But $\lambda_\ell > 0$ while $\mu_{\ell+1} = 0$, so the condition is not satisfied. □

In addition, since stars are examples of spiders, we can use Theorem 3.12 to write the star basis in terms of the path basis and vice versa.

Corollary 3.14. For all integers $n \geq 3$, we have

$$\mathrm{st}_n = \sum_{\substack{\lambda \vdash n \\ \ell(\lambda) \leq n-2}} (-1)^{m_{>1}(\lambda)+1} \frac{m_{>1}(\lambda)!(n - m_{>1}(\lambda) - 2)!}{\mathrm{aut}(\lambda)(n - \ell(\lambda) - 2)!} \mathrm{pa}_\lambda \tag{3.14}$$

and

$$\mathrm{pa}_n = \sum_{\substack{\lambda \vdash n \\ \ell(\lambda) \leq n-2}} (-1)^{m_1(\lambda)} \frac{m_{>1}(\lambda)!(n - m_{>1}(\lambda) - 2)!}{\mathrm{aut}(\lambda)(n - \ell(\lambda) - 2)!} \mathrm{st}_\lambda \tag{3.15}$$

where $m_{>1}(\lambda) = \#\{i : \lambda_i > 1\} = \ell(\lambda) - m_1(\lambda)$. (Also, by definition, $\mathrm{st}_n = \mathrm{pa}_n$ for $n \leq 3$.)

Proof. We apply Theorem 3.12 with $d = 0$ and $\mu = 1^{n-1}$, so that $\mathbf{X}_{\mu,d} = \mathrm{St}_n$. Fix $\lambda = (\lambda_1, \ldots, \lambda_k) \vdash n$. If $k \geq n - 1$, then all parts of λ after the first are 1, so one can check that Theorem 3.12(i) gives $[\mathrm{pa}_\lambda] \mathrm{st}_n = 0$.

Suppose $k < n - 1$. Then $j = m_{>1}(\lambda) + 1$, so that $\lambda_1 \geq \cdots \geq \lambda_{j-1} > 1 = \lambda_j = \ldots = \lambda_k$.

Therefore, Theorem 3.12(ii) gives

$$\begin{aligned}[\mathrm{pa}_\lambda]\,\mathrm{st}_n &= (-1)^j \operatorname{perm}_>(\lambda_1 \dots \lambda_{j-1}; 1^{j-1}) \operatorname{perm}_\leq(\lambda_j \cdots \lambda_k 0^{n-k-2}; 1^{n-j-1}) \\ &= (-1)^j \operatorname{perm}(\lambda_1 \dots \lambda_{j-1}) \operatorname{perm}(1^{k-j+1} 0^{n-k-2}) \\ &= (-1)^j \frac{(j-1)!}{\operatorname{aut}(\lambda_1 \dots \lambda_{j-1})} \binom{n-j-1}{k-j+1}.\end{aligned}$$

Substituting $j = m_{>1}(\lambda)+1$ and noting that $\operatorname{aut}(\lambda_1 \dots \lambda_{j-1}) = \operatorname{aut}(\lambda)/m_1(\lambda)!$ yields (3.14) after some effort. To obtain (3.15), simply apply the involution ψ of Theorem 3.6. □

We note that one may obtain (3.15) directly by applying the deletion–near-contraction recurrence (Proposition 2.7) to a path graph, although the relevant computations are likely to be rather involved; also see [5, Theorem 3.2].

3.3 Spider space

Next, we study the subspace of Sym spanned by the symmetric functions $\mathbf{X}_{\mu,d}$. Let

$$\mathcal{S}_{n,d} = \operatorname{span}\{\mathbf{X}_{\mu,d} : \mu \vdash n-1\} \subseteq \mathrm{Sym}_{n+d}.$$

When $d = 0$, we get $\mathcal{S}_{n,0} = \operatorname{span}\{\mathbf{X}_S : S \in \mathrm{Sp}'_n\}$, the span of the CSFs of the spiders and path of order n. (Later on, we will be able to modify our results to exclude the path.) Our goal in this section is to give a basis for and compute the dimension of $\mathcal{S}_{n,d}$, as well as characterize the linear relations between the $\mathbf{X}_{\mu,d}$.

3.3.1 Linearly independent spiders

We begin by giving a linearly independent set of functions $\mathbf{X}_{\mu,d}$, which will turn out to be the desired basis for $\mathcal{S}_{n,d}$. To describe this set, we need the following definition:

Definition 3.15. Let μ be a partition and k be an integer. A *k-fixed point* of μ is an index i such that $\mu_i = i+k$. We write $\mathrm{Par}_m^{(k)}$ for the set of all partitions of m with a k-fixed point.

In this definition, we consider μ to be padded with infinitely many zeros: thus, if $\ell(\mu) < r$ for some r, then r is a $(-r)$-fixed point of μ. Observe that a k-fixed point of μ is unique if it exists.

Example 3.16. We have $\mathrm{Par}_7^{(-3)} = \{7, 61, 52, 43, 4111, 3211, 2221\}$. Note that the first four partitions μ in this set have $\mu_3 = 0$, while the last three have $\mu_4 = 1$. $\Diamond$

As another example of working with fixed points of partitions, we record the following lemma for future use.

Lemma 3.17. Let a and b be integers with $a + b = -1$. Given any partition μ, *exactly one* of the following is true: (i) μ has a a-fixed point; (ii) μ' has a b-fixed point.

Proof. In this proof, we pad all partitions with infinitely many zeroes; we also use the convention $\lambda_0 = +\infty$ for partitions λ (and do not consider this to be a part of λ).

Suppose that μ' has a b-fixed point, say $\mu'_j = j + b$. Then $\mu_{j+b} \geq j > \mu_{j+b+1}$, so $\mu_{j+b} - (j+b) \geq -b$ while $\mu_{j+b+1} - (j+b+1) < -b-1$. The sequence $(\mu_k - k)_{k \geq 0}$ is strictly decreasing, so it never takes the value $-b - 1 = a$. Thus μ does not have an a-fixed point.

Conversely, suppose that μ' does not have a b-fixed point. The sequence $(\mu'_k - k)_{k \geq 0}$ is strictly decreasing and does not take the value b, so there exists a unique k such that $\mu'_k - k > b > \mu'_{k+1} - (k+1)$. Then $\mu'_k > b+k$, which means $\mu_{b+k+1} \geq k$, while $\mu'_{k+1} < b+k+1$, which means $\mu_{b+k+1} < k+1$. Thus $\mu_{b+k+1} = k$, or $\mu_j = j + a$ where $j = b + k + 1$, so μ has an a-fixed point. $\square$

We now seek to prove the following result.

Theorem 3.18. The functions $\mathbf{X}_{\mu,d}$ for $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$ are linearly independent. Thus,

$$\dim \mathcal{S}_{n,d} \geq \left|\mathrm{Par}_{n-1}^{(-d-2)}\right|.$$

To prove this, we study the expansions of these symmetric functions in the path basis. Suppose that $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$, and let i be the unique $(-d-2)$-fixed point of μ. We define a

corresponding partition $\hat{\mu}$ by

$$\hat{\mu} = (\mu_1 + 1, \dots, \mu_{i-1} + 1, \mu_{i+1}, \mu_{i+2}, \dots).$$

Then the size of $\hat{\mu}$ is $|\hat{\mu}| = |\mu| + (i-1) - \mu_i = (n-1) + (i-1) - (i-d-2) = n + d$.

Lemma 3.19. The map $\mu \mapsto \hat{\mu}$ is a bijection from $\mathrm{Par}_{n-1}^{(-d-2)}$ to $\{\lambda : \lambda' \in \mathrm{Par}_{n+d}^{(d)}\}$.

Proof. Let $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$ be arbitrary with $\mu_i = i - d - 2$. Then $\hat{\mu}$ satisfies

$$\hat{\mu}_{i-1} = \mu_{i-1} + 1 \geq \mu_i + 1 = i - d - 1 \qquad \text{and} \qquad \hat{\mu}_i = \mu_{i+1} \leq \mu_i = i - d - 2.$$

It follows that $\hat{\mu}'_{i-d-1} = i - 1$, so $\hat{\mu}'$ has a d-fixed point as required.

We now construct the inverse map. Let $\lambda \vdash n + d$ be a partition such that λ' has a d-fixed point, say $\lambda'_j = j + d$ for some j. In other words, we have $\lambda_{j+d} \geq j > \lambda_{j+d+1}$. The fixed point j is unique, so we can define a partition $\hat{\lambda}$ by

$$\hat{\lambda} = (\lambda_1 - 1, \dots, \lambda_{j+d} - 1, j - 1, \lambda_{j+d+1}, \lambda_{j+d+2}, \dots).$$

Then $\hat{\lambda}_{j+d+1} = j - 1$, so $\hat{\lambda}$ has a $(-d-2)$-fixed point at $j + d + 1$. It is easy to check that the map $\lambda \mapsto \hat{\lambda}$ is the desired inverse. □

Lemma 3.20. If $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$, then $\hat{\mu}$ is the lexicographically least partition $\lambda \vdash n + d$ such that $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} \neq 0$.

Proof. Let $\mu_i = i - d - 2$, so that $\hat{\mu} = (\mu_1 + 1, \dots, \mu_{i-1} + 1, \mu_{i+1}, \mu_{i+2}, \dots)$. Taking $\lambda = \hat{\mu}$ in Corollary 3.13, we get $[\mathrm{pa}_{\hat{\mu}}]\mathbf{X}_{\mu,d} \neq 0$.

Now, choose an arbitrary $\lambda \vdash n + d$ that satisfies $[\mathrm{pa}_\lambda]\mathbf{X}_{\mu,d} \neq 0$. We show that $\hat{\mu}$ weakly precedes λ in lexicographic order. By Corollary 3.13, there exists j such that $\lambda_k > \mu_k$ for $1 \leq k \leq j - 1$ and $\lambda_k \leq \mu_{k+1}$ for $k \geq j$. If $j > i$, then for all $1 \leq k \leq i - 1$ we have $\lambda_k \geq \mu_k + 1 = \hat{\mu}_k$, while $\lambda_i > \mu_i \geq \mu_{i+1} = \hat{\mu}_i$; thus, $\hat{\mu}$ precedes λ in lexicographic order.

Suppose instead that $j \leq i$. Then $\lambda_k \geq \mu_k + 1 = \hat{\mu}_k$ for all $1 \leq k \leq j-1$. If this inequality is strict for any k, then $\hat{\mu}$ again precedes λ in lexicographic order. Otherwise, we have $(\lambda_1, \ldots, \lambda_{j-1}) = (\mu_1 + 1, \ldots, \mu_{j-1} + 1)$, so

$$\begin{aligned}
|\lambda| &= (\mu_1 + 1) + \cdots + (\mu_{j-1} + 1) + \lambda_j + \lambda_{j+1} + \cdots \\
&\leq (\mu_1 + 1) + \cdots + (\mu_{j-1} + 1) + \mu_{j+1} + \mu_{j+2} + \cdots \\
&= |\mu| + (j-1) - \mu_j \\
&\leq (n-1) + (i-1) - \mu_i \\
&= n + d.
\end{aligned}$$

But $|\lambda| = n + d$, so equality must hold throughout. Thus, $j = i$ and $\lambda_k = \mu_{k+1}$ for all $k \geq j$, which means $\lambda = \hat{\mu}$. □

By Lemma 3.19, the partitions $\hat{\mu}$ are all distinct for $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$. Hence, Lemma 3.20 implies Theorem 3.18, as desired.

3.3.2 The Bernstein creation operators

We now seek to prove the corresponding upper bound on $\dim \mathcal{S}_{n,d}$. To do this, we relate the symmetric functions $\mathbf{X}_{\mu,d}$ with a family of linear endomorphisms of Sym called the *Bernstein creation operators*. First, we give some background on these operators, following [7, Section 2.2] or [18, Exercise 2.9.1].

Let g be a symmetric function. The *skewing* operator associated to g is the map $g^{\perp} \colon \mathrm{Sym} \to \mathrm{Sym}$ defined by $g^{\perp}(f) = \sum \langle g, f_{(1)} \rangle f_{(2)}$. (Here we are using Sweedler notation (2.6); recall that $\langle \cdot, \cdot \rangle$ is the Hall inner product on Sym.)

Definition 3.21. For an integer m, the *Bernstein creation operator* of order m is given by

$$B_m(g) = \sum_{r \geq 0} (-1)^r h_{m+r} e_r^{\perp}(g).$$

Observe that B_m raises degree by m: if $g \in \mathrm{Sym}_k$, then $B_m(g) \in \mathrm{Sym}_{m+k}$.

The importance of the B_m, and the reason that they are called "creation" operators, is that they can be used to generate the Schur basis $\{s_\lambda\}$ of Sym. Given any tuple of (not necessarily positive) integers $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_\ell)$, we define a symmetric function s_α by

$$s_\alpha = \det\left[h_{\alpha_i - i + j}\right]_{1 \le i,j \le \ell} = \det \begin{bmatrix} h_{\alpha_1} & h_{\alpha_1+1} & \cdots & h_{\alpha_1+\ell-1} \\ h_{\alpha_2-1} & h_{\alpha_2} & \cdots & h_{\alpha_2+\ell-2} \\ \vdots & \vdots & \ddots & \vdots \\ h_{\alpha_\ell-\ell+1} & h_{\alpha_\ell-\ell+2} & \cdots & h_{\alpha_\ell} \end{bmatrix}. \tag{3.16}$$

When α is a partition, s_α is the usual Schur function by the Jacobi-Trudi formula (2.2). In general, by appropriately permuting the rows of the determinant in (3.16), one can show that s_α is either 0 or $\pm s_\lambda$ for some partition λ: for instance,

$$s_{136} = \det \begin{bmatrix} h_1 & h_2 & h_3 \\ h_2 & h_3 & h_4 \\ h_4 & h_5 & h_6 \end{bmatrix} = -\det \begin{bmatrix} h_4 & h_5 & h_6 \\ h_2 & h_3 & h_4 \\ h_1 & h_2 & h_3 \end{bmatrix} = -s_{433}.$$

Bernstein observed [24, pp. 96] that for all $m \in \mathbb{Z}$ and $\alpha = (\alpha_1, \ldots, \alpha_\ell) \in \mathbb{Z}^\ell$, we have

$$B_m(s_\alpha) = s_{(m,\alpha_1,\ldots,\alpha_\ell)}. \tag{3.17}$$

It follows that for any α, we have $s_\alpha = (B_{\alpha_1} \circ \cdots \circ B_{\alpha_\ell})(1)$; thus the Schur basis can be "created" starting with 1 by iteratively applying the Bernstein operators.

The key property of the B_m that we will need is the relation

$$B_m \circ B_{m+1} = 0 \tag{3.18}$$

for all m. To see why this is true, simply evaluate $(B_m \circ B_{m+1})(s_\alpha) = B_{(m,m+1,\alpha_1,\ldots,\alpha_\ell)}$ and note that the corresponding determinant in (3.16) has its first two rows equal.

We now present two more facts, which we then use to state and prove our theorem relating spiders and the Bernstein operators.

Proposition 3.22 ([2, Remark 6.5])**.** There is an algebra map $U\colon \mathrm{Sym} \to \mathrm{Sym}$ that satisfies $U(h_\mu) = p_\mu$ and $U(e_\mu) = \mathrm{pa}_\mu$ for all μ.

Lemma 3.23. For all μ, we have $e_r^\perp(e_\mu) = \sum_{\substack{0 \le \gamma \le \mu \\ |\gamma| = |\mu| - r}} e_\gamma$. (Here γ is a weak composition.)

Proof. Recall that $\Delta(e_n) = \sum_{i+j=n} e_i \otimes e_j$ for all n. Because Δ is multiplicative, it follows that $\Delta(e_\mu) = \sum_{\beta+\gamma=\mu} e_\beta \otimes e_\gamma$, where β and γ range over all pairs of weak compositions such that $\beta_i + \gamma_i = \mu_i$ for all i. Hence, $e_r^\perp(e_\mu) = \sum_{\beta+\gamma=\mu} \langle e_r, e_\beta \rangle\, e_\gamma$.

To compute $\langle e_r, e_\beta \rangle$, we apply the involution ω, giving $\langle e_r, e_\beta \rangle = \langle h_r, h_\beta \rangle$. We have $h_r = \sum_{\lambda \vdash r} m_\lambda$, so it follows from the definition of the Hall inner product that $\langle h_r, h_\beta \rangle$ equals 1 if $|\beta| = r$ and 0 otherwise. The conclusion follows, since $\beta + \gamma = \mu$ and $|\beta| = r$ is equivalent to $0 \le \gamma \le \mu$ and $|\gamma| = |\mu| - r$. □

Theorem 3.24. For an integer d, define a linear map $F_d\colon \mathrm{Sym} \to \mathrm{Sym}$ by $F_d(\mathrm{pa}_\mu) = \mathbf{X}_{\mu,d}$ for all partitions μ. Then

$$(-1)^d F_d = U \circ B_{d+1} \circ U^{-1}.$$

Proof. It suffices to check this equality on the path basis. We fix $\mu \vdash n-1$ and compute

$$\begin{aligned}
(B_{d+1} \circ U^{-1})(\mathrm{pa}_\mu) = B_{d+1}(e_\mu) &= \sum_{r \ge 0} (-1)^r h_{d+1+r}\, e_r^\perp(e_\mu) \\
&= \sum_{r \ge 0} (-1)^r h_{d+1+r} \sum_{\substack{0 \le \gamma \le \mu \\ |\gamma| = n-1-r}} e_\gamma \\
&= \sum_{0 \le \gamma \le \mu} (-1)^{n-1-|\gamma|} h_{n+d-|\gamma|} e_\gamma,
\end{aligned}$$

where in the last step we swapped the order of summation and substituted $r = n - 1 - |\gamma|$.

We now apply U to obtain

$$\begin{aligned}(U \circ B_{d+1} \circ U^{-1})(\mathrm{pa}_\mu) &= \sum_{0 \le \gamma \le \mu} (-1)^{n-1-|\gamma|} p_{n+d-|\gamma|}\, \mathrm{pa}_\gamma \\ &= \sum_{0 \le \gamma \le \mu} (-1)^{n-1-|\gamma|} \left(\sum_{\alpha \vDash n+d-|\gamma|} (-1)^{n+d-|\gamma|-\ell(\alpha)}\, \mathrm{pa}_\alpha \right) \mathrm{pa}_\gamma \\ &= \sum_{0 \le \gamma \le \mu} \sum_{\alpha \vDash n+d-|\gamma|} (-1)^{d-\ell(\alpha)-1}\, \mathrm{pa}_{\alpha\gamma} \\ &= (-1)^d \mathbf{X}_{\mu,d},\end{aligned}$$

where we used (3.4) and (3.11). This is exactly $(-1)^d F_d(\mathrm{pa}_\mu)$, so the theorem is proved. □

It follows from the above theorem and (3.18) that $F_{d-1} \circ F_d = 0$ for all d. This is the fact we need to prove our main result.

Theorem 3.25. For all integers $n \ge 1$ and d, the subspace $\mathcal{S}_{n,d} = \mathrm{span}\{\mathbf{X}_{\mu,d} : \mu \vdash n-1\}$ has basis given by$\mathbf{X}_{\mu,d}$ for $\mu \in \mathrm{Par}_{n-1}^{(-d-2)}$, and so its dimension is $\left|\mathrm{Par}_{n-1}^{(-d-2)}\right|$.

Proof. We have that the composition of maps

$$\mathrm{Sym}_{n-1} \xrightarrow{F_d} \mathrm{Sym}_{n+d} \xrightarrow{F_{d-1}} \mathrm{Sym}_{n+2d}$$

is zero. Thus, $\mathcal{S}_{n,d} = \mathrm{im}\, F_d \subseteq \ker F_{d-1}$, so by the rank-nullity theorem,

$$\mathrm{rank}\, F_d + \mathrm{rank}\, F_{d-1} \le \dim \mathrm{Sym}_{n+d} = |\mathrm{Par}_{n+d}|$$

(where we restrict the domains of F_d and F_{d-1} to Sym_{n-1} and Sym_{n+d}, respectively). On the other hand, by the definition of F_d, we have

$$\mathrm{rank}\, F_d = \dim \mathcal{S}_{n,d} \ge \left|\mathrm{Par}_{n-1}^{(-d-2)}\right|$$

by Theorem 3.18. Similarly,

$$\operatorname{rank} F_{d-1} = \dim \mathcal{S}_{n+d,d-1} \geq \left|\mathrm{Par}_{n+d}^{(-d-1)}\right|.$$

Thus

$$\begin{aligned}\operatorname{rank} F_d + \operatorname{rank} F_{d-1} &\geq \left|\mathrm{Par}_{n-1}^{(-d-2)}\right| + \left|\mathrm{Par}_{n+d}^{(-d-1)}\right| \\ &= \left|\mathrm{Par}_{n+d}^{(d)}\right| + \left|\mathrm{Par}_{n+d}^{(-d-1)}\right| \\ &= |\mathrm{Par}_{n+d}|,\end{aligned}$$

where we used Lemma 3.19 and then Lemma 3.17. Hence equality must hold everywhere: in particular, $\dim \mathcal{S}_{n,d} = \left|\mathrm{Par}_{n-1}^{(-d-2)}\right|$, as desired. □

The dimension-counting in the previous proof implies the following fact:

Corollary 3.26. The sequence $\mathrm{Sym}_{n-1} \xrightarrow{F_d} \mathrm{Sym}_{n+d} \xrightarrow{F_{d-1}} \mathrm{Sym}_{n+2d}$ is exact.

To arrive at these results, one could instead show that the corresponding sequence of Bernstein operators is exact, or directly compute the rank of the Bernstein operators, and then apply Theorem 3.24. These methods do not require first constructing an explicit basis for $\mathcal{S}_{n,d}$. (This can all be done by analyzing the effect of the Bernstein operators on the Schur basis, as is done in [18, Exercise 2.9.1].)

We can now take $d = 0$ to recover the span of the CSFs of the order-n spiders and path. It is known that $\mathbf{X}_{\mathrm{Pa}_n} = \mathrm{pa}_n$ does not lie in the span of $\mathbf{X}_T$ for non-path trees T [27, Corollary 4.7], so excluding the path simply decreases the dimension of the span by 1. The following corollary is therefore immediate from Theorem 3.25:

Corollary 3.27. Let n be a positive integer.

(i) The span of $\{\mathbf{X}_S : S \in \mathrm{Sp}'_n\}$ has basis $\{\mathbf{X}_{\mathrm{Sp}(\mu)} : \mu \in \mathrm{Par}_{n-1}^{(-2)}\}$ and dimension $\left|\mathrm{Par}_{n-1}^{(-2)}\right|$.

(ii) The span of $\{\mathbf{X}_S : S \in \mathrm{Sp}_n\}$ has basis $\{\mathbf{X}_{\mathrm{Sp}(\mu)} : \mu \in \mathrm{Par}_{n-1}^{(-2)}, \ell(\mu) \geq 3\}$ and dimension $\left|\mathrm{Par}_{n-1}^{(-2)}\right| - 1$.

3.3.3 Linear relations between spiders

Finally, we use our results to study the linear relations between the CSFs of spiders. To our knowledge, this is the first nontrivial family of trees for which all linear relations between their CSFs have been classified.

The key fact is the exactness described in Corollary 3.26. By the definition of F_d, the linear relations between the symmetric functions $\mathbf{X}_{\mu,d}$ correspond to the elements of $\ker F_d$ (when written in the path basis), which are exactly the elements of $\operatorname{im} F_{d+1}$. But $\operatorname{im} F_{d+1}$ is spanned by the symmetric functions $\mathbf{X}_{\mu,d+1}$, which we know how to write in the path basis. Thus we have the following result:

Theorem 3.28. Fix integers $n \geq 1$ and d. Choose any partition $\nu \vdash n-d-3$, and expand $\mathbf{X}_{\nu,d+1}$ in the path basis as $\mathbf{X}_{\nu,d+1} = \sum_{\mu \vdash n-1} a_\mu \operatorname{pa}_\mu$ for some coefficients a_μ (which can be computed using Theorem 3.12, equation (3.11), or equation (3.13)). Then

$$\sum_{\mu \vdash n-1} a_\mu \mathbf{X}_{\mu,d} = 0.$$

Moreover, by varying ν, these relations generate all linear relations between the $\mathbf{X}_{\mu,d}$ for $\mu \vdash n-1$. (If we require that ν has a $(-d-3)$-fixed point, then the corresponding linear relations are themselves linearly independent.)

Proof. We have $F_{d+1}(\operatorname{pa}_\mu) = \mathbf{X}_{\nu,d+1} = \sum_{\mu \vdash n-1} a_\mu \operatorname{pa}_\mu$, and applying F_d gives the desired relation. The rest follows from exactness and Theorem 3.18. □

Example 3.29. Here is an example application of Theorem 3.28. Let $n = 10$, $d = 0$, and

$\nu = 3211 \vdash n-d-3$. Using (3.13), we have

$$\mathbf{X}_{\nu,d+1} = \mathbf{X}_{3211,1} = -[t^9]f_{211} + [t^9]\tilde{f}_3(t)f_{11}(t) - [t^9]\tilde{f}_{32}(t)f_1(t) + [t^9]\tilde{f}_{321}(t) - [t^9]\frac{\tilde{f}_{3211}(t)}{f(t)}$$
$$= \mathrm{pa}_9 + 2\,\mathrm{pa}_{81} + \mathrm{pa}_{711} - \mathrm{pa}_{63} - 2\,\mathrm{pa}_{54} - \mathrm{pa}_{531} - \mathrm{pa}_{441} + \mathrm{pa}_{432}\,.$$

Therefore, Theorem 3.28 implies that

$$\mathbf{X}_{\mathrm{Sp}(9)} + 2\mathbf{X}_{\mathrm{Sp}(81)} + \mathbf{X}_{\mathrm{Sp}(711)} - \mathbf{X}_{\mathrm{Sp}(63)} - 2\mathbf{X}_{\mathrm{Sp}(54)} - \mathbf{X}_{\mathrm{Sp}(531)} - \mathbf{X}_{\mathrm{Sp}(441)} + \mathbf{X}_{\mathrm{Sp}(432)} = 0$$

(where we used the fact that $\mathbf{X}_{\mu,0} = \mathbf{X}_{\mathrm{Sp}(\mu)}$). Note that $\mathbf{X}_{\mathrm{Sp}(\mu)} = \mathrm{pa}_9$ for $\mu \vdash 9$ and $\ell(\mu) \leq 2$, so we can cancel out all such terms. This gives

$$\mathbf{X}_{\mathrm{Sp}(711)} - \mathbf{X}_{\mathrm{Sp}(531)} - \mathbf{X}_{\mathrm{Sp}(441)} + \mathbf{X}_{\mathrm{Sp}(432)} = 0,$$

a nontrivial linear relation between CSFs of spiders of order 10.[5] ◊

As an edge case, note that if $n-d-3 < 0$, then no partitions ν exist in the statement of Theorem 3.28. Correspondingly, when $n-d-3 < 0$, the $\mathbf{X}_{\mu,d}$ for $\mu \vdash n-1$ are linearly independent. (This is due to Theorem 3.25: we have $-d-2 < -(n-1)$, so every partition of $n-1$ has a $(-d-2)$-fixed point.)

Here is another way to generate relations between the $\mathbf{X}_{\mu,d}$, this time by using some additional symmetric function theory.

Theorem 3.30. Fix integers $n \geq 1$ and d. Choose any partition $\lambda = (\lambda_1, \ldots, \lambda_\ell) \vdash n-1$ without a $(-d-2)$-fixed point. Then we have[6]

$$\sum_{\mu \vdash n-1} \langle s_\lambda, m_\mu \rangle \, \mathbf{X}_{\mu,d} = 0.$$

[5] In any such relation (with $d = 0$), all the occurrences of the path must cancel out because, as previously noted, pa_n is not in the span of the $\mathbf{X}_{\mathrm{Sp}(\mu)}$.

[6] The coefficients $\langle s_\lambda, m_\mu \rangle$ are called the *inverse Kostka numbers* and are sometimes denoted $K^{-1}_{\mu\lambda}$.

Furthermore, by varying λ, this generates all linear relations between $\mathbf{X}_{\mu,d}$ for $\mu \vdash n-1$, and these relations are themselves linearly independent.

Proof. First, we show that $s_{\lambda'}$ is in the image of the Bernstein operator B_{d+2}. Since λ does not have a $(-d-2)$-fixed point, by Lemma 3.17, λ' has a $(d+1)$-fixed point. Let $\ell(\lambda') = k$ and $\lambda'_i = i+d+1$ for some $i \in [k]$. Then in the determinantal formula (3.16) for $s_{\lambda'}$, the ith row is $(h_{d+2}, h_{d+3}, \ldots, h_{d+k+1})$. By swapping rows to put this row on top while keeping the other rows in the same relative order, we obtain

$$s_{\lambda'} = (-1)^{i-1} s_{(d+2,\lambda'_1+1,\ldots,\lambda'_{i-1}+1,\lambda'_{i+1},\ldots,\lambda'_k)}.$$

Therefore, Bernstein's result (3.17) shows that $s_{\lambda'}$ is in the image of B_{d+2}. It follows that $(F_d \circ U)(s_{\lambda'}) = 0$ by Theorem 3.24 and (3.18). Now we have

$$s_{\lambda'} = \omega(s_\lambda) = \omega\left(\sum_\mu \langle s_\lambda, m_\mu\rangle\, h_\mu\right) = \sum_\mu \langle s_\lambda, m_\mu\rangle\, e_\mu$$

since the m-basis and h-basis are dual with respect to the Hall inner product. Applying U sends e_μ to pa_μ, and then applying F_d sends pa_μ to $\mathbf{X}_{\mu,d}$, giving the desired relation.

To conclude, note that the number of possible $s_{\lambda'}$ is exactly $\left|\mathrm{Par}_{n-1}^{(d+1)}\right|$. By Lemma 3.19, Theorem 3.25, and Theorem 3.24, this is equal to $\left|\mathrm{Par}_{n-d-3}^{(-d-3)}\right| = \operatorname{rank} F_{d+1} = \operatorname{rank} B_{d+2}$ (where we consider F_{d+1} and B_{d+2} as maps $\mathrm{Sym}_{n-d-3} \to \mathrm{Sym}_{n-1}$). The $s_{\lambda'}$ are obviously linearly independent, so it follows that they form a basis for $\operatorname{im} B_{d+2}$, or equivalently, the $U(s_{\lambda'})$ form a basis for $\operatorname{im} F_{d+1} = \ker F_d$. □

In practice, the relations in Theorem 3.30 can be computed using the dual Jacobi-Trudi formula (2.3), as in the following example.

Example 3.31. Here is an example application of Theorem 3.30. Let $n = 10$, $d = 0$, and $\lambda = 432$ (which lacks a (-2)-fixed point). We compute $s_{\lambda'}$ by expanding the Jacobi-Trudi

determinant

$$s_{\lambda'} = \det\left[e_{\lambda_i - i + j}\right]_{1 \le i,j \le \ell} = \det \begin{bmatrix} e_4 & e_5 & e_6 \\ e_2 & e_3 & e_4 \\ e_0 & e_1 & e_2 \end{bmatrix}$$
$$= e_{432} + e_{54} + e_{621} - e_{63} - e_{522} - e_{441}.$$

Applying $F_0 \circ U$ amounts to replacing each e_μ with $\mathbf{X}_\mu$, so it follows that

$$\mathbf{X}_{\mathrm{Sp}(432)} + \mathbf{X}_{\mathrm{Sp}(54)} + \mathbf{X}_{\mathrm{Sp}(621)} - \mathbf{X}_{\mathrm{Sp}(63)} - \mathbf{X}_{\mathrm{Sp}(522)} - \mathbf{X}_{\mathrm{Sp}(441)} = 0,$$

or, after cancelling $\mathbf{X}_{\mathrm{Sp}(54)}$ and $\mathbf{X}_{\mathrm{Sp}(63)}$,

$$\mathbf{X}_{\mathrm{Sp}(432)} + \mathbf{X}_{\mathrm{Sp}(621)} - \mathbf{X}_{\mathrm{Sp}(522)} - \mathbf{X}_{\mathrm{Sp}(441)} = 0,$$

another nontrivial relation between CSFs of spiders of order 10. ◊

3.4 Future directions

To conclude, we discuss two possible directions for future work based on our results.

First, it would be interesting to use the star and path bases to study the CSFs of other subclasses of trees. In [17], the star basis was used in part to prove several structural results regarding the family of caterpillars. Namely, by using the star basis, the authors showed that the space spanned by the CSFs of caterpillars is exactly the same as that spanned by the CSFs of all trees (unlike in the case of spiders, whose CSFs span a strict subspace), and constructed a basis of caterpillars for this space. To our knowledge, however, it is not precisely understood what kinds of linear relations exist between the CSFs of caterpillars. Finding ways to generate or characterize such linear relations, whether for caterpillars or other families of trees, could be a fruitful approach to attack Question 1.1 (since, after all, an equality between two CSFs of trees is simply a linear relation of the form $\mathbf{X}_T - \mathbf{X}_{T'} = 0$).

The star and path bases may be particularly well suited for these purposes, since they are defined using CSFs and possess many appealing properties (including our Theorem 3.6).

It would also be interesting to revisit the problem of recovering caterpillars from their chromatic symmetric functions. In [3], it was shown that *proper* caterpillars are distinguished by the CSF; the method of this paper relied on a very difficult result, namely the characterization of when two ribbon Schur functions are equal [9]. (Note that by applying Theorem 3.6, we immediately get that the CSF distinguishes among the "dual" class of caterpillars whose complements are proper. Unfortunately, most caterpillars are neither proper nor have proper complements.) In addition, a proof is claimed in [22] that *all* caterpillars are distinguished by the CSF, but their method is rather *ad hoc* and does not appear to easily generalize to other kinds of trees.[7] It is possible that by studying the CSFs of caterpillars when expanded in chromatic bases, one could find more elementary or elegant proofs of these results.

Second, we present a group of related conjectures regarding the space of spiders, which we studied in Section 3.3. Based on computer calculations, there appear to exist other "nice" bases for the subspace $\mathcal{S}_{n,d} = \operatorname{span}\{\mathbf{X}_{\mu,d} : \mu \vdash n-1\}$ besides the one given in Theorem 3.25. These correspond to additional enumerative results regarding k-fixed points of partitions, which we now explain.

Given a partition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$ with parts written in weakly decreasing order as usual, its *crank* is an integer defined in the following way:

$$\operatorname{crank}(\lambda) = \begin{cases} \lambda_1 & \text{if } m_1(\lambda) = 0, \\ \#\{i : \lambda_i > m_1(\lambda)\} - m_1(\lambda) & \text{if } m_1(\lambda) > 0. \end{cases}$$

For instance, we have $\operatorname{crank}(6532111) = 2 - 3 = -1$ because 6532111 has three 1s and two parts greater than 3.

---

[7] Actually, I have been unable to verify that the argument in [22] is correct. In the future, I plan to post a discussion of this matter to my webpage.

It can be shown for $n \geq 2$ that the number of partitions of n with crank k is equal to the number of partitions of n with crank $-k$; see [8, Appendix A] for a bijective proof. Additionally, the number of partitions of n with a k-fixed point is equal to the number of partitions of n with crank greater than k [20, Theorem 3.2] (or, equivalently, the number with crank less than $-k$). This means that, at least in terms of dimension-counting, the following conjectures are plausible:

Conjecture 3.32. For all integers $n \geq 3$ and d, the symmetric functions

$$\{\mathbf{X}_{\mu,d} : \mu \vdash n-1, \operatorname{crank}(\mu) < d+2\}$$

form a basis for $\mathcal{S}_{n,d}$.

Conjecture 3.33. For all integers $n \geq 3$ and d, the symmetric functions

$$\{\mathbf{X}_{\mu,d} : \mu \vdash n-1, \operatorname{crank}(\mu^c) > -(d+2)\}$$

form a basis for $\mathcal{S}_{n,d}$.

Conjecture 3.34. For all integers $n \geq 3$, the symmetric functions

$$\{\mathbf{X}_{\mathrm{Sp}(\mu)} : \mu \vdash n-1, \operatorname{crank}(\mu, 1) < 0\}$$

form a basis for $\mathcal{S}_{n,0} = \operatorname{span}\{\mathbf{X}_S : S \in \mathrm{Sp}'_n\}$ (where $\operatorname{crank}(\mu, 1)$ means the crank of the partition $(\mu_1, \ldots, \mu_{\ell(\mu)}, 1)$).[8]

Proving any of these conjectures would give an interesting algebraic parallel to the enumerative results we have just discussed.

[8] To check the size of this conjectured basis: note that the number of $\mu \vdash n-1$ with $\operatorname{crank}(\mu, 1) < 0$ is equal to the number of $\lambda \vdash n$ with $\operatorname{crank}(\lambda) < 0$, since any partition with negative crank must have at least one 1. This is equal to the number of partitions of n with a 0-fixed point, which by Lemma 3.19 is equal to the number of partitions of $n-1$ with a (-2)-fixed point, and that is the dimension of $\mathcal{S}_{n,0}$.

Chapter 4

POLYNOMIAL TREE INVARIANTS

We now shift our attention to studying other tree invariants that are determined by the chromatic symmetric function. Unlike the chromatic symmetric function, these invariants are polynomials in finitely many variables and are therefore easier to work with (as we will try to demonstrate).

The advantages of studying these invariants are twofold. First, because they are determined by the chromatic symmetric function, any data about a tree that one can recover from such an invariant is also data that can be recovered from the chromatic symmetric function. Collecting such data could be a fruitful approach for obtaining a positive answer to Stanley's question. Second, by finding instances in which these invariants fail to distinguish trees (that is, pairs of trees with the same value of a particular invariant), we can construct good candidates for pairs of trees that may have the same chromatic symmetric function and then see whether in fact they do. Considering such candidates could be a fruitful approach for obtaining a negative answer to Stanley's question.

Our primary focus for the remainder of this thesis will be a three-variable tree invariant introduced by Logan Crew [11] called the *generalized degree polynomial* (GDP), denoted $\mathrm{G}_T(x, y, z)$, which enumerates subsets of the vertices of a tree T according to size and number of internal and boundary edges. In this chapter, we define the GDP, give a simple proof that it is determined by the chromatic symmetric function, and present several types of tree data that it determines. We also consider some specializations and further generalizations of the GDP. Then, in Chapter 5 and Chapter 6, we build the theory of the GDPs of rooted trees and polarized trees, which we use to derive a recurrence relation for the (ordinary) GDP and give several constructions for families of trees with the same GDP.

Besides the first section, the content of this chapter is largely drawn from the paper [21], which is joint work with Ricky Liu.

4.1 Two-variable tree invariants

Before introducing the generalized degree polynomial, we first consider a trio of weaker graph invariants: the *subtree polynomial*, *connector polynomial*, and *half-generalized degree polynomial*. These are two-variable polynomial invariants that are determined by the CSF (in fact, we will see that they are also determined by the GDP) and which were among the first to be used to recover information about a tree from its CSF. In addition to providing some historical context, we present a new formula relating these invariants, and prove a conjecture on the relationship between these polynomials and the CSF.

Let T be a tree. A *subtree* of T is a connected subgraph of T (with nonempty vertex set); we denote the collection of subtrees of T by $\mathcal{S}(T)$. For a subtree $S \in \mathcal{S}(T)$, a *boundary edge* of S is an edge of T with exactly one endpoint in S, and an *internal edge* of S is an edge of T with both endpoints in S (that is, an edge of S when considered as a tree in its own right). Also recall that a leaf edge of S is an edge of S incident to at least one leaf of S. Let $d(S)$, $e(S)$, and $\ell(S)$ be the number of boundary edges, internal edges, and leaf edges of S, respectively. (Note that $d(S)$ depends on T as well as S, but $e(S)$ and $\ell(S)$ can be computed only knowing the isomorphism type of S.)

We need one more definition. For every nonempty set of edges $\varnothing \neq A \subseteq E(T)$, there is a unique minimal subset $K(A) \subseteq E(T) \setminus A$ such that $A \cup K(A)$ is the edge set of a subtree of T. We call $K(A)$ the *connector* of A.

Definition 4.1 ([4, 13, 25])**.** Let T be a tree.

(i) The *subtree polynomial* of T, denoted $\mathrm{S}_T(q, r)$, is defined by

$$\mathrm{S}_T(q, r) = \sum_{S \in \mathcal{S}(T)} q^{e(s)} r^{\ell(S)}.$$

(ii) The *connector polynomial* of T, denoted $\mathrm{K}_T(x, y)$, is defined by

$$\mathrm{K}_T(x, y) = \sum_{\varnothing \neq A \subseteq E(T)} x^{|A|} y^{|K(A)|}.$$

(iii) The *half-generalized degree polynomial* of T, denoted $\mathrm{H}_T(y, z)$, is defined by

$$\mathrm{H}_T(y, z) = \sum_{S \in \mathcal{S}(T)} y^{d(S)} z^{e(S)}.$$

We note that our definition of S_T differs slightly from the one given in [25] because we allow the subtree S to have one vertex and no edges. (However, our definition agrees with the later definition of S_T given in [4].)

Example 4.2. Let T be the five-vertex tree depicted in Figure 4.1. We label the vertices and edges of T as shown so that we can refer to them. Then

$$\begin{aligned}
\mathrm{S}_T(q, r) &= 5 + 4qr + 4q^2r^2 + 2q^3r^2 + q^3r^3 + q^4r^3, \\
\mathrm{K}_T(x, y) &= 4x + 4x^2 + 2x^2y + 3x^3 + x^3y + x^4, \\
\mathrm{H}_T(y, z) &= 3y + y^2 + y^3 + yz + 2y^2z + y^3z + yz^2 + 3y^2z^2 + 3yz^3 + z^4.
\end{aligned}$$

For instance, if S is the subtree of T with vertex set $\{a, b, c\}$, then S has one boundary edge, namely h, and two internal edges, f and g, both of which are leaf edges. Thus, S contributes q^2r^2 to $\mathrm{S}_T(q, r)$ and yz^2 to $\mathrm{H}_T(y, z)$. As another example, for the subset of edges $A = \{f, i\}$, we have $K(A) = \{h\}$, so A contributes x^2y to $\mathrm{K}_T(x, y)$. ◊

The subtree polynomial was introduced by Eisenstat and Gordon [13]. Later, Martin, Morin, and Wagner studied the subtree and connector polynomials and their relationship with the chromatic symmetric function [25]. Specifically, they showed that the subtree and connector polynomials contain equivalent information about a tree, in the sense that each

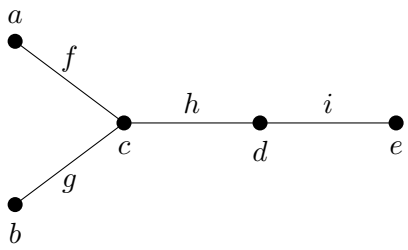

Figure 4.1: The tree T used in Example 4.2.

one can be obtained from the other via the pair of identities

$$\mathrm{S}_T(q,r) = n + \mathrm{K}_T(qr, q(1-r)) \quad \text{and} \quad \mathrm{K}_T(x,y) = -n + \mathrm{S}_T(x+y, x/(x+y)) \tag{4.1}$$

for $|T| = n$ (slightly modified from [25, Proposition 4] to match our definition of S_T). Furthermore, they found formulas expressing $\mathrm{S}_T(q,r)$ and $\mathrm{K}_T(q,r)$ as linear combinations of the coefficients of $\mathbf{X}_T$ in the power sum basis [25, Theorem 1], thus showing that S_T and K_T are determined linearly by the CSF. Using this fact, they then showed that $\mathbf{X}_T$ determines the *degree sequence* of T (the number of vertices of each degree) and the *path sequence* of T (the number of paths of each length).

The half-generalized degree polynomial H_T was introduced after Crew conjectured that the CSF $\mathbf{X}_T$ determines the generalized degree polynomial G_T [11]. As the name may suggest, H_T is a specialization of G_T, and it was proved by Wang, Xu, and Zhang [35] that $\mathbf{X}_T$ determines H_T. Later, Aliste-Prieto et al. showed (in the same paper settling Crew's conjecture) that H_T is equivalent to S_T and K_T [4, Theorem 8]. However, their method did not directly lead to any identities relating H_T to either S_T or K_T (analogous to (4.1)). We now present a pair of such identities.

Theorem 4.3. For any tree T of order n, we have

$$\mathrm{H}_T(y,z) = (n-1)(y-1) + \mathrm{S}_T(z, 1 + (y-1)/z) \tag{4.2}$$

and

$$\mathrm{S}_T(q,r) = -(n-1)q(r-1) + \mathrm{H}_T(qr-q+1,q). \tag{4.3}$$

Proof. The proof idea is similar in spirit to that of [4, Theorem 8], but we present it in a way that allows us to deduce the claimed identities. Let $D(S)$ denote the set of boundary edges of a subtree $S \in \mathcal{S}(T)$. We manipulate the sum defining H_T as follows:

$$\mathrm{H}_T(y,z) = \sum_{S\in\mathcal{S}(T)} y^{d(S)} z^{e(S)} = \sum_{S\in\mathcal{S}(T)} \sum_{A\subseteq D(S)} (y-1)^{|A|} z^{e(S)}.$$

Next, we reparametrize this double summation. Let $\hat{S}$ be the subtree with edge set $E(S)\cup A$. Then the elements of A are leaf edges of $\hat{S}$. Thus, we get a map

$$\{(S,A) : S \in \mathcal{S}(T),\, A \subseteq D(S)\} \to \{(\hat{S},A) : \hat{S} \in \mathcal{S}(T),\, A \subseteq L(\hat{S})\},$$

where $L(\hat{S})$ is the set of leaf edges of $\hat{S}$, such that $e(S) = e(\hat{S}) - |A|$. This map is surjective, but not quite injective. The only instance of non-injectivity is that, given an edge $e = \{u,v\}$ of T, the inputs $(\{u\},\{e\})$ and $(\{v\},\{e\})$ have the same output $(\{u,v\},\{e\})$. (Here we are representing a subtree by its vertex set.) Therefore, to obtain a bijection, we should exclude one pair of the form $(\{u\},\{e\})$ from the domain for each edge $e = \{u,v\}$. There are $n-1$ such pairs, and each contributes $(y-1)^1 z^0 = y-1$ to the double summation, so we get

$$\begin{aligned}
\mathrm{H}_T(y,z) &= (n-1)(y-1) + \sum_{\hat{S}\in\mathcal{S}(T)} \sum_{A\subseteq L(\hat{S})} (y-1)^{|A|} z^{e(\hat{S})-|A|} \\
&= (n-1)(y-1) + \sum_{\hat{S}\in\mathcal{S}(T)} \sum_{A\subseteq L(\hat{S})} ((y-1)/z)^{|A|} z^{e(\hat{S})} \\
&= (n-1)(y-1) + \sum_{\hat{S}\in\mathcal{S}(T)} (1+(y-1)/z)^{\ell(\hat{S})} z^{e(\hat{S})} \\
&= (n-1)(y-1) + \mathrm{S}_T(z, 1+(y-1)/z).
\end{aligned}$$

This is exactly (4.2), and (4.3) follows by putting $(y, z) = (qr - q + 1, q)$. □

4.1.1 Proof of a positivity conjecture

We now prove a conjecture of Martin, Morin, and Wagner regarding their formula expressing $\mathrm{K}_T(x, y)$ in terms of $\mathbf{X}_T$ [25, Theorem 1]. First, let us write this formula explicitly in terms of the Hall inner product on Sym, as follows:

Theorem 4.4 ([25, p. 246])**.** Fix a positive integer n, and define

$$\Psi_n(x, y) = \sum_{a \geq 1} \sum_{b \geq 0} x^a y^b \sum_{\lambda \vdash n} \psi(\lambda, a, b) \frac{p_\lambda}{z_\lambda},$$

where the coefficients $\psi(\lambda, a, b)$ are given by

$$\psi(\lambda, a, b) = (-1)^{a+b} \binom{\ell(\lambda) - 1}{n - a - b - 1} \sum_{k=1}^{\ell(\lambda)} \binom{\lambda_k - 1}{a}.$$

Then for all trees T of order n, we have $\mathrm{K}_T(x, y) = \langle \Psi_n(x, y), \mathbf{X}_T \rangle$.

(Note that x and y are treated as constants when computing the Hall inner product.)

The conjecture concerns the expansion of $\Psi_n(x, y)$ into the complete homogeneous basis $\{h_\lambda\}$ of Sym. Define the rational numbers $\xi(\mu, a, b)$ for $\mu \vdash n$ such that

$$\Psi_n(x, y) = \sum_{a \geq 1} \sum_{b \geq 0} x^a y^b \sum_{\mu \vdash n} \xi(\mu, a, b) h_\mu.$$

(In other words, $\sum_{\lambda \vdash n} \psi(\lambda, a, b) \frac{p_\lambda}{z_\lambda} = \sum_{\mu \vdash n} \xi(\mu, a, b) h_\mu$ for all a and b.) Then the conjecture states that $(-1)^{n-\ell(\mu)} \xi(\mu, a, b) z_\mu$ is a nonnegative integer for all a, b, and μ [25, Conjectures 6 and 7]. We will prove this conjecture but with z_μ replaced by $\mathrm{aut}(\mu)$:

Theorem 4.5. For all a, b, and $\mu \vdash n$, $(-1)^{n-\ell(\mu)} \xi(\mu, a, b) \, \mathrm{aut}(\mu)$ is a nonnegative integer.

(Because $\mathrm{aut}(\mu)$ always divides z_μ, this is a slightly stronger version of the conjecture.)

The proof proceeds in several steps. We first find a formula for $(-1)^{n-\ell(\mu)}\xi(\mu,a,b)\operatorname{aut}(\mu)$ that shows it is an integer, and then prove that it is nonnegative.

We will use the following notation: if α is a composition (or partition) and I is a subset of $[\ell(\alpha)]$ (or a cycle of a permutation on $[\ell(\alpha)]$), let $\alpha_I = \sum_{i\in I}\alpha_i$. For any permutation $w \in \mathfrak{S}_n$, let $\operatorname{Cyc}(w)$ be the set of cycles of w and $\operatorname{cyc}(w) = |\operatorname{Cyc}(w)|$. We use $\ell(C)$ to denote the length (number of elements) of a cycle C.

Here are two computational lemmas we will need.

Lemma 4.6. Given $\lambda, \mu \vdash n$, the coefficient of h_μ in the h-expansion of p_λ/z_λ is

$$(-1)^{\ell(\lambda)-\ell(\mu)}\frac{|\mathfrak{S}_{\lambda,\mu}|}{\operatorname{aut}(\mu)},$$

where $\mathfrak{S}_{\lambda,\mu}$ is the set of all $w \in \mathfrak{S}_{\ell(\mu)}$ such that the numbers μ_C for $C \in \operatorname{Cyc}(w)$ are the parts of λ in some order. (In particular, $\operatorname{cyc}(w) = \ell(\lambda)$ for all such w.)

Proof. Let $\ell(\lambda) = k$ and $\ell(\mu) = \ell$.

We use the identity [6, (11)], which expresses a noncommutative analog of the power sum basis in terms of the noncommutative homogeneous basis. (See [16] for background on noncommutative symmetric functions.) The symmetric version of this identity is

$$p_\lambda = \operatorname{prod}(\lambda)\sum_{\beta\preceq\lambda}(-1)^{\ell(\beta)-\ell(\lambda)}\frac{1}{\ell(\beta,\lambda)}h_\beta.$$

Here, $\beta \preceq \lambda$ (β *refines* λ) means that one can combine groups of adjacent parts of β to obtain λ (or equivalently, $\operatorname{Set}(\lambda) \subseteq \operatorname{Set}(\beta)$). For a composition $\beta = \beta^{(1)}\cdots\beta^{(k)}$ refining λ, where $|\beta^{(i)}| = \lambda_i$, we define $\ell(\beta,\lambda) = \ell(\beta^{(1)})\cdots\ell(\beta^{(k)})$. Then the coefficient of h_μ in p_λ is

$$[h_\mu]p_\lambda = \operatorname{prod}(\lambda)\sum_{\substack{\beta\preceq\lambda\\ \beta\sim\mu}}(-1)^{\ell(\beta)-\ell(\lambda)}\frac{1}{\ell(\beta,\lambda)} = (-1)^{k-\ell}\operatorname{prod}(\lambda)\sum_{\substack{\beta\preceq\lambda\\ \beta\sim\mu}}\frac{1}{\ell(\beta,\lambda)}.$$

Let $B_{\lambda,\mu}$ be the set of all compositions β such that $\beta \preceq \lambda$ and $\beta \sim \mu$. (For instance, we have $B_{24,2211} = \{2211, 2121, 2112, 1122\}$.) We want to relate $B_{\lambda,\mu}$ with $\mathfrak{S}_{\lambda,\mu}$ so that we can

use a double-counting argument. In order to be precise in doing this, we define a third set with surjections to both $B_{\lambda,\mu}$ and $\mathfrak{S}_{\lambda,\mu}$.

Let $U_{\lambda,\mu}$ be the set of all k-tuples $\mathbf{u} = (\mathbf{u}^{(1)}, \ldots, \mathbf{u}^{(k)})$ satisfying the following conditions:

(i) The $\mathbf{u}^{(i)}$ are sequences of integers that, together, contain every integer in $[\ell]$ exactly once (and no others).

(ii) For $i = 1, \ldots, k$, we have $\mu_{\mathbf{u}^{(i)}} = \lambda_i$.

There is a surjection $U_{\lambda,\mu} \to B_{\lambda,\mu}$ that sends $\mathbf{u} = (\mathbf{u}^{(1)}, \ldots, \mathbf{u}^{(k)})$ to the β obtained by concatenating the compositions $(\mu_j)_{j \in \mathbf{u}^{(i)}}$ for $i = 1, \ldots, k$. (In other words, to get β, we combine the sequences $\mathbf{u}^{(i)}$, giving a permutation of $[\ell]$, and read off the corresponding parts of μ in order.) Note that then $\ell(\beta, \lambda) = \ell(\mathbf{u}^{(1)}) \cdots \ell(\mathbf{u}^{(k)})$. Under this surjection, the fiber of any $\beta \in B_{\lambda,\mu}$ has size $\mathrm{aut}(\mu)$, since indices corresponding to equal parts of μ can be permuted without changing β.

Also, there is a surjection $U_{\lambda,\mu} \to \mathfrak{S}_{\lambda,\mu}$ that sends $\mathbf{u} = (\mathbf{u}^{(1)}, \ldots, \mathbf{u}^{(k)})$ to the permutation of $[\ell]$ whose cycles are the sequences $\mathbf{u}^{(i)}$, wrapping around from the last element to the first. Under this surjection, the fiber of any $w \in \mathfrak{S}_{\lambda,\mu}$ has size $\mathrm{aut}(\lambda) \prod_{C \in \mathrm{Cyc}(w)} \ell(C)$, since there are $\mathrm{aut}(\lambda)$ ways to permute cycles C with equal values of μ_C, and for each cycle C there are $\ell(C)$ sequences (linear orders) corresponding to it. Alternatively, for any $\mathbf{u} = (\mathbf{u}^{(1)}, \ldots, \mathbf{u}^{(k)})$ in the fiber of w, the size of the fiber is $\mathrm{aut}(\lambda) \cdot \ell(\mathbf{u}^{(1)}) \cdots \ell(\mathbf{u}^{(k)})$.

Using these two maps, we obtain

$$\begin{aligned}
[h_\mu]p_\lambda &= (-1)^{k-\ell} \operatorname{prod}(\lambda) \sum_{\beta \in B_{\lambda,\mu}} \frac{1}{\ell(\beta, \lambda)} \\
&= (-1)^{k-\ell} \operatorname{prod}(\lambda) \cdot \frac{1}{\mathrm{aut}(\mu)} \sum_{\mathbf{u} \in U_{\lambda,\mu}} \frac{1}{\ell(\mathbf{u}^{(1)}) \cdots \ell(\mathbf{u}^{(k)})} \\
&= (-1)^{k-\ell} \operatorname{prod}(\lambda) \cdot \frac{1}{\mathrm{aut}(\mu)} \sum_{w \in \mathfrak{S}_{\lambda,\mu}} \mathrm{aut}(\lambda) \\
&= (-1)^{k-\ell} z_\lambda \frac{|\mathfrak{S}_{\lambda,\mu}|}{\mathrm{aut}(\mu)}.
\end{aligned}$$

Dividing by z_λ gives the result. □

Lemma 4.7. For nonnegative integers n and k, let $T(n,k) = \sum_{w\in\mathfrak{S}_n}(-1)^{\mathrm{cyc}(w)-k}\binom{\mathrm{cyc}(w)}{k}$. Then

$$T(n,k) = \begin{bmatrix} n \\ k \end{bmatrix} - n\begin{bmatrix} n-1 \\ k \end{bmatrix}$$

where $\begin{bmatrix} n \\ k \end{bmatrix} = \#\{w \in \mathfrak{S}_n : \mathrm{cyc}(w) = k\}$ is an unsigned Stirling number of the first kind.

Proof. We have

$$T(n,k) = \sum_{w\in\mathfrak{S}_n}(-1)^{\mathrm{cyc}(w)-k}\binom{\mathrm{cyc}(w)}{k} = \sum_{w\in\mathfrak{S}_n}\sum_{A\in\binom{\mathrm{Cyc}(w)}{k}}(-1)^{\mathrm{cyc}(w)-k}.$$

Given w and A, let w' be the permutation on a subset of $[n]$ whose cycles are the elements of A. Then the remaining cycles of w form another permutation w'' with $\mathrm{cyc}(w'') = \mathrm{cyc}(w)-k$. Therefore, we can write

$$T(n,k) = \sum_{w'}\sum_{w''}(-1)^{\mathrm{cyc}(w'')},$$

where w' ranges over permutations of some $B \subseteq [n]$ with k cycles and w'' ranges over permutations of $[n] \setminus B$. It is well-known that the inner sum is zero unless the ground set $[n] \setminus B$ has size 0 or 1 (in which case the sum equals 1 or -1, respectively). If $[n] \setminus B$ has size 0, then $B = [n]$ and w' is any permutation on $[n]$ with k cycles. If $[n] \setminus B$ has size 1, then there are n possible sets B, and w' is a permutation on a set of size $n-1$ with k cycles. The conclusion follows. □

We can now give a formula for the quantity we are considering. Fix $a \geq 1$, $b \geq 0$, and $n \geq 1$ such that $a + b \leq n - 1$. Let $c = n - a - b - 1 \geq 0$. Let $\mu = (\mu_1, \ldots, \mu_\ell) \vdash n$.

Theorem 4.8. We have

$$(-1)^{n-\ell}\xi(\mu,a,b)\,\mathrm{aut}(\mu) = \sum_{m=1}^{\ell}(m-1)!\,\sigma_m\,T(\ell-m, n-a-b-1), \tag{4.4}$$

where $\sigma_m = \sum_{I \in \binom{[\ell]}{m}} \binom{|\mu_I|-1}{a}$.

Proof. By definition, $\xi(\mu, a, b)$ is the coefficient of h_μ in the h-expansion of $\sum_{\lambda \vdash n} \psi(\lambda, a, b) \frac{p_\lambda}{z_\lambda}$. By Lemma 4.6 and the definition of $\psi(\lambda, a, b)$, this equals

$$\begin{aligned}\xi(\mu, a, b) &= \sum_{\lambda \vdash n} \psi(\lambda, a, b)(-1)^{\ell(\lambda)-\ell} \frac{|\mathfrak{S}_{\lambda,\mu}|}{\operatorname{aut}(\mu)} \\ &= \frac{(-1)^{a+b-\ell}}{\operatorname{aut}(\mu)} \sum_{\lambda \vdash n} (-1)^{\ell(\lambda)} \binom{\ell(\lambda)-1}{n-a-b-1} \sum_{k=1}^{\ell(\lambda)} \binom{\lambda_k - 1}{a} \sum_{w \in \mathfrak{S}_{\lambda,\mu}} 1.\end{aligned}$$

We now interchange the order of summation, so that $w \in \mathfrak{S}_\ell$ and λ is the partition of n whose parts are the numbers μ_C for $C \in \operatorname{Cyc}(w)$. Then

$$\xi(\mu, a, b) = \frac{(-1)^{a+b-\ell}}{\operatorname{aut}(\mu)} \sum_{w \in \mathfrak{S}_\ell} (-1)^{\operatorname{cyc}(w)} \binom{\operatorname{cyc}(w)-1}{n-a-b-1} \sum_{C \in \operatorname{Cyc}(w)} \binom{\mu_C - 1}{a}.$$

Given w and C, let w' be the permutation of $[\ell] \setminus C$ obtained by deleting C, so that $\operatorname{Cyc}(w') = \operatorname{Cyc}(w) \setminus \{C\}$. Interchanging the order of summation once again, we get

$$\xi(\mu, a, b) = \frac{(-1)^{a+b-\ell}}{\operatorname{aut}(\mu)} \sum_{C} \binom{\mu_C - 1}{a} \sum_{w' \in \mathfrak{S}([\ell] \setminus C)} (-1)^{\operatorname{cyc}(w')+1} \binom{\operatorname{cyc}(w')}{n-a-b-1},$$

where C ranges over all cycles on nonempty subsets of $[\ell]$. Let I be the ground set of C and let $m = |I|$. Then each I corresponds to $(m-1)!$ cycles C. Note that the inner sum only depends on the value of m. Hence, we can instead sum over m, as follows:

$$\begin{aligned}\xi(\mu, a, b) &= \frac{(-1)^{a+b-\ell}}{\operatorname{aut}(\mu)} \sum_{m=1}^{\ell} (m-1)! \sum_{I \in \binom{[\ell]}{m}} \binom{\mu_I - 1}{a} \sum_{w' \in \mathfrak{S}_{\ell-m}} (-1)^{\operatorname{cyc}(w')+1} \binom{\operatorname{cyc}(w')}{n-a-b-1} \\ &= \frac{(-1)^{a+b-\ell}}{\operatorname{aut}(\mu)} \sum_{m=1}^{\ell} (m-1)! \cdot \sigma_m \cdot (-1)^{n-a-b} T(\ell-m, n-a-b-1) \\ &= \frac{(-1)^{n-\ell}}{\operatorname{aut}(\mu)} \sum_{m=1}^{\ell} (m-1)! \, \sigma_m \, T(\ell-m, n-a-b-1)\end{aligned}$$

using Lemma 4.7. □

The preceding theorem makes it clear that $(-1)^{n-\ell}\xi(\mu, a, b)\operatorname{aut}(\mu)$ is an integer. This proves half of Theorem 4.5. For the positivity, we need some estimates on the size of the numbers σ_m.

Lemma 4.9. For all $1 \leq m \leq \ell$, we have $\sigma_m \leq \binom{\ell-1}{m-1}\binom{n-1}{a}$.

Proof. With a, ℓ, m, and n all fixed, consider the left-hand side as a function $\sigma_m(\mu)$, where μ ranges over compositions of size n and length ℓ. We claim that $\sigma(\mu)$ is maximized for $\mu = (n-\ell+1, 1^{\ell-1})$. Suppose μ is a composition with at least two parts greater than one. Without loss of generality, assume that $\mu_1 \geq \mu_2 \geq 2$, and let $\nu = (\mu_1+1, \mu_2-1, \mu_3, \ldots, \mu_\ell)$. Then it will suffice to show $\sigma_m(\mu) \leq \sigma_m(\nu)$. In the difference $\sigma_m(\nu) - \sigma_m(\mu)$, corresponding terms $\binom{\nu_I - 1}{a}$ and $\binom{\mu_I - 1}{a}$ cancel except when $|I \cap \{1,2\}| = 1$. Therefore

$$\sigma_m(\nu) - \sigma_m(\mu) = \sum_{\substack{1\in I\\ 2\notin I}} \left[\binom{\mu_I}{a} - \binom{\mu_I - 1}{a}\right] + \sum_{\substack{1\notin I\\ 2\in I}} \left[\binom{\mu_I - 2}{a} - \binom{\mu_I - 1}{a}\right] \tag{4.5}$$

$$= \sum_{\substack{1\in I\\ 2\notin I}} \binom{\mu_I - 1}{a-1} - \sum_{\substack{1\notin I\\ 2\in I}} \binom{\mu_I - 2}{a-1}. \tag{4.6}$$

Setting $J = I \setminus \{1,2\}$ in both sums (so J is a size-$(m-1)$ subset of $[3,\ell]$), we get

$$\sigma_m(\nu) - \sigma_m(\mu) = \sum_J \left[\binom{\mu_J + \mu_1 - 1}{a-1} - \binom{\mu_J + \mu_2 - 2}{a-1}\right], \tag{4.7}$$

and each bracketed difference is nonnegative because $\mu_1 \geq \mu_2 \geq 2$. This proves the claim. Thus,

$$\sigma_m(\mu) \leq \sigma_m(n-\ell+1, 1^{\ell-1}) = \binom{\ell-1}{m-1}\binom{n-\ell+m-1}{a} + \binom{\ell-1}{m}\binom{m-1}{a}. \tag{4.8}$$

Since $a \geq 1$, the function $N \mapsto \frac{1}{N}\binom{N-1}{a}$ is (weakly) increasing for positive integers N. Thus,

$$\begin{aligned}\sigma_m(\mu) &\leq \binom{\ell-1}{m-1}\binom{n-\ell+m-1}{a} + \binom{\ell-1}{m}\binom{m-1}{a} \\ &\leq \binom{\ell-1}{m-1}\frac{n-\ell+m}{n}\binom{n-1}{a} + \frac{\ell-m}{m}\binom{\ell-1}{m-1}\cdot\frac{m}{n}\binom{n-1}{a} \\ &= \binom{\ell-1}{m-1}\binom{n-1}{a}\end{aligned}$$

because $m \leq \ell \leq n$, completing the proof. □

Corollary 4.10. If integers m and M satisfy $1 \leq m \leq M \leq \ell$, then

$$(m-1)!(\ell-m)!\sigma_m \leq (M-1)!(\ell-M)!\sigma_M. \tag{4.9}$$

Proof. First, fix $S \in \binom{[\ell]}{M}$. Applying the result of Lemma 4.9 but with μ replaced by the partition whose parts are μ_i for $i \in S$, we obtain

$$\sum_{T\in\binom{S}{m}} \binom{\mu_T - 1}{d} \leq \binom{M-1}{m-1}\binom{\mu_S - 1}{d}. \tag{4.10}$$

We now sum this over all S. On the right-hand side, we obtain $\binom{M-1}{m-1}\sigma_M$, while on the left-hand side, we get

$$\sum_{S\in\binom{[\ell]}{M}} \sum_{T\in\binom{S}{m}} \binom{\mu_T-1}{d} = \sum_{T\in\binom{[\ell]}{m}} \sum_{S:\, T\subseteq S\in\binom{[\ell]}{M}} \binom{\mu_T-1}{d} = \binom{\ell-m}{M-m}\sigma_m. \tag{4.11}$$

Hence $\binom{\ell-m}{M-m}\sigma_m \leq \binom{M-1}{m-1}\sigma_M$, which rearranges to the desired result. □

We are almost ready to prove the theorem, after one more (easy) lemma.

Lemma 4.11. Let $x_1, \ldots, x_\ell, y_1, \ldots, y_\ell$ be real numbers such that $0 \leq x_1 \leq \cdots \leq x_\ell$, and suppose there exist $z_1, \ldots, z_\ell \geq 0$ such that $y_i = z_i - z_{i+1}$ for all i (where $z_{\ell+1} = 0$). Then

$$x_1y_1 + \cdots + x_\ell y_\ell \geq 0.$$

Proof. We have

$$\begin{aligned} x_1y_1 + \cdots + x_\ell y_\ell &= x_1(z_1 - z_2) + x_2(z_2 - z_3) + \cdots + x_{\ell-1}(z_{\ell-1} - z_\ell) + x_\ell z_\ell \\ &= x_1z_1 + (x_2 - x_1)z_2 + (x_3 - x_2)z_3 + \cdots + (x_\ell - x_{\ell-1})z_\ell, \end{aligned}$$

whence the conclusion follows, as all the terms are nonnegative. □

Proof of Theorem 4.5. Theorem 4.8 states that

$$(-1)^{n-\ell}\xi(\mu, a, b)\operatorname{aut}(\mu) = \sum_{m=1}^{\ell}(m-1)!\,\sigma_m\,T(\ell - m, n - a - b - 1),$$

which is clearly an integer. To show that it is nonnegative, we claim that we can take

$$x_m = (m-1)!(\ell-m)!\sigma_m, \quad y_m = \frac{T(\ell - m, n - a - b - 1)}{(\ell - m)!}, \quad z_m = \frac{1}{(\ell - m)!}\begin{bmatrix} \ell - m \\ n - a - b - 1 \end{bmatrix}$$

for $1 \le m \le \ell$ in Lemma 4.11. The condition $0 \le x_1 \le \cdots \le x_\ell$ is Corollary 4.10, and the condition $y_i = z_i - z_{i+1}$ follows from Lemma 4.7. Therefore the lemma applies, giving

$$x_1y_1 + \cdots + x_\ell y_\ell = \sum_{m=1}^{\ell}(m-1)!\,\sigma_m\,T(\ell - m, n - a - b - 1) \ge 0,$$

as desired. □

4.2 The generalized degree polynomial

We now move to the generalized degree polynomial, our main focus. We begin by defining the GDP of a forest F, following [4, Section 2.2]. Let $A \subseteq V(F)$. An *internal edge* of A is an edge of F with both endpoints in A, and a *boundary edge* of A is an edge of F with exactly one endpoint in A. (Note that an internal edge of A is the same as an edge of $F|_A$.) Let $E_F(A)$ and $D_F(A)$ denote the sets of internal edges and boundary edges of A, respectively; we write $e_F(A) = |E_F(A)|$ and $d_F(A) = |D_F(A)|$. (When there is no risk of confusion, we

will omit the subscripts.) Observe that

$$d_F(A) + 2e_F(A) = \sum_{v \in A} \deg_F(v) \tag{4.12}$$

by a double-counting argument. Finally, for nonnegative integers a, b, c, let

$$g_F(a, b, c) = \#\{A \subseteq V(F) : \ |A| = a, \ d(A) = b, \ e(A) = c\}.$$

Definition 4.12. The **generalized degree polynomial (GDP)** of F is defined by

$$\mathrm{G}_F(x, y, z) = \sum_{A \subseteq V(F)} x^{|A|} y^{d(A)} z^{e(A)} = \sum_{a,b,c} g_F(a, b, c) x^a y^b z^c.$$

Thus the generalized degree polynomial enumerates subsets $A \subseteq V(F)$ by size and the number of internal and boundary edges. In particular, if $|A| = 1$, then A consists of a single vertex v such that $d(A) = \deg(v)$ and $e(A) = 0$. This implies that G_F determines the *degree sequence* of F: specifically, the number of vertices of F with degree d is just $g_F(1, d, 0)$.

Example 4.13. Suppose $F = \mathrm{St}_4$, the star with 4 vertices. Then

$$\mathrm{G}_F(x, y, z) = x^4z^3 + x^3y^3 + 3x^3yz^2 + 3x^2y^2z + 3x^2y^2 + xy^3 + 3xy + 1.$$

For example, the coefficient of x^3yz^2 in $\mathrm{G}_F(x, y, z)$ is $g_F(3, 1, 2) = 3$ because there are three subsets of vertices $A \subseteq V(F)$ such that $|A| = 3$, $d(A) = 1$, and $e(A) = 2$. (Specifically, these subsets consist of the center vertex of F and any two of its outer vertices.) $\lozenge$

From the definitions, it is clear that the generalized degree polynomial G_T of a tree T determines its half-generalized degree polynomial H_T, thus also its subtree and connector polynomials S_T and K_T (defined in Section 4.1). However, the converse is not true: there exist trees T_1 and T_2 such that $\mathrm{H}_{T_1} = \mathrm{H}_{T_2}$ but $\mathrm{G}_{T_1} \neq \mathrm{G}_{T_2}$. One such pair of trees is depicted in Figure 4.2 (and also appears in [4, Section 7.2]). We note that the results of

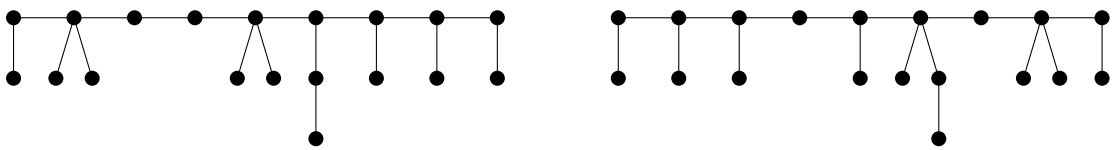

Figure 4.2: Two trees with the same half-generalized degree polynomial but different GDPs.

this chapter will allow us to easily show that these trees have different GDPs by considering their *double-degree sequences*.

Aliste-Prieto et al. [4] showed that for a tree T, the generalized degree polynomial G_T is determined linearly by the chromatic symmetric function $\mathbf{X}_T$, that is, they exhibit an explicit $\mathbb{Q}$-linear map $\mathrm{Sym} \to \mathbb{Q}(x, y, z)$ that sends $\mathbf{X}_T$ to G_T for every tree T. We now prove a slight generalization of this statement by using the Hopf algebra structure of Sym.

Proposition 4.14. There exists a $\mathbb{Q}$-algebra map $\gamma\colon \mathrm{Sym} \to \mathbb{Q}(x, y, z)$ such that for every forest F, we have

$$\gamma(\mathbf{X}_F) = y^{c(F)}\mathrm{G}_F(x, y, z).$$

(In particular, the chromatic symmetric function also linearly determines the generalized degree polynomial when restricted to forests with a fixed number of components.) To prove this statement, we will use the following lemma.

Lemma 4.15. There exists a $\mathbb{Q}$-algebra map $\varphi_{t,u}\colon \mathrm{Sym} \to \mathbb{Q}[t, u]$ such that for every forest F, we have

$$\varphi_{t,u}(\mathbf{X}_F) = t^{|F|}u^{e(F)}. \tag{4.13}$$

Proof. We claim that we can define $\varphi_{t,u}$ on the power sums p_k by $\varphi_{t,u}(p_k) = t^k(1-u)^{k-1}$ (for $k \geq 1$) and hence on the power sum basis $\{p_\lambda\}$ by

$$\varphi_{t,u}(p_\lambda) = t^{|\lambda|}(1-u)^{|\lambda|-\ell(\lambda)}.$$

To show (4.13), let F be a forest and let $n = |F|$. By Proposition 2.5 we have

$$\begin{aligned}\varphi_{t,u}(\mathbf{X}_F) &= \sum_{\lambda \vdash n} b_\lambda(F)(-1)^{n-\ell(\lambda)} \cdot t^n (1-u)^{n-\ell(\lambda)} \\ &= t^n \sum_{\lambda \vdash n} b_\lambda(F)(u-1)^{n-\ell(\lambda)}. \end{aligned} \tag{4.14}$$

Note that connected partitions of F are in bijection with subsets of $E(F)$, where $S \subseteq E(F)$ corresponds to the connected partition $\mathcal{C}$ giving the components of the spanning subgraph of F with edge set S. If the type of $\mathcal{C}$ is $\lambda \vdash n$, then we have $|S| = n - \ell(\lambda)$, so it follows that $\sum_{\ell(\lambda)=i} b_\lambda(F) = \binom{e(F)}{n-i}$. Therefore, letting $i = \ell(\lambda)$ in (4.14), we have

$$\varphi_{t,u}(\mathbf{X}_F) = t^n \sum_i \binom{e(F)}{n-i} (u-1)^{n-i} = t^n u^{e(F)},$$

as claimed. □

We now use Lemma 4.15 together with (2.9) to prove Proposition 4.14.

Proof of Proposition 4.14. Take $\gamma = \varphi_{xy,y^{-1}z} * \varphi_{y,y^{-1}}$. Given a forest F, we compute:

$$\begin{aligned}(\varphi_{xy,y^{-1}z} * \varphi_{y,y^{-1}})(\mathbf{X}_F) &= \sum_{A \sqcup B = V(F)} \varphi_{xy,y^{-1}z}(\mathbf{X}_{F|A}) \cdot \varphi_{y,y^{-1}}(\mathbf{X}_{F|B}) \\ &= \sum_{A \sqcup B = V(F)} (xy)^{|A|} (y^{-1}z)^{e(A)} \cdot y^{|B|} (y^{-1})^{e(B)} \\ &= \sum_{A \sqcup B = V(F)} x^{|A|}\, y^{|F|-e(A)-e(B)}\, z^{e(A)}.\end{aligned}$$

Each edge of F is either an internal edge of A, an internal edge of B, or a boundary edge of A (and these are mutually exclusive), so $e(A) + e(B) + d(A) = |E(F)| = |F| - c(F)$. Thus,

$$(\varphi_{xy,y^{-1}z} * \varphi_{y,y^{-1}})(\mathbf{X}_F) = \sum_{A \subseteq V(F)} x^{|A|}\, y^{c(F)+d(A)}\, z^{e(A)} = y^{c(F)} \mathrm{G}_F(x,y,z)$$

as desired. □

A couple of remarks are in order regarding the map γ constructed above. First, since γ is multiplicative, it is determined by the values of $\gamma(p_n)$ for $n \geq 1$. Since the comultiplication of Sym obeys $\Delta(p_n) = p_n \otimes 1 + 1 \otimes p_n$, we have

$$\begin{aligned}\gamma(p_n) &= (\varphi_{xy,y^{-1}z} * \varphi_{y,y^{-1}})(p_n) \\ &= \varphi_{xy,y^{-1}z}(p_n) \cdot \varphi_{y,y^{-1}}(1) + \varphi_{xy,y^{-1}z}(1) \cdot \varphi_{y,y^{-1}}(p_n) \\ &= (xy)^n(1 - y^{-1}z)^{n-1} + y^n(1 - y^{-1})^{n-1} \\ &= x^n y(y - z)^{n-1} + y(y - 1)^{n-1},\end{aligned}$$

which provides an alternative characterization of γ. Second, it can be shown by a direct (albeit tedious) calculation that γ agrees, up to scaling, with the $\mathbb{Q}$-linear map from [4] that sends $\mathbf{X}_T$ to G_T for trees T. Specifically, in their notation, one has

$$\gamma(p_\lambda) = y \cdot \sum_{a,b,c} \omega(\lambda, a, b, c) x^a y^b z^c.$$

In Section 4.5, we will discuss a further generalization of the generalized degree polynomial that is obtained by a similar construction.

4.2.1 Tree augmentation

By contrast with Question 1.1, the generalized degree polynomial does not distinguish trees. Figure 3.1 from the previous chapter shows the smallest pair of non-isomorphic trees with the same GDP; both trees have order 11.

The entirety of Chapter 6 will be devoted to presenting constructions that produce pairs (and larger families) of trees with the same GDP. For now, though, we present one possible such construction that does not require any special machinery. It takes a known pair of trees with the same GDP and modifies them in the same way, producing a pair of larger trees which also have the same GDP.

Let T be a tree and let k be a positive integer such that $k \geq \deg(v)$ for every $v \in V(T)$.

We define the *k-augmentation* of T, denoted $T^{[k]}$, to be the tree created by attaching $k - \deg(v)$ leaves to every vertex v of T. Thus, every internal vertex of $T^{[k]}$ has degree k, and the subtree of $T^{[k]}$ induced by the internal vertices is isomorphic to T. We therefore view T as embedded in $T^{[k]}$.

Proposition 4.16. Fix integers $n, k \geq 2$. Then there exists a $\mathbb{Q}$-linear map $\mathbb{Q}(x, y, z) \to \mathbb{Q}(x, y, z)$ that sends G_T to $\mathrm{G}_{T^{[k]}}$ for trees T of order n.

Proof. For convenience, let $U = T^{[k]}$. We show that each coefficient of G_U is a linear combination of coefficients of G_T. Let $a, b, c \geq 0$ and consider

$$g_U(a, b, c) = \#\{A \subseteq V(U) : \ |A| = a, \ d_U(A) = b, \ e_U(A) = c\}.$$

By (4.12), we have $b + 2c = d_U(A) + 2e_U(A) = \sum_{v \in A} \deg_U(v)$. But every vertex of U has degree either 1 (if it is a leaf of U) or k (if it is not a leaf of U and hence a vertex of T). Let $\ell_U(A)$ and $i_U(A)$ denote the number of vertices in A of each type, respectively. Then,

$$\begin{cases} \ell_U(A) + i_U(A) = a, \\ \ell_U(A) + k \cdot i_U(A) = b + 2c, \end{cases}$$

which has a unique solution $(\ell_U(A), i_U(A)) = (\ell, i)$ in terms of a, b, c. It follows that

$$g_U(a, b, c) = \#\{A \subseteq V(U) : \ \ell_U(A) = \ell, \ i_U(A) = i, \ e_U(A) = c\}.$$

Now let $A_0 = A \cap V(T)$. This set must have size i, and it has some number of internal and boundary edges, say $e_T(A_0) = e$ and $d_T(A_0) = d$. (See Figure 4.3 for an example, in which the parameters have values $(n, k) = (6, 5)$, $(a, b, c) = (8, 11, 4)$, $(\ell, i) = (4, 3)$, and $(d, e) = (2, 1)$.) We compute the number of choices for $A \setminus A_0$ (consisting of leaves of U)

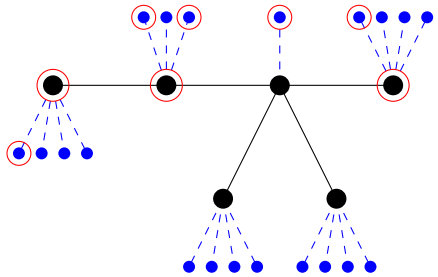

Figure 4.3: An example illustrating the proof of Proposition 4.16. Solid black edges and black vertices are in T, dashed blue edges and smaller blue vertices are added to get U, and circled vertices are in A.

such that $\ell_U(A) = \ell$ and $e_U(A) = c$. First, observe that the number of leaves of U is

$$\ell^* = \sum_{v \in V(T)} (k - \deg_T(v)) = kn - (2n - 2),$$

which is a constant. Of those leaves, the number that are adjacent to a vertex in A_0 is

$$\sum_{v \in A_0} (k - \deg_T(v)) = ki - (d + 2e),$$

again using (4.12). Now $A \setminus A_0$ must consist of ℓ of these leaves, and in order to make $e_U(A) = c$, exactly $c - e$ of the leaves must be adjacent to vertices in A_0. Therefore we have

$$g_U(a, b, c) = \sum_{d,\, e} g_T(i, d, e) \binom{ki - (d + 2e)}{c - e} \binom{\ell^* - (ki - (d + 2e))}{\ell - (c - e)}.$$

Hence $g_U(a, b, c)$ is a linear combination of the coefficients of G_T, as desired. □

Thus if T_1 and T_2 are the trees in Figure 3.1, for example, then for all $k \geq 4$, the trees $T_1^{[k]}$ and $T_2^{[k]}$ also have the same generalized degree polynomial. (In fact, we can take repeated augmentations: for example, $(T_1^{[k]})^{[\ell]}$ and $(T_2^{[k]})^{[\ell]}$ have the same generalized degree polynomial for $\ell \geq k \geq 4$.)

4.3 Degree embeddings

Our main result concerns several classes of information about a tree T that are determined by G_T. This information can be organized using the idea of *degree embeddings*, which we now define.

Definition 4.17. Let D be a tree whose vertices u are labeled with positive integers a_u. A *degree embedding* of D into T is an isomorphism $\varphi\colon D \to T'$, where T' is a subtree of T, such that $a_u = \deg_T(\varphi(u))$ for all $u \in V(D)$. We define the *degree embedding number* of D in T, denoted $N_D(T)$, to be the number of subtrees $T' \subseteq T$ for which a degree embedding exists with image T'.

We adopt a simplified notation when D is a path, say with vertex labels $a_1, \ldots, a_m$ in that order. In this case, $N_D(T)$ is the number of paths $(v_1, \ldots, v_m)$ in T such that $a_i = \deg(v_i)$ for $i = 1, \ldots, m$. (Note that if the sequence $a_1, \ldots, a_m$ is palindromic, then these paths should only be counted up to reversal.) Then we will write $N_{a_1,\ldots,a_m}(T)$ to mean $N_D(T)$. In particular, $N_a(T)$ is the number of vertices of T with degree a.

Example 4.18. For the tree T depicted in Figure 4.4, we have

$$N_1(T) = 9, \qquad N_{2,2}(T) = 1, \qquad N_{1,3,1}(T) = 1, \quad \text{and} \quad N_D(T) = 4,$$

where D is the star on four vertices with center labeled 4 and outer vertices labeled $1, 1, 4$. ◊

Clearly, a tree T can be recovered from all its degree embedding numbers because T (with vertices labeled by degree) is the unique largest tree D with $N_D(T) \neq 0$. However, knowing any "bounded" collection of degree embedding numbers is not sufficient in general to recover T, as we now show.

Proposition 4.19. For every $m \geq 1$, there exist non-isomorphic trees T_1 and T_2 such that $N_D(T_1) = N_D(T_2)$ for all D with $|D| \leq m$.

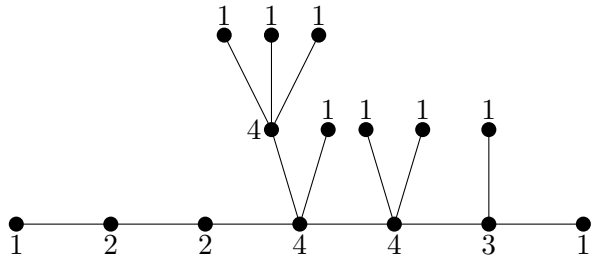

Figure 4.4: The tree T of Example 4.18. Vertices are labeled with their degrees.

Proof. Let T_1 be the tree created by attaching leaves to a path with $2m+4$ vertices so that the path vertices have degrees

$$4, \underbrace{2, \ldots, 2}_{m}, 3, \underbrace{2, \ldots, 2}_{m+1}, 5$$

in that order. Let T_2 be the tree constructed analogously but using degrees

$$4, \underbrace{2, \ldots, 2}_{m+1}, 3, \underbrace{2, \ldots, 2}_{m}, 5.$$

These trees are not isomorphic; however, we can give a bijection between subtrees $S_1 \subseteq T_1$ and $S_2 \subseteq T_2$ of order $|S_1| = |S_2| \leq m$ that preserves vertex degrees (with respect to T_1 and T_2). Suppose S_1 is a subtree of T_1 with $|S_1| \leq m$. If S_1 contains the degree-4 vertex of T_1, then it contains at most $m-1$ degree-2 vertices on the left side of T_1. It follows that there exists a corresponding subtree S_2 on the left side of T_2 with the same vertex degrees. A similar argument holds for the case in which S_1 contains the degree-5 vertex of T_1.

Now the remaining subtrees of T_1 and T_2 are those not containing the degree-4 or degree-5 vertices. But when the degree-4 and degree-5 vertices of T_1 and T_2 are deleted, the resulting graphs are isomorphic (under an isomorphism that preserves the original vertex degrees), so a bijection between subtrees exists as desired. Thus, $N_D(T_1) = N_D(T_2)$ whenever $|D| \leq m$. □

As a preview of the next section, we now present a list of degree embedding numbers that we will be able to recover from G_T. (Note, however, that many degree embedding numbers cannot be determined by G_T. For example, in Figure 3.1, let T_1 be the tree on the left and let T_2 be the tree on the right. Then $\mathrm{G}_{T_1} = \mathrm{G}_{T_2}$, but one can check that $N_{1,2,3}(T_1) = 1 \neq 0 = N_{1,2,3}(T_2)$ and $N_{1,3,2}(T_1) = 2 \neq 3 = N_{1,3,2}(T_2)$.)

Theorem 4.20. The generalized degree polynomial G_T determines the degree embedding numbers $N_D(T)$ when D is any of the following:

(i) A path of order 2.

(ii) A path of order 3 with all vertex labels equal.

(iii) A path of order 3 with vertex labels a, a, b in that order, where $b \geq 2a$.

(iv) A path in which all vertices are labeled 2.

(v) A star in which the outer vertices are all labeled 1.

For completeness, we note that G_T determines $N_D(T)$ if D is a single vertex, since we already know that G_T determines the degree sequence of T. Additionally, G_T determines $N_D(T)$ if D is a path with vertex labels $1, 2, 2, \ldots, 2$ in that order; this is a consequence of a result of Crew [11, Theorem 3] about the *twigs* of a tree, as mentioned in the introduction.

4.4 Degree embedding numbers from the generalized degree polynomial

To prove Theorem 4.20, we first construct a sequence of two-variable rational functions $(\Gamma_T^{(k)})_{k\geq 1}$ which, taken together, are equivalent to G_T. Then, by working with these rational functions, we recover all the desired degree embedding numbers $N_D(T)$.

4.4.1 The rational functions $\Gamma_T^{(k)}$

To define $\Gamma_T^{(k)}$, we will need the following notion. For a subset of edges $S \subseteq E(T)$, let $\overline{S}$ be the subset of $V(T)$ consisting of all vertices incident to at least one edge in S. (That is, considering each edge as a set of size 2, we have $\overline{S} = \bigcup_{e \in S} e$.)

Definition 4.21. Let $\Gamma_T^{(0)}(x, y) = 1$, and for positive integers k, define

$$\Gamma_T^{(k)}(x, y) = \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \prod_{v \in \overline{S}} \frac{xy^{\deg(v)}}{1 + xy^{\deg(v)}}.$$

Example 4.22. Let T be the path on 6 vertices. For brevity, let $q = \frac{xy}{1+xy}$ and $r = \frac{xy^2}{1+xy^2}$. Then $\Gamma_T^{(1)}(x, y) = 2qr + 3r^2$ because two edges of T connect a leaf and a vertex of degree 2, and the other three edges connect two vertices of degree 2. We also have

$$\Gamma_T^{(2)}(x, y) = q^2r^2 + 4qr^3 + r^4 + 2qr^2 + 2r^3.$$

For example, the $2r^3$ term arises because there are two sets $S \subseteq E(T)$ with $|S| = 2$ such that $\overline{S}$ consists of three vertices of degree 2. ◊

We now show how $\Gamma_T^{(k)}$ can be computed from G_T. The key is to take partial derivatives of $\mathrm{G}_T(x, y, z)$ with respect to z.

Proposition 4.23. For each k, $\Gamma_T^{(k)}(x, y)$ is determined by $\mathrm{G}_T(x, y, z)$. Specifically, we have

$$\Gamma_T^{(k)}(x, y) = \frac{y^{2k}}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, y^2) \cdot \mathrm{G}_T(x, y, y^2)^{-1}. \tag{4.15}$$

Proof. Since $\mathrm{G}_T(x, y, z) = \sum_{A \subseteq V(T)} x^{|A|} y^{d(A)} z^{e(A)}$, we have

$$\frac{z^k}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, z) = \sum_{A \subseteq V(T)} \binom{e(A)}{k} x^{|A|} y^{d(A)} z^{e(A)} = \sum_{A \subseteq V(T)} \sum_{\substack{S \subseteq E(A) \\ |S|=k}} x^{|A|} y^{d(A)} z^{e(A)}.$$

Next, we set $z = y^2$, which yields

$$\frac{y^{2k}}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, y^2) = \sum_{A \subseteq V(T)} \sum_{\substack{S \subseteq E(A) \\ |S|=k}} x^{|A|} y^{d(A)+2e(A)} = \sum_{A \subseteq V(T)} \sum_{\substack{S \subseteq E(A) \\ |S|=k}} \prod_{v \in A} xy^{\deg(v)}$$

by the identity (4.12). We now reverse the order of summation. Observe that for $A \subseteq V(T)$ and $S \subseteq E(T)$, we have $S \subseteq E(A)$ if and only if $\overline{S} \subseteq A$. Therefore we can write

$$\begin{aligned}
\frac{y^{2k}}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, y^2) &= \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \sum_{\overline{S} \subseteq A \subseteq V(T)} \prod_{v \in A} xy^{\deg(v)} \\
&= \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \left(\prod_{v \in \overline{S}} xy^{\deg(v)} \prod_{v \in V(T) \setminus \overline{S}} \left(1 + xy^{\deg(v)}\right) \right) \\
&= \prod_{v \in V(T)} (1 + xy^{\deg(v)}) \cdot \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \prod_{v \in \overline{S}} \frac{xy^{\deg(v)}}{1 + xy^{\deg(v)}},
\end{aligned}$$

which shows that

$$\frac{y^{2k}}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, y^2) = \prod_{v \in V(T)} (1 + xy^{\deg(v)}) \cdot \Gamma_T^{(k)}(x, y).$$

By taking $k = 0$, we see that $\mathrm{G}_T(x, y, y^2) = \prod_{v \in V(T)} (1 + xy^{\deg(v)})$, so (4.15) follows. □

Note that, conversely, G_T can be recovered from all of the $\Gamma_T^{(k)}$. Indeed, the Taylor series expansion of $\mathrm{G}_T(x, y, z)$ about $z = y^2$ is

$$\begin{aligned}
\mathrm{G}_T(x, y, z) &= \sum_{k \geq 0} \frac{1}{k!} \cdot \frac{\partial^k \mathrm{G}_T}{\partial z^k}(x, y, y^2) \cdot (z - y^2)^k \\
&= \mathrm{G}_T(x, y, y^2) \sum_{k \geq 0} \frac{1}{y^{2k}} \cdot \Gamma_T^{(k)}(x, y) \cdot (z - y^2)^k,
\end{aligned}$$

so the $\Gamma_T^{(k)}$ determine the product $\mathrm{G}_T(x, y, z) \cdot \mathrm{G}_T(x, y, y^2)^{-1}$. We will see as a consequence of the results of Section 4.4.2 that the $\Gamma_T^{(k)}$ determine the degree sequence of T, so they

determine $\mathrm{G}_T(x, y, y^2)$ and thus $\mathrm{G}_T(x, y, z)$.

Before proceeding, we comment on the relationship between G_T and the $\Gamma_T^{(k)}$. Whereas $\mathbf{X}_T$ determines G_T linearly, it is not true that G_T determines $\Gamma_T^{(k)}$ linearly. In other words, treating $\Gamma_T^{(k)}$ as a formal power series in x and y, its coefficients are generally not linear functions of the coefficients $g_T(a, b, c)$ of G_T (as one can check computationally); rather, the formula (4.15) implies that each coefficient of $\Gamma_T^{(k)}$ can be written as a polynomial in the $g_T(a, b, c)$. We also remark that some of the methods we will use to extract degree embedding numbers from the $\Gamma_T^{(k)}$ are highly nonlinear (particularly those in Section 4.4.3 and Section 4.4.5).

4.4.2 Double-degree numbers

We now prove the first part of Theorem 4.20 by showing that $\Gamma_T^{(1)}$, and thus G_T, determines all the degree embedding numbers $N_{a,b}(T)$. (We call these *double-degree numbers*.) Observe that by a double-counting argument,

$$N_a(T) = \frac{1}{a}\left(\sum_{b \neq a} N_{a,b}(T) + 2N_{a,a}(T)\right).$$

Thus the double-degree numbers determine the degree sequence of T, as promised above.

Theorem 4.24. The rational function $\Gamma_T^{(1)}$ determines all double-degree numbers $N_{a,b}(T)$.

Proof. By symmetry, we may consider only the double-degree numbers $N_{a,b}(T)$ where $a \leq b$. We proceed by induction on the pairs (a, b) with $a \leq b$ in lexicographical order. The base case $(a, b) = (1, 1)$ is trivial because $N_{1,1}(T) = 1$ if $|T| = 2$ and $N_{1,1}(T) = 0$ otherwise. For the induction step, fix positive integers $a_0 \leq b_0$ with $b_0 \geq 2$; assume that $\Gamma_T^{(1)}$ determines all $N_{a,b}(T)$ (with $a \leq b$) where (a, b) precedes (a_0, b_0) lexicographically. We expand $\Gamma_T^{(1)}$ as

a formal power series as follows:

$$\begin{aligned}\Gamma_T^{(1)}(x,y) &= \sum_{\{v,w\}\in E(T)} \frac{xy^{\deg(v)}}{1+xy^{\deg(v)}} \frac{xy^{\deg(w)}}{1+xy^{\deg(w)}} \\ &= \sum_{a\le b} \frac{xy^a}{1+xy^a} \frac{xy^b}{1+xy^b} N_{a,b}(T) \\ &= \sum_{a\le b} \sum_{k,\ell\ge 1} (-1)^{k+\ell} (xy^a)^k (xy^b)^\ell N_{a,b}(T). \end{aligned} \tag{4.16}$$

Consider the coefficient C of the monomial $x^{b_0}y^{a_0b_0-a_0+b_0}$ in this power series. From (4.16), we see that C is a linear combination of terms $N_{a,b}(T)$, where $a \le b$ satisfy

$$k+\ell = b_0 \qquad \text{and} \qquad ka+\ell b = a_0b_0 - a_0 + b_0$$

for some $k, \ell \ge 1$. (The sign $(-1)^{k+\ell} = (-1)^{b_0}$ is constant for all such terms, so there is no cancellation.) These equations are satisfied when $(a,b,k,\ell) = (a_0, b_0, b_0-1, 1)$, so $N_{a_0,b_0}(T)$ appears as a term in the expression for C. We claim that for every other term $N_{a,b}(T)$ that appears, (a,b) precedes (a_0,b_0) lexicographically. Indeed, we have

$$ab_0 = (k+\ell)a \le ka+\ell b = a_0b_0 - a_0 + b_0 < (a_0+1)b_0,$$

so $a < a_0+1$ and thus $a \le a_0$. Furthermore, if $a = a_0$, then we have

$$\ell b = (a_0b_0 - a_0 + b_0) - ka = (a_0b_0 - a_0 + b_0) - (b_0-\ell)a_0 = (\ell-1)a_0 + b_0 \le \ell b_0,$$

so $b \le b_0$. This proves the claim. It then follows from the induction hypothesis that $\Gamma_T^{(1)}$ determines $N_{a_0,b_0}(T)$. □

This establishes Theorem 4.20(i).

Example 4.25. As an illustration of the above proof, we derive a formula for $N_{2,3}(T)$ in terms of coefficients of $\Gamma_T^{(1)}(x,y)$. For $(a_0,b_0) = (2,3)$, we have $x^{b_0}y^{a_0b_0-a_0+b_0} = x^3y^7$, and

so using (4.16) we compute

$$[x^3y^7]\Gamma_T^{(1)} = N_{1,3}(T) + N_{1,5}(T) + N_{2,3}(T).$$

It then remains to compute $N_{1,3}(T)$ and $N_{1,5}(T)$. For $(a_0, b_0) = (1, 5)$, we get

$$[x^5y^9]\Gamma_T^{(1)} = N_{1,2}(T) + N_{1,3}(T) + N_{1,5}(T),$$

so (by coincidence) the $N_{1,3}(T)$ terms will cancel and we need only compute $N_{1,2}(T)$. For $(a_0, b_0) = (1, 2)$, we have just

$$[x^2y^3]\Gamma_T^{(1)} = N_{1,2}(T).$$

Therefore, we get the formula

$$N_{2,3}(T) = [x^3y^7]\Gamma_T^{(1)} - [x^5y^9]\Gamma_T^{(1)} + [x^2y^3]\Gamma_T^{(1)}.$$

$\Diamond$

Using this procedure, any double-degree number $N_{a,b}(T)$ can be expressed as a linear combination of coefficients of $\Gamma_T^{(1)}$ and hence as a polynomial in the coefficients of G_T.

4.4.3 Matchings and paths

We now prove parts (ii) and (iv) of Theorem 4.20. To do so, recall that for a graph G, the *matching-generating polynomial* of G is the univariate polynomial

$$M_G(t) = \sum_S t^{|S|},$$

where S ranges over all matchings (collections of disjoint edges) of G. For $d \geq 2$, denote by $T_{(d)}$ the induced subgraph of T consisting of all degree-d vertices of T.

Theorem 4.26. The polynomial G_T determines $M_{T_{(d)}}(t)$ for all d.

Proof. Let k be an arbitrary positive integer. Consider the product

$$(1+xy^d)^{2k} \cdot \Gamma_T^{(k)}(x,y) = (1+xy^d)^{2k} \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \prod_{v \in \overline{S}} \frac{xy^{\deg(v)}}{1+xy^{\deg(v)}}.$$

For any set of edges $S \subseteq E$ with $|S| = k$, we have $\left|\overline{S}\right| \leq 2k$, with equality if and only S is a matching. Therefore, if we set $x = -y^{-d}$ (so $1 + xy^d = 0$), the only terms that survive are those having $2k$ factors of $1 + xy^d$ in the denominator. These terms occur exactly when S is a matching of size k of $T_{(d)}$. Each such term evaluates to $(-1)^{2k} = 1$, so we recover the number of such matchings for each k. Since k was arbitrary and G_T determines all $\Gamma_T^{(k)}$, it follows that G_T determines $M_{T_{(d)}}(t)$. □

Now, observe that the number of matchings of size 2 in $T_{(d)}$ is given by $\binom{N_{d,d}(T)}{2} - N_{d,d,d}(T)$. Because G_T determines $N_{d,d}(T)$ from Section 4.4.2, it follows that G_T determines $N_{d,d,d}(T)$ as well. This proves Theorem 4.20(ii).

To establish Theorem 4.20(iv), we will show that G_T determines $T_{(2)}$ (up to isomorphism). Note that $T_{(2)}$ is a disjoint union of paths since all of its vertices have degree at most 2. Furthermore, G_T determines the order of $T_{(2)}$ (because this is simply $N_2(T)$) as well as the matching-generating polynomial $M_{T_{(2)}}(t)$. This information suffices to determine $T_{(2)}$ because of the following elementary result.

Proposition 4.27. If G is a disjoint union of paths, then G is determined by $|G|$ and $M_G(t)$.

Proof. For $j \geq 0$, let $f_j(t)$ be the matching-generating polynomial of a path on j vertices. Because the matching-generating polynomial multiplies over disjoint unions, it follows that

$$M_G(t) = f_2(t)^{c_2} \cdots f_n(t)^{c_n} \tag{4.17}$$

where $n = |G|$ and $c_j \geq 0$ is the number of components of G of order j. (Note that $f_1(t) = 1$, so there is no $f_1(t)$ term.) We will show that all products of this form are distinct. The

combinatorial definition of $f_j(t)$ yields the recurrence relation $f_j(t) = f_{j-1}(t) + tf_{j-2}(t)$ for $j \geq 2$, with $f_0(t) = f_1(t) = 1$. One can then show by induction that

$$U_j(x) = (2x)^j f_j\left(-\tfrac{1}{4x^2}\right),$$

where $U_j(x)$ denotes the jth Chebyshev polynomial of the second kind. The roots of $U_j(x)$ are the numbers $x = \cos(\frac{k\pi}{j+1})$ for $k = 1, 2, \ldots, j$, so it follows that $f_j(t)$ has distinct roots

$$t = -\frac{1}{4\cos^2(\frac{k\pi}{j+1})}$$

for $1 \leq k < \frac{j+1}{2}$, a total of $\lfloor \frac{j}{2} \rfloor$ roots. Since f_j has degree $\lfloor \frac{j}{2} \rfloor$, these are all of its roots. In particular, the largest root of $f_j(t)$ is $t_j = -1/(4\cos^2 \frac{\pi}{j+1})$, which increases with j, so t_j is not a root of any of $f_2(t), \ldots, f_{j-1}(t)$. It then follows that all products of the form (4.17) are distinct: if

$$f_2(t)^{c_2} \cdots f_n(t)^{c_n} = f_2(t)^{d_2} \cdots f_n(t)^{d_n}$$

for some $c_j, d_j \geq 0$, then considering the order of the root t_n on both sides forces $c_n = d_n$, and one can continue by induction to show that $c_j = d_j$ for all j.

Therefore, $M_G(t)$ determines the number of components of G of order j for all $j \geq 2$. Then, $|G|$ determines the number of components of order 1, so G is fully determined. □

It now follows that G_T determines $N_{\underbrace{2,\ldots,2}_{k}}(T)$ for all k, because this is the number of paths of order k in $T_{(2)}$. Thus, we have established Theorem 4.20(iv).

4.4.4 Triple-degree numbers

We now consider degree embedding numbers of the form $N_{a,b,c}(T)$, which we call *triple-degree numbers*. We will recover certain linear combinations of triple-degree numbers from G_T, including $N_{a,a,b}(T)$ for $b \geq 2a$, which proves Theorem 4.20(iii).

Theorem 4.28. Let $a \geq 2$ be an integer. Then G_T determines all of the following:

(i) $N_{a,a,1}(T) + N_{a,a,2}(T) + \cdots + N_{a,a,a-1}(T)$,

(ii) $N_{a,a,a-k}(T) - N_{a,a,a+k}(T)$ for all $1 \leq k \leq a-1$, and

(iii) $N_{a,a,b}(T)$ for all $b \geq 2a$.

Proof. Throughout this proof, we write N_D instead of $N_D(T)$ for brevity. We know that G_T determines the rational function

$$\Gamma_T^{(2)}(x,y) = \sum_{\substack{S \subseteq E(T) \\ |S|=2}} \prod_{v \in \overline{S}} \frac{xy^{\deg(v)}}{1+xy^{\deg(v)}}.$$

For some of the terms in this sum, we have $S \subseteq E(T_{(a)})$, that is, all the vertices in $\overline{S}$ have degree a. Each such term equals either $(\frac{xy^a}{1+xy^a})^4$ or $(\frac{xy^a}{1+xy^a})^3$ depending on whether S is a matching or not, respectively. By Theorem 4.26, G_T determines the number of terms of each type. Therefore, it follows that G_T determines the sum

$$\widetilde{\Gamma}(x,y) = \sum_{\substack{S \subseteq E(T) \\ S \nsubseteq E(T_{(a)}) \\ |S|=2}} \prod_{v \in \overline{S}} \frac{xy^{\deg(v)}}{1+xy^{\deg(v)}}$$

obtained by deleting all terms in $\Gamma_T^{(2)}(x,y)$ where $S \subseteq E(T_{(a)})$.

In each term of $\widetilde{\Gamma}(x,y)$, the number of factors of $1+xy^a$ in the denominator is at most 3. Furthermore, equality holds exactly when the two edges of S are disjoint, the endpoints of one edge both have degree a, and the endpoints of the other edge have degrees a and b for some $b \neq a$. Then the corresponding term of $\widetilde{\Gamma}(x,y)$ is $(\frac{xy^a}{1+xy^a})^3 \frac{xy^b}{1+xy^b}$, and the number of such terms for a given b is $N_{a,a}N_{a,b} - N_{a,a,b}$. Therefore if we take $x = -y^{-a}$ in the product

$(1+xy^a)^3 \cdot \widetilde{\Gamma}(x,y)$, then the result is

$$\sum_{b\neq a}(xy^a)^3\frac{xy^b}{1+xy^b}\cdot(N_{a,a}N_{a,b}-N_{a,a,b})\Big|_{x=-y^{-a}}$$
$$=\sum_{b\neq a}\frac{y^{b-a}}{1-y^{b-a}}\cdot(N_{a,a}N_{a,b}-N_{a,a,b})\,.$$

Because G_T determines all $N_{a,a}$ and $N_{a,b}$ by Theorem 4.24, it follows that G_T also determines

$$\sum_{b\neq a}\frac{y^{b-a}}{1-y^{b-a}}N_{a,a,b}.$$

Now we write

$$\begin{aligned}
&\sum_{b\neq a}\frac{y^{b-a}}{1-y^{b-a}}N_{a,a,b}\\
&=-\sum_{b=1}^{a-1}\frac{1}{1-y^{a-b}}N_{a,a,b}+\sum_{b=a+1}^{\infty}\frac{y^{b-a}}{1-y^{b-a}}N_{a,a,b}\\
&=-\sum_{k=1}^{a-1}\frac{1}{1-y^{k}}N_{a,a,a-k}+\sum_{k=1}^{\infty}\frac{y^{k}}{1-y^{k}}N_{a,a,a+k}\\
&=-\sum_{k=1}^{a-1}\left(1+\frac{y^k}{1-y^k}\right)N_{a,a,a-k}+\sum_{k=1}^{\infty}\frac{y^{k}}{1-y^{k}}N_{a,a,a+k}\\
&=-\sum_{k=1}^{a-1}N_{a,a,a-k}-\sum_{k=1}^{a-1}\frac{y^k}{1-y^k}(N_{a,a,a-k}-N_{a,a,a+k})+\sum_{k=a}^{\infty}\frac{y^{k}}{1-y^{k}}N_{a,a,a+k}.
\end{aligned}$$

The rational functions 1 and $\frac{y^k}{1-y^k}$ for $k\geq 1$ are linearly independent over $\mathbb{Q}$ because they have roots of different orders at $y=0$. Thus, G_T determines the coefficients of these rational functions in the above expression, which are precisely the desired linear combinations. □

4.4.5 Leaf adjacency

Finally, we prove Theorem 4.20(v). Assume that $|T|\geq 3$, so that no two leaves of T are adjacent. For a vertex $v\in V(T)$, let $\ell(v)$ be the number of leaves of T adjacent to v, and let $L_{d,\ell}(T)$ be the number of vertices v with $\deg(v)=d$ and $\ell(v)=\ell$. The numbers $L_{d,\ell}(T)$

form the *leaf adjacency sequence* of T.

Theorem 4.29. Fix integers $d \geq 2$ and $\ell \geq 0$. Then G_T determines $L_{d,\ell}(T)$.

Proof. Fix a positive integer k. Instead of $\Gamma_T^{(k)}(x, y)$, it will be more convenient to make the substitution

$$\Gamma_T^{(k)}(x^{-1}, y^{-1}) = \sum_{\substack{S \subseteq E(T) \\ |S|=k}} \prod_{v \in \overline{S}} \frac{1}{1 + xy^{\deg(v)}}.$$

For each term in the above sum, the number of factors of $1 + xy$ in the denominator is at most k, with equality if and only if $S \subseteq L(T)$, where $L(T)$ is the set of all edges of T adjacent to a leaf. Therefore, we have

$$(1 + xy)^k \cdot \Gamma_T^{(k)}(x^{-1}, y^{-1})\Big|_{x=-y^{-1}} = \sum_{\substack{S \subseteq L(T) \\ |S|=k}} \prod_{v \in \widehat{S}} \frac{1}{1 - y^{\deg(v)-1}},$$

where $\widehat{S}$ denotes the set of all internal vertices in $\overline{S}$. (Note that $\widehat{S}$ has size at most k.) Now fix an integer $a \geq 2$, and consider the product

$$(1 - y^{a-1})^k \sum_{\substack{S \subseteq L(T) \\ |S|=k}} \prod_{v \in \widehat{S}} \frac{1}{1 - y^{\deg(v)-1}}.$$

Let $\omega = e^{2\pi i/(a-1)}$, which is a root of $1 - y^{d-1}$ if and only if $(a-1) \mid (d-1)$, in which case it is a simple root. If we set $y = \omega$ in the above product, the only terms that will survive are those for which $\widehat{S}$ has size k (thus, all the edges of S are disjoint), and every vertex $v \in \widehat{S}$ satisfies $(a-1) \mid (\deg(v) - 1)$. Let $V^{(a)}$ be the set of all such vertices:

$$V^{(a)} = \{v \in V(T) : \deg(v) > 1 \text{ and } (a-1) \mid (\deg(v) - 1)\}.$$

Then each surviving term corresponds to a choice of k vertices from $V^{(a)}$ and a leaf edge

incident to each chosen vertex. Therefore, setting $y = \omega$ yields

$$\begin{aligned}(1-y^{a-1})^k \sum_{\substack{A\subseteq V^{(a)}\\|A|=k}} \prod_{v\in A}\left(\ell(v)\frac{1}{1-y^{\deg(v)-1}}\right)\Bigg|_{y=\omega} &= \sum_{\substack{A\subseteq V^{(a)}\\|A|=k}} \prod_{v\in A}\left(\ell(v)\lim_{y\to\omega}\frac{1-y^{a-1}}{1-y^{\deg(v)-1}}\right)\\ &= \sum_{\substack{A\subseteq V^{(a)}\\|A|=k}} \prod_{v\in A}\left(\ell(v)\frac{a-1}{\deg(v)-1}\right)\\ &= (a-1)^k \sum_{\substack{A\subseteq V^{(a)}\\|A|=k}} \prod_{v\in A}\frac{\ell(v)}{\deg(v)-1}.\end{aligned}$$

Ignoring the factor of $(a-1)^k$, this is the kth elementary symmetric sum of the multiset

$$M^{(a)} = \left\{\frac{\ell(v)}{\deg(v)-1} : v \in V^{(a)}\right\}.$$

Since k was arbitrary, it follows that G_T determines the multiset $M^{(a)}$ for each $a \geq 2$.

We now show how the leaf adjacency sequence of T can be recovered from the multisets $M^{(a)}$. For each $d \geq 2$, let V_d be the set of vertices of T with degree d, so $V^{(a)} = \bigcup_{i\geq 1} V_{1+i(a-1)}$. Let M_d be the multiset

$$M_d = \left\{\frac{\ell(v)}{d-1} : v \in V_d\right\},$$

so $M^{(a)} = \bigsqcup_{i\geq 1} M_{1+i(a-1)}$. In particular, if Δ is the largest degree of a vertex of T, then $M^{(\Delta)} = M_\Delta$, so the multiset M_Δ is determined by G_T. Then, for $2 \leq d < \Delta$, we have

$$M_d = M^{(d)} \setminus \bigsqcup_{i\geq 2} M_{1+i(d-1)},$$

so by an inductive process, we can recover all the M_d from the $M^{(a)}$ and thus from G_T. Now M_d is clearly equivalent to the multiset of numbers $\ell(v)$ for $v \in V_d$, so the M_d together determine the leaf adjacency sequence of T, as desired. □

This result is equivalent to Theorem 4.20(v). Indeed, let $S_{d,k}$ be the star of order $k+1$

with center vertex labeled d and outer vertices all labeled 1. Then we have

$$N_{S_{d,k}}(T) = \sum_{\ell \geq k} \binom{\ell}{k} L_{d,\ell}(T),$$

so by triangularity, the numbers $(N_{S_{d,k}}(T))$ and $(L_{d,\ell}(T))$ determine each other. (Note that $N_{S_{d,2}}(T) = N_{1,d,1}(T)$, so this recovers another sequence of triple-degree numbers of T.)

4.5 Higher-order generalized degree polynomials

In the final section of this chapter, we discuss a further generalization of the GDP: the *generalized degree polynomial of order m* for positive integers m.

To define this polynomial, we first extend the notion of *boundary edge* from Section 4.2 as follows. Let F be a forest and let $A_1, \ldots, A_m$ be disjoint subsets of $V(F)$. A *boundary edge* of $(A_1, \ldots, A_m)$ is an edge of F that is a boundary edge of at least one A_i, that is, the set of boundary edges of $(A_1, \ldots, A_m)$ is simply

$$D(A_1, \ldots, A_m) = D(A_1) \cup \cdots \cup D(A_m).$$

We write $d(A_1, \ldots, A_m) = |D(A_1, \ldots, A_m)|$. Thus, $d(A_1, \ldots, A_m) \leq d(A_1) + \cdots + d(A_m)$.

Definition 4.30. The *generalized degree polynomial of order m* of a forest F is

$$\mathrm{G}_F^{(m)}(x_1, \ldots, x_m, y, z_1, \ldots, z_m) = \sum_{A_1 \sqcup \cdots \sqcup A_m \subseteq V(F)} x_1^{|A_1|} \cdots x_m^{|A_m|} \cdot y^{d(A_1,\ldots,A_m)} \cdot z_1^{e(A_1)} \cdots z_m^{e(A_m)},$$

where the sum is over tuples $(A_1, \ldots, A_m)$ of disjoint subsets of $V(F)$.

Note that $\mathrm{G}_F^{(1)}$ is equivalent to the usual generalized degree polynomial G_F. As in Section 4.2, we will show that $\mathrm{G}_F^{(m)}$ is a linear function of $\mathbf{X}_F$. To do this, we first consider a variant of $\mathbf{X}_F$ itself, the *bad chromatic symmetric function* $\mathbf{B}_F$, and show that in fact it is equivalent to $\mathbf{X}_F$. We will then show how $\mathbf{B}_F$ naturally gives rise to $\mathrm{G}_F^{(m)}$.

Definition 4.31. For a forest F, the *bad chromatic symmetric function* of F is

$$\mathbf{B}_F(x_1, x_2, \dots; z_1, z_2, \dots) = \sum_{\kappa} \prod_i x_i^{|\kappa^{-1}(i)|} z_i^{b_\kappa(i)},$$

where $\kappa\colon V(F) \to \mathbb{Z}_{>0}$ is any vertex coloring (not necessarily proper) of F, and $b_\kappa(i)$ is the number of edges $\{v, w\} \in E(F)$ such that $\kappa(v) = \kappa(w) = i$.

(This can be seen as a further generalization of Stanley's symmetric function generalization of the bad coloring polynomial by setting $z_i = 1 + t$; see [32, Section 3.1].)

If we let A_i denote the set of vertices colored i, then we can alternatively write

$$\mathbf{B}_F(x_1, x_2, \dots; z_1, z_2, \dots) = \sum_{A_1 \sqcup A_2 \sqcup \cdots = V(F)} \prod_i x_i^{|A_i|} z_i^{e(A_i)},$$

where the sum ranges over all ways to express $V(F)$ as a disjoint union of (possibly empty) sets A_i. We denote by $\mathbf{B}_F^{(m)}$ the polynomial in variables $x_1, \dots, x_m, z_1, \dots, z_m$ obtained by setting $x_i = z_i = 0$ in $\mathbf{B}_F$ for $i > m$. Observe that $\mathbf{B}_F$ can be recovered as an inverse limit over the polynomials $\mathbf{B}_F^{(m)}$ because any particular monomial of $\mathbf{B}_F$ appears in $\mathbf{B}_F^{(m)}$ for all sufficiently large m.

As a formal power series, $\mathbf{B}_F$ is a *MacMahon symmetric function*, as defined in [28]: that is, it is invariant under the diagonal action of the symmetric group of the positive integers simultaneously permuting the x-variables and the z-variables. In addition, setting all $z_i = 0$ in $\mathbf{B}_F$ recovers the chromatic symmetric function $\mathbf{X}_F$. Note also that if A_i is a nonempty subset of $V(F)$ inducing a connected subgraph of F, then $e(A_i) = |A_i| - 1$. Thus, for any partition $\lambda \vdash |F|$, the coefficient of $\prod_i x_i^{\lambda_i} z_i^{\lambda_i - 1}$ in $\mathbf{B}_F$ is (up to a factor corresponding to permuting equal parts of λ) exactly $b_\lambda(F)$, the number of connected partitions of F of type λ. It follows that one can also recover the power sum expansion of $\mathbf{X}_F$ from $\mathbf{B}_F$.

While it might seem that $\mathbf{B}_F$ contains strictly more information than $\mathbf{X}_F$, the following proposition shows that $\mathbf{X}_F$ determines $\mathbf{B}_F^{(m)}$ for each m, so $\mathbf{B}_F$ and $\mathbf{X}_F$ are in fact equivalent.

Proposition 4.32. For each m, there exists a $\mathbb{Q}$-algebra map $\beta^{(m)}$ such that $\beta^{(m)}(\mathbf{X}_F) = \mathbf{B}_F^{(m)}$ for all forests F.

Proof. Recall the functions $\varphi_{t,u}$ from Lemma 4.15. We consider the m-fold convolution

$$\beta^{(m)} = \varphi_{x_1,z_1} * \cdots * \varphi_{x_m,z_m}.$$

For a forest F, iterating Proposition 2.6 yields

$$\Delta^{m-1}(\mathbf{X}_F) = \sum_{A_1 \sqcup \cdots \sqcup A_m = V(F)} \mathbf{X}_{F|A_1} \otimes \cdots \otimes \mathbf{X}_{F|A_m},$$

where $A_1, \ldots, A_m$ partition the vertices of F. Therefore, we have

$$\begin{aligned} \beta^{(m)}(\mathbf{X}_F) &= \sum_{A_1 \sqcup \cdots \sqcup A_m = V(F)} \varphi_{x_1,z_1}(\mathbf{X}_{F|A_1}) \cdots \varphi_{x_m,z_m}(\mathbf{X}_{F|A_m}) \\ &= \sum_{A_1 \sqcup \cdots \sqcup A_m = V(F)} x_1^{|A_1|} \cdots x_m^{|A_m|} \cdot z_1^{e(A_1)} \cdots z_m^{e(A_m)}, \end{aligned}$$

which is precisely $\mathbf{B}_F^{(m)}$. □

We now show that $\mathrm{G}_F^{(m)}$ is essentially equivalent to $\mathbf{B}_F^{(m+1)}$. If $A_1 \sqcup \cdots \sqcup A_{m+1} = V(F)$, then the boundary edges of $(A_1, \ldots, A_m)$ are precisely the edges whose endpoints lie in different subsets A_i and A_j for some $1 \leq i \neq j \leq m+1$. Hence

$$d(A_1, \ldots, A_m) + \sum_{i=1}^{m+1} e(A_i) = \sum_{i=1}^{m+1} |A_i| - c(F)$$

since both sides count the total number of edges in F. Therefore

$$\begin{aligned}
&\mathbf{B}_F^{(m+1)}(x_1y,\dots,x_my,y;\, y^{-1}z_1,\dots,y^{-1}z_m,y^{-1}) \\
&= \sum_{A_1\sqcup\cdots\sqcup A_m\sqcup A_{m+1}=V(F)} x_1^{|A_1|}\cdots x_m^{|A_m|}\cdot y^{\sum_{i=1}^{m+1}(|A_i|-e(A_i))}\cdot z_1^{e(A_1)}\cdots z_m^{e(A_m)} \\
&= \sum_{A_1\sqcup\cdots\sqcup A_m\subseteq V(F)} x_1^{|A_1|}\cdots x_m^{|A_m|}\cdot y^{d(A_1,\dots,A_m)+c(F)}\cdot z_1^{e(A_1)}\cdots z_m^{e(A_m)} \\
&= y^{c(F)}\mathrm{G}_F^{(m)}(x_1,\dots,x_m,y,z_1,\dots,z_m).
\end{aligned}$$

It follows that there is a $\mathbb{Q}$-algebra map $\gamma^{(m)}$ that sends $\mathbf{X}_F$ to $y^{c(F)}\mathrm{G}_F^{(m)}$. Specifically, the above computations show that we can take

$$\gamma^{(m)} = \varphi_{x_1y,y^{-1}z_1} * \cdots * \varphi_{x_my,y^{-1}z_m} * \varphi_{y,y^{-1}}.$$

(When $m = 1$, this reduces to the map from Proposition 4.14.) We note that $\gamma^{(m)}$ is multiplicative and satisfies

$$\gamma^{(m)}(p_n) = x_1^n y(y-z_1)^{n-1} + \cdots + x_m^n y(y-z_m)^{n-1} + y(y-1)^{n-1}.$$

4.5.1 Future work

A natural direction for future work is to investigate what information about a tree T is determined by $\mathrm{G}_T^{(m)}$ for each positive integer m. In particular, we pose the following question.

Question 4.33. Does there exist a positive integer m such that $\mathrm{G}_T^{(m)}$ distinguishes trees?

As we already know, it is not true that $\mathrm{G}_T^{(1)}$ (i.e., G_T) distinguishes trees; the smallest pair of non-isomorphic trees with the same generalized degree polynomial is depicted in Figure 3.1. However, we have not been able to find a pair of non-isomorphic trees T_1 and T_2 with $\mathrm{G}_{T_1}^{(2)} = \mathrm{G}_{T_2}^{(2)}$. A computer search indicates that no such trees exist with $|T_1| = |T_2| \leq 23$.

A positive answer to Question 4.33 would immediately imply a positive answer to Ques-

tion 1.1. However, the converse does not necessarily hold, that is, it is possible that the chromatic symmetric function distinguishes trees but no single $\mathrm{G}_T^{(m)}$ does. (This is because our method to recover $\mathbf{X}_T$ from $\mathbf{B}_T$ and thus from the $\mathrm{G}_T^{(m)}$ requires $m = |T| - 1$, which is unbounded.) If this is indeed the case, then a construction to produce pairs (or larger sets) of non-isomorphic trees with the same $\mathrm{G}_T^{(m)}$ could be valuable in studying Question 1.1.

Chapter 5

GENERALIZED DEGREE POLYNOMIALS OF ROOTED AND POLARIZED GRAPHS

In this chapter, we consider variations on the generalized degree polynomial for graphs with distinguished vertices, specifically *rooted graphs* (with one distinguished vertex called the *root*) and *polarized graphs* (with two distinguished vertices called the *left pole* and *right pole*). These variations on the generalized degree polynomial enumerate the same triple of statistics over vertex subsets as the ordinary GDP, but do so while also tracking whether the distinguished vertices lie in that subset.

Unlike ordinary trees, rooted trees and polarized trees are distinguished by their respective GDPs, as we show by adapting work of Loehr and Warrington [23]. However, studying the GDPs of rooted and polarized trees will yield results about the GDPs of ordinary trees: specifically, we present a recurrence relation satisfied by the ordinary GDP and prove some divisibility results. In the next chapter, we will use the machinery of rooted and polarized GDPs to give several constructions for families of ordinary trees that have the same GDP.

5.1 Rooted graphs

Recall that a *rooted graph* is a graph P together with a distinguished vertex $r(P)$ called its root.[1] Two rooted graphs P and Q are considered isomorphic, written $P \cong Q$, if there is an edge-preserving bijection $V(P) \to V(Q)$ that sends $r(P)$ to $r(Q)$.

Given a rooted graph P, we define two polynomials $\mathrm{G}_P^1(x, y, z)$ and $\mathrm{G}_P^r(x, y, z)$ by

$$\mathrm{G}_P^1 = \sum_{r(P) \notin S} x^{|S|} y^{d(S)} z^{e(S)} \qquad \text{and} \qquad \mathrm{G}_P^r = \sum_{r(P) \in S} x^{|S|} y^{d(S)} z^{e(S)},$$

[1] We will use the letters P and Q to represent rooted graphs when possible.

where in the first sum, S ranges over all subsets of $V(P)$ not containing $r(P)$, and in the second sum, S ranges over all subsets of $V(P)$ containing $r(P)$.

Definition 5.1. The *generalized degree polynomial* (GDP) of a rooted graph P is the vector

$$\mathbf{G}_P = \begin{bmatrix} \mathrm{G}_P^1 \\ \mathrm{G}_P^r \end{bmatrix}.$$

If P is a rooted graph, then by default, "the GDP" of P should be understood to mean $\mathbf{G}_P$ as just defined; we will also refer to $\mathbf{G}_P$ as "the rooted GDP" of P for clarity. However, on occasion we may want to refer to "the ordinary GDP" of P, denoted G_P, which is the GDP of P considered as an ordinary graph (forgetting about the root). This satisfies

$$\mathrm{G}_P = \mathrm{G}_P^1 + \mathrm{G}_P^r = \begin{bmatrix} 1 & 1 \end{bmatrix} \mathbf{G}_P.$$

Example 5.2. If P is the path Pa_3 with its root at one end, then we have

$$\mathbf{G}_P = \begin{bmatrix} 1 + xy + xy^2 + x^2yz \\ xy + x^2y^2 + x^2yz + x^3z^2 \end{bmatrix}.$$

$\Diamond$

Before continuing, we remark that it is not entirely necessary to include both G_P^1 and G_P^r in our definition of $\mathbf{G}_P$ because, in fact, G_P^1 and G_P^r determine each other. To see why, first, let n and m be the number of vertices and edges of P, respectively. Note that the unique term of G_P^1 with highest x-degree is $x^{n-1}y^{\deg r(P)}z^{m-\deg r(P)}$, obtained by choosing the vertex subset S to contain all vertices of P except for $r(P)$. Thus, n and m can be recovered from G_P^1. Similarly, the unique term of G_P^r with highest x-degree is x^nz^m, obtained by choosing S to be all of $V(P)$, so n and m can also be recovered from G_P^r.

Next, suppose that $S \subseteq V(P)$ and let $S' = V(P) \setminus S$ be the complement of S. Then

$$\begin{aligned} |S'| &= n - |S|, \\ d(S') &= d(S), \\ e(S') &= m - d(S) - e(S). \end{aligned}$$

Hence, the linear involution on $\mathbb{Q}[x, y, z]$ that sends $x^a y^b z^c$ to $x^{n-a} y^b z^{m-b-c}$ (if all the exponents are nonnegative, otherwise zero) sends the monomial $x^{|S|} y^{d(S)} z^{e(S)}$ to $x^{|S'|} y^{d(S')} z^{e(S')}$ for all $S \subseteq V(P)$. Since $r(P) \notin S$ if and only if $r(P) \in S'$, it follows that this involution swaps G_P^1 and G_P^r. (We can apply this involution given G_P^1 or G_P^r because we can first recover the values of n and m.) Thus, knowing either G_P^1 or G_P^r is enough to determine the rooted GDP of P. Despite this, our choice of definition for $\mathbf{G}_P$ is more symmetric and will prove useful later on in our analysis.

5.1.1 Distinguishing rooted trees

Unlike for ordinary (unrooted) trees, rooted trees are completely distinguished by their GDPs, as we now prove. To do so, we rely on a rooted version of the chromatic symmetric function [23], defined by Loehr and Warrington, which also distinguishes rooted trees. In fact, they proved a much stronger result, namely, that there is a univariate specialization of the rooted CSF that distinguishes rooted trees.

Definition 5.3 ([23, Section 3.4])**.** Let P be a rooted graph and $i \in \{0, 1, 2, 3\}$. We define a polynomial

$$f_{P,i}(q) = \sum_{\kappa \in C_i(P)} q^{|\kappa^{-1}(\{0,1\})|},$$

where $C_i(P)$ is the set of all proper colorings $\kappa \colon V(P) \to \{0, 1, 2, 3\}$ such that $\kappa(r(P)) = i$.

Note that $f_{P,0}(q) = f_{P,1}(q)$ because, given a proper coloring $\kappa \in C_0(P)$, replacing all 0s with 1s and 1s with 0s gives a corresponding proper coloring $\kappa' \in C_1(P)$. By a similar

argument, $f_{P,2}(q) = f_{P,3}(q)$. Furthermore, $f_{P,0}(q)$ and $f_{P,2}(q)$ determine each other because their coefficient sequences are reverses of each other [23, Proposition 16(a)]. Specifically, if P has n vertices, then there are nonnegative integers $c_1, \ldots, c_n$ (depending on P) such that

$$f_{P,0}(q) = c_n q^n + c_{n-1} q^{n-1} + \cdots + c_2 q^2 + c_1 q$$

and

$$f_{P,2}(q) = c_1 q^{n-1} + c_2 q^{n-2} + \cdots + c_{n-1} q + c_n.$$

(This can be proved by considering the effect of switching 0 with 2 and 1 with 3 in proper colorings of P.) Thus, the invariants $f_{P,i}(q)$ are all equivalent. Loehr and Warrington showed that they distinguish rooted trees [23, Theorem 17].

To bring the rooted GDP into this picture, we show that $f_{P,0}(q)$ (and thus each $f_{P,i}(q)$) is a specialization of $\mathbf{G}_P$ when P is a rooted tree.

Theorem 5.4. For any rooted tree P, $f_{P,0}(q) = \mathrm{G}^r_P(q, 2, 1)$ and $f_{P,2}(q) = \mathrm{G}^1_P(q, 2, 1)$.

Proof. Given a proper coloring $\kappa \in C_0(P)$, let $S = \kappa^{-1}(\{0, 1\})$ be the set of all vertices colored 0 or 1. Then S is a subset of $V(P)$ containing $r(P)$, and the contribution of κ to the sum defining $f_{P,0}(q)$ is $q^{|S|}$. Therefore

$$f_{P,0}(q) = \sum_{r(P) \in S} q^{|S|} \, \#\{\kappa \in C_0(P) \colon \kappa^{-1}(\{0, 1\}) = S\}.$$

Now fix some $S \subseteq V(P)$ with $r(P) \in S$, and let $S' = V(P) \setminus S$. Consider a proper coloring $\kappa \in C_0(P)$ such that $\kappa^{-1}(\{0, 1\}) = S$. Under κ, each component of $P|S$ must be colored with colors 0 and 1, while each component of $P|S'$ must be colored with colors 2 and 3. These components are all subtrees of P and can be colored independently. Each has exactly two valid colorings, except for the component of $P|S$ containing $r(P)$, which has only one

(since $r(P)$ must get color 0). Thus the number of possible κ is

$$2^{c(P|S)+c(P|S')-1} = 2^{|S|-e(S)+|S'|-e(S')-1} = 2^{d(S)}$$

(since $|S| + |S'| = |P|$). We conclude that

$$f_{P,0}(q) = \sum_{r(P)\in S} q^{|S|} 2^{d(S)} = \mathrm{G}^r_P(q, 2, 1),$$

as claimed. To prove $f_{P,2}(q) = \mathrm{G}^1_P(q, 2, 1)$, we use the same argument: for each S, the number of proper colorings $\kappa \in C_2(P)$ such that $\kappa^{-1}(\{0, 1\}) = S$ is still $2^{d(S)}$, but now S must not contain $r(P)$. □

When this is combined with Loehr and Warrington's result, we immediately obtain:

Corollary 5.5. Non-isomorphic rooted trees have different rooted GDPs.

In fact, we can prove a stronger result: non-isomorphic rooted trees must have *linearly independent* rooted GDPs. To do so, we will need to perform a somewhat deeper analysis of the polynomials $f_{P,i}(q)$ and the structure of rooted trees.

We say that a rooted tree is a *flower* if its root has degree 1. (In particular, the rooted tree with one vertex is not a flower.) Then, every rooted tree P can be written uniquely as a union of flowers with the same root as P and that only intersect at their roots. If $P = P_1 \cup \cdots \cup P_k$ is such a decomposition, then we can naturally identify the set of colorings $C_i(P)$ with the Cartesian product $C_i(P_1) \times \cdots \times C_i(P_k)$, for each $i \in \{0, 1, 2, 3\}$. It follows that

$$f_{P,2}(q) = f_{P_1,2}(q) \cdots f_{P_k,2}(q) \qquad \text{and} \qquad \frac{f_{P,0}(q)}{q} = \frac{f_{P_1,0}(q)}{q} \cdots \frac{f_{P_k,0}(q)}{q},$$

where in the latter equation we divide each polynomial by q to avoid overcounting the root. (The rooted GDP factors similarly: $\mathrm{G}^1_P = \mathrm{G}^1_{P_1} \cdots \mathrm{G}^1_{P_k}$ and $\frac{\mathrm{G}^r_P}{x} = \frac{\mathrm{G}^r_{P_1}}{x} \cdots \frac{\mathrm{G}^r_{P_k}}{x}$.)

Lemma 5.6. If P is a flower, then $f_{P,0}(q)/q$ and $f_{P,2}(q)$ are irreducible over $\mathbb{Q}[q]$.

Proof. By Theorem 5.4, we have

$$f_{P,0}(q) = \mathrm{G}^r_P(q, 2, 1) = \sum_{r(P)\in S} q^{|S|} 2^{d(S)}.$$

Since P is a flower, $d(\{r(P)\}) = \deg r(P) = 1$. Furthermore, $d(S) > 0$ for all nonempty proper subsets S of $V(P)$. This means that

$$f_{P,0}(q) = q^n + c_{n-1}q^{n-1} + \cdots + c_2 q^2 + 2q$$

where $n = |P|$ and $c_{n-1}, \ldots, c_2$ are even integers. Thus, Eisenstein's criterion with respect to the prime 2 shows that $f_{P,0}(q)/q$ is irreducible. Since $f_{P,2}(q)$ is obtained by reversing the coefficient sequence of $f_{P,0}(q)$, its irreducibility follows immediately. □

We can now prove the desired strengthening of Corollary 5.5.

Theorem 5.7. If P and Q are non-isomorphic rooted trees, then $\mathbf{G}_P$ and $\mathbf{G}_Q$ are linearly independent.

Proof. Suppose that $\mathbf{G}_P$ and $\mathbf{G}_Q$ are linearly dependent. Then $\mathrm{G}^1_P \mathrm{G}^r_Q = \mathrm{G}^1_Q \mathrm{G}^r_P$. Making the substitution $(x, y, z) = (q, 2, 1)$ and dividing by q, we get

$$f_{P,2}\frac{f_{Q,0}}{q} = f_{Q,2}\frac{f_{P,0}}{q}.$$

Let $P = P_1 \cup \cdots \cup P_k$ and $Q = Q_1 \cup \cdots \cup Q_\ell$ be the flower decompositions of P and Q. Then we have

$$f_{P_1,2} \cdots f_{P_k,2} \cdot \frac{f_{Q_1,0}}{q} \cdots \frac{f_{Q_\ell,0}}{q} = f_{Q_1,2} \cdots f_{Q_\ell,2} \cdot \frac{f_{P_1,0}}{q} \cdots \frac{f_{P_k,0}}{q}.$$

By Lemma 5.6, all these terms are irreducible in $\mathbb{Q}[q]$, so we must be able to pair up terms between the left and right sides such that the terms in each pair are rational multiples of each other.

However, observe that if C is a flower, then $f_{C,2}(q)$ has leading coefficient 2 and constant coefficient 1, while $f_{C,0}(q)/q$ has leading coefficient 1 and constant coefficient 2. Thus, for this pairing to exist, we must have $k = \ell$, and $f_{P_i,2}(q) = f_{Q_i,2}(q)$ and $f_{P_i,0}(q)/q = f_{Q_i,0}(q)/q$ for all i (relabeling the flowers P_i as necessary). By [23, Theorem 17], this means that $P_i \cong Q_i$ for all i, and hence $P \cong Q$. □

5.2 Polarized graphs

Next, we consider a kind of graph with two distinguished vertices, which is lesser-known but nonetheless will prove very useful in analyzing the generalized degree polynomial.

Definition 5.8. A *polarized graph* is a graph A together with two distinguished vertices, called the *left pole* and *right pole*, and denoted by $L(A)$ and $R(A)$, respectively.[2]

We consider polarized graphs A and B to be isomorphic, written $A \cong B$, if there is an edge-preserving bijection $V(A) \to V(B)$ that sends $L(A)$ to $L(B)$ and $R(A)$ to $R(B)$.

Note that the left and right poles of a polarized graph are allowed to be the same vertex; such polarized graphs are essentially equivalent to rooted graphs. However, if they are distinct, then it matters which pole is which: switching them generally produces a non-isomorphic polarized graph.

It will be useful to have a standard notation for two small polarized graphs:

(i) We denote by $\underset{\text{LR}}{\bullet}$ the polarized graph with one vertex (and no edges).

(ii) We denote by $\underset{\text{L}}{\bullet}\ \underset{\text{R}}{\bullet}$ the polarized graph with two vertices, no edges, and distinct poles.

Let us now define the generalized degree polynomial of a polarized graph.

Definition 5.9. Let A be a polarized graph. We define four polynomials G_A^1, G_A^R, G_A^L, and

[2]We will use the first few capital letters A, B, C, etc., to represent polarized graphs when possible.

G_A^{LR} in x, y, z as follows:

$$\mathrm{G}_A^1 = \sum_{\substack{L(A)\notin S\\ R(A)\notin S}} x^{|S|}y^{d(S)}z^{e(S)}, \qquad \mathrm{G}_A^R = \sum_{\substack{L(A)\notin S\\ R(A)\in S}} x^{|S|}y^{d(S)}z^{e(S)},$$

$$\mathrm{G}_A^L = \sum_{\substack{L(A)\in S\\ R(A)\notin S}} x^{|S|}y^{d(S)}z^{e(S)}, \qquad \mathrm{G}_A^{LR} = \sum_{\substack{L(A)\in S\\ R(A)\in S}} x^{|S|}y^{d(S)}z^{e(S)}.$$

(In each sum, S ranges over subsets of $V(A)$ that either do or do not contain $L(A)$, and either do or do not contain $R(A)$.) Then the *generalized degree polynomial* of A is the matrix

$$\mathbf{G}_A = \begin{bmatrix} \mathrm{G}_A^1 & \mathrm{G}_A^R \\ \mathrm{G}_A^L & \mathrm{G}_A^{LR} \end{bmatrix}.$$

If A is a polarized graph, then "the GDP of A" should be understood to mean the matrix $\mathbf{G}_A$ just defined; we will also use the term "polarized GDP of A" for clarity. Note that if A is considered as an ordinary graph, then its ordinary GDP G_A is

$$\mathrm{G}_A = \mathrm{G}_A^1 + \mathrm{G}_A^L + \mathrm{G}_A^R + \mathrm{G}_A^{LR} = \begin{bmatrix} 1 & 1 \end{bmatrix} \mathbf{G}_A \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Example 5.10. For the polarized graph A depicted in Figure 5.1, we have

$$\mathbf{G}_A = \begin{bmatrix} 1+xy & xy+x^2y^2 \\ xy^2+x^2yz & x^2yz+x^3z^2 \end{bmatrix}.$$

Note that when drawing a polarized graph, we label the left and right poles with L and R, respectively.

$\Diamond$

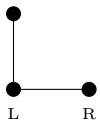

Figure 5.1: The polarized graph A of Example 5.10.

Example 5.11. The polarized GDPs of $\underset{\text{LR}}{\bullet}$ and $\underset{\text{L}}{\bullet}\,\underset{\text{R}}{\bullet}$ are respectively

$$\mathbf{G}_{\bullet} = \begin{bmatrix} 1 & 0 \\ 0 & x \end{bmatrix} \qquad \text{and} \qquad \mathbf{G}_{\bullet\bullet} = \begin{bmatrix} 1 & x \\ x & x^2 \end{bmatrix}.$$

Note that we write $\mathbf{G}_{\bullet}$ and $\mathbf{G}_{\bullet\bullet}$ instead of the more cumbersome $\mathbf{G}_{\underset{\text{LR}}{\bullet}}$ and $\mathbf{G}_{\underset{\text{L}}{\bullet}\,\underset{\text{R}}{\bullet}}$. ◊

If A is a polarized graph and one of its components C contains neither $L(A)$ nor $R(A)$, then we have $\mathbf{G}_A = \mathrm{G}_C \mathbf{G}_{A \setminus C}$, where $A \setminus C$ is the union of the other components. That is, the presence of a "non-polarized" component only multiplies the polarized GDP by a scalar. (The analogous fact for rooted graphs is also true.)

Remark 5.12. It may seem strange to define the polarized GDP as a matrix, rather than (say) a 4-dimensional vector. One reason for doing so is that this definition generalizes readily to graphs with any number of distinguished vertices, so that the GDP of a graph with an n-tuple of distinguished vertices is a tensor of order n over the field $\mathbb{Q}(x, y, z)$. In addition, we will see several ways in which one can analyze polarized graphs by applying linear algebra to their GDPs, and this is made much more elegant by using matrices. (In particular, *concatenation* of polarized graphs will be essential to stating and proving our main results, and its effect on GDPs is naturally described using matrix multiplication.)

5.2.1 Properties of polarized graphs

We now define some basic concepts related to polarized graphs, beginning with some properties they may possess.

Definition 5.13. A polarized graph A is called

(i) *diagonal* if $L(A) = R(A)$;

(ii) *symmetric* if there is an automorphism of A (as an ordinary graph) that swaps $L(A)$ and $R(A)$;

(iii) *singular* if $L(A)$ and $R(A)$ are in different connected components of A.

As we are about to show, these properties of polarized graphs are so named because they imply the corresponding property (for matrices) of the polarized GDP.

Proposition 5.14. Let A be a polarized graph.

(i) If A is diagonal, then $\mathbf{G}_A$ is diagonal.

(ii) If A is symmetric, then $\mathbf{G}_A$ is symmetric.

(iii) If A is singular, then $\mathbf{G}_A$ is singular.

Proof. If A is diagonal, then clearly there are no subsets $S \subseteq V(A)$ with $L(A) \in S$ and $R(A) \notin S$, or $L(A) \notin S$ and $R(A) \in S$. Thus $\mathrm{G}_A^L = \mathrm{G}_A^R = 0$, so $\mathbf{G}_A$ is diagonal, proving (i).

If A is symmetric, let φ be an automorphism of A swapping $L(A)$ and $R(A)$. Then φ induces a bijection between subsets of $V(A)$ that contain $L(A)$ but not $R(A)$, and those that contain $R(A)$ but not $L(A)$. Therefore $\mathrm{G}_A^L = \mathrm{G}_A^R$, so $\mathbf{G}_A$ is symmetric, proving (ii).

Finally, suppose that A is singular. Let P be the component of A containing $L(A)$, considered as a rooted tree with root $L(A)$. Let Q be the union of all the other components, considered as a rooted tree with root $R(A)$. Then we have the equalities

$$\mathrm{G}_A^1 = \mathrm{G}_P^1\mathrm{G}_Q^1, \qquad \mathrm{G}_A^L = \mathrm{G}_P^r\mathrm{G}_Q^1, \qquad \mathrm{G}_A^R = \mathrm{G}_P^1\mathrm{G}_Q^r, \qquad \text{and} \qquad \mathrm{G}_A^{LR} = \mathrm{G}_P^r\mathrm{G}_Q^r,$$

which implies $\mathrm{G}_A^1\mathrm{G}_A^{LR} = \mathrm{G}_A^L\mathrm{G}_A^R$. Thus $\mathbf{G}_A$ is singular, proving (iii). (In the next section, we will give a name to this decomposition of singular polarized graphs.) $\square$

Observe that the converse of (1) also holds, since if $L(A) \neq R(A)$, then G_A^L and G_A^R are both nonzero. Later in this chapter, we will be able to prove the converses of (2) and (3) as well (assuming that A is a forest, which is the case we are most concerned with).

5.3 Operations on polarized and rooted graphs

We now consider some operations that can be performed on polarized and rooted graphs, together with their effects on the polarized and rooted GDP.

Definition 5.15. Let A be a polarized graph. The *opposite* of A, denoted A^*, is the polarized graph derived from A by switching its poles: $L(A^*) = R(A)$ and $R(A^*) = L(A)$.

Thus, a polarized graph A is symmetric if and only if $A \cong A^*$ as polarized graphs.

From the definition of A^*, it is clear that $\mathrm{G}_{A^*}^1 = \mathrm{G}_A^1$ and $\mathrm{G}_{A^*}^{LR} = \mathrm{G}_A^{LR}$, while $\mathrm{G}_{A^*}^L = \mathrm{G}_A^R$ and $\mathrm{G}_{A^*}^R = \mathrm{G}_A^L$. Therefore

$$\mathbf{G}_{A^*} = \mathbf{G}_A^T, \tag{5.1}$$

that is, $\mathbf{G}_{A^*}$ and $\mathbf{G}_A$ are transposes.

Next, we consider two ways to turn polarized graphs into rooted graphs.

Definition 5.16. Let A be a polarized graph. The *left-rooting* of A, denoted A_L, is the rooted graph obtained from A by choosing its left pole $L(A)$ as the root. Analogously, the *right-rooting* of A, denoted A_R, is obtained from A by choosing $R(A)$ as the root.

Observe that

$$\mathrm{G}_{A_L}^1 = \mathrm{G}_A^1 + \mathrm{G}_A^R \qquad \text{and} \qquad \mathrm{G}_{A_L}^r = \mathrm{G}_A^L + \mathrm{G}_A^{LR},$$

which means that the rooted GDP of A_L satisfies

$$\mathbf{G}_{A_L} = \mathbf{G}_A \begin{bmatrix} 1 \\ \\ 1 \end{bmatrix}. \tag{5.2}$$

Similarly, the rooted GDP of the right-rooting of A satisfies

$$\mathbf{G}_{A_R} = \begin{bmatrix} 1 & 1 \end{bmatrix} \mathbf{G}_A \tag{5.3}$$

(treating $\mathbf{G}_{A_R}$ as a row vector for the purpose of this formula).

Conversely, here are some ways to turn rooted graphs into polarized graphs.

Definition 5.17. If P is a rooted graph, its *polarization* $\hat{P}$ is the polarized graph obtained by P by taking $L(\hat{P}) = R(\hat{P}) = r(P)$.

Thus, the polarization of a rooted graph is always diagonal, and its polarized GDP satisfies

$$\mathbf{G}_{\hat{P}} = \begin{bmatrix} \mathrm{G}^1_{\hat{P}} & 0 \\ 0 & \mathrm{G}^{LR}_{\hat{P}} \end{bmatrix} = \begin{bmatrix} \mathrm{G}^1_P & 0 \\ 0 & \mathrm{G}^r_P \end{bmatrix}. \tag{5.4}$$

In other words, $\mathbf{G}_{\hat{P}}$ is the diagonal matrix whose diagonal entries are the entries of $\mathbf{G}_P$.

Definition 5.18. If P and Q are any rooted graphs, we can consider their disjoint union $P \sqcup Q$ as a polarized graph with left pole $r(P)$ and right pole $r(Q)$. We call this polarized graph the *tensor product* of P and Q, and denote it by $P \otimes Q$.[3]

Observe that the tensor product of rooted graphs is always a singular polarized graph. Furthermore, its polarized GDP satisfies

$$\mathbf{G}_{P\otimes Q} = \begin{bmatrix} \mathrm{G}^1_{P\otimes Q} & \mathrm{G}^R_{P\otimes Q} \\ \mathrm{G}^L_{P\otimes Q} & \mathrm{G}^{LR}_{P\otimes Q} \end{bmatrix} = \begin{bmatrix} \mathrm{G}^1_P \mathrm{G}^1_Q & \mathrm{G}^1_P \mathrm{G}^r_Q \\ \mathrm{G}^r_P \mathrm{G}^1_Q & \mathrm{G}^r_P \mathrm{G}^r_Q \end{bmatrix} = \begin{bmatrix} \mathrm{G}^1_P \\ \mathrm{G}^r_P \end{bmatrix} \begin{bmatrix} \mathrm{G}^1_Q & \mathrm{G}^r_Q \end{bmatrix} = \mathbf{G}_P \, \mathbf{G}_Q^T, \tag{5.5}$$

which justifies the name "tensor product" for this operation.

Conversely, for any singular polarized graph A, we can always find P and Q such that $A = P \otimes Q$, say, by taking P to be the component of A containing $L(A)$ and taking Q to

[3]Note that the operation $\otimes$ on rooted graphs is not symmetric; rather, $P \otimes Q$ and $Q \otimes P$ are opposites of each other as defined above.

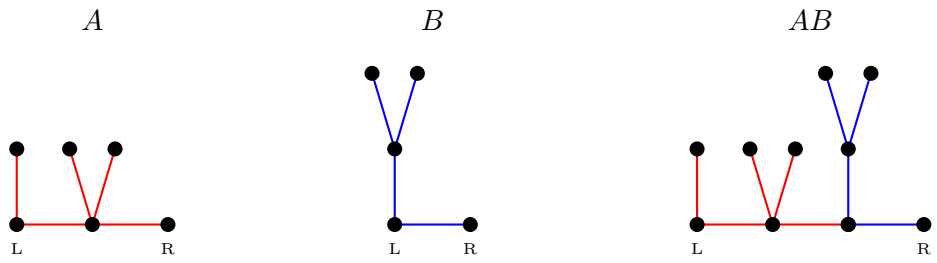

Figure 5.2: Two polarized graphs and their concatenation.

be the union of all the other components. If A only has two components, then the pair (P, Q) is unique. Instead, if A has components containing neither $L(A)$ nor $R(A)$, then these components can be chosen to be part of either P or Q. (However, in applications, our graphs will generally not have any such components.)

5.3.1 Concatenation and cross-join

Now, we define the two most important operations on polarized graphs. We begin with *concatenation*, which combines two polarized graphs by joining them end to end.

Definition 5.19. Let A and B be polarized graphs. The *concatenation* of A and B, denoted AB, is the polarized graph obtained from the disjoint union $A \sqcup B$ by identifying $R(A)$ and $L(B)$, with poles $L(AB) = L(A)$ and $R(AB) = R(B)$. (Here we view A and B as embedded inside AB.)

We say that polarized graphs A and B *commute* if $AB \cong BA$.

Example 5.20. Figure 5.2 depicts the concatenation of two polarized graphs. ◊

Example 5.21. If P and Q are rooted graphs, then their tensor product $P \otimes Q$ can also be written as the concatenation $\hat{P}\left(\underset{\text{L}}{\bullet}\,\underset{\text{R}}{\bullet}\right)\hat{Q}$. ◊

Note that the polarized graph $\underset{\text{LR}}{\bullet}$ is the (two-sided) identity element for concatenation.

As promised, here is the main reason for defining the polarized GDP as a matrix: concatenation of polarized graphs corresponds (very nearly) to multiplication of their GDPs.

Theorem 5.22. For any two polarized graphs A and B, we have

$$\mathbf{G}_{AB} = \mathbf{G}_A M \mathbf{G}_B \tag{5.6}$$

where $M = \left[\begin{smallmatrix} 1 & 0 \\ 0 & x^{-1} \end{smallmatrix}\right]$.

Proof. The claimed equality is

$$\begin{bmatrix} \mathrm{G}^1_{AB} & \mathrm{G}^R_{AB} \\ \mathrm{G}^L_{AB} & \mathrm{G}^{LR}_{AB} \end{bmatrix} = \begin{bmatrix} \mathrm{G}^1_{A} & \mathrm{G}^R_{A} \\ \mathrm{G}^L_{A} & \mathrm{G}^{LR}_{A} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & x^{-1} \end{bmatrix} \begin{bmatrix} \mathrm{G}^1_{B} & \mathrm{G}^R_{B} \\ \mathrm{G}^L_{B} & \mathrm{G}^{LR}_{B} \end{bmatrix}. \tag{5.7}$$

For each entry on the left-hand side, a bijective argument shows that it is equal to the corresponding entry of the product on the right-hand side. For example, taking the top-left entry of both sides, we see that we want to prove

$$\mathrm{G}^1_{AB} = \mathrm{G}^1_A \mathrm{G}^1_B + x^{-1} \mathrm{G}^R_A \mathrm{G}^L_B. \tag{5.8}$$

Suppose that $S \subseteq V(AB)$ is a subset containing neither $L(AB)$ nor $R(AB)$ (which are the same as $L(A)$ and $R(B)$, respectively). There are two cases to consider:

(i) If S does not contain the vertex $R(A) = L(B)$ either, then we can uniquely write S as a union $S = S_1 \cup S_2$, where $S_1 \subseteq V(A)$ and $S_2 \subseteq V(B)$ such that $L(A), R(A) \notin S_1$ and $L(B), R(B) \notin S_2$. Then the contribution to the sum defining G^1_{AB} over such S is exactly $\mathrm{G}^1_A \mathrm{G}^1_B$, giving the first term on the right-hand side.

(ii) If instead S contains $R(A) = L(B)$, then we have $S = S_1 \cup S_2$, where $S_1 \subseteq V(A)$ and $S_2 \subseteq V(B)$ such that $L(A) \notin S_1$, $R(A) \in S_1$, $L(B) \in S_2$, and $R(B) \notin S_2$. The contribution to G^1_{AB} over such S is $x^{-1}\mathrm{G}^R_A \mathrm{G}^L_B$, where the factor of x^{-1} is necessary to avoid double-counting the shared vertex $R(A) = L(B)$. (That is, $|S| = |S_1| + |S_2| - 1$ in this case, whereas in the previous case we had just $|S| = |S_1| + |S_2|$.) This gives the second term on the right-hand side.

By combining these cases, we obtain (5.8), as desired. The equations for the other entries of (5.7) follow similarly. □

We can also define concatenation when one or both of the graphs are rooted. Specifically:

(i) If A is rooted and B is polarized, then AB is a rooted graph with root $R(B)$, obtained from $A \sqcup B$ by identifying $r(A)$ with $L(B)$.

(ii) If A is polarized and B is rooted, then AB is a rooted graph with root $L(A)$, obtained from $A \sqcup B$ by identifying $R(A)$ with $r(B)$.

(iii) If A and B are both rooted, then AB is an ordinary graph, obtained from $A \sqcup B$ by identifying $r(A)$ with $r(B)$.

In all of these cases, equation (5.6) still holds (interpreting the GDPs so that the matrix multiplication makes sense[4]). Here is one way to see this. Consider, for example, case (i), in which A is rooted and B is polarized. Let $\hat{A}$ be the polarization of A. Then (5.6) gives $\mathbf{G}_{\hat{A}B} = \mathbf{G}_{\hat{A}} M \mathbf{G}_B$. Multiplying on the left by $[1\ 1]$, we obtain $\mathbf{G}_{AB} = \mathbf{G}_A M \mathbf{G}_B$ by (5.3), since AB is the right-rooting of $\hat{A}B$ and A is the right-rooting of $\hat{A}$. A similar argument applies for cases (ii) and (iii).

Remark 5.23. Observe that the matrix M of Theorem 5.22 satisfies $M = \mathbf{G}_{\bullet}^{-1}$. Therefore, taking $A = \underset{\text{LR}}{\bullet}$ in (5.6) gives $\mathbf{G}_{\bullet B} = \mathbf{G}_B$, as expected because $\underset{\text{LR}}{\bullet} B$ is isomorphic to B. (We get a similar conclusion by taking $B = \underset{\text{LR}}{\bullet}$.) In addition, multiplying (5.6) on the right by M gives

$$\mathbf{G}_{AB} M = \mathbf{G}_A M \mathbf{G}_B M$$

for all polarized graphs A and B. This implies that concatenation corresponds exactly to matrix multiplication under the map $A \mapsto \mathbf{G}_A M$.

[4]Namely, we should assume that in case (i) above, $\mathbf{G}_{AB}$ and $\mathbf{G}_A$ are row vectors; in case (ii), $\mathbf{G}_{AB}$ and $\mathbf{G}_B$ are column vectors; and in case (iii), $\mathbf{G}_{AB}$ is a scalar, $\mathbf{G}_A$ is a row vector, and $\mathbf{G}_B$ is a column vector.

Remark 5.24. Theorem 5.22 can also be derived by considering a (vertex-)weighted version of the polarized GDP. Let us sketch the details. A *weighted polarized graph* is a pair (A, w), where A is a polarized graph and w is a function $V(A) \to \mathbb{Z}$. (When w is a constant function, say $w(v) = c$ for all v, we write (A, c) instead of (A, w).) We extend w to subsets of $V(A)$ additively and call $w(S)$ the *weight* of S. Then the polarized GDP of (A, w), denoted $\mathbf{G}_{A,w}$, is defined by replacing $|S|$ with $w(S)$ throughout Definition 5.9; for example, $\mathbf{G}_{A,w}$ has top-left entry

$$\mathrm{G}^1_{A,w} = \sum_{\substack{L(A)\notin S \\ R(A)\notin S}} x^{w(S)} y^{d(S)} z^{e(S)}.$$

(The entries of $\mathbf{G}_{A,w}$ are *Laurent* polynomials in x, y, z, since w can take negative values.) Note that $\mathbf{G}_{A,1} = \mathbf{G}_A$, and in general, $\mathbf{G}_{A,c}$ is obtained from $\mathbf{G}_A$ by replacing x with x^c.

We define the *concatenation* of two weighted polarized graphs just like in the unweighted case, except that we must specify the weight function. The most natural choice is to sum the weights on the identified vertex v: that is, the concatenation of (A, w_A) and (B, w_B) is (AB, w_{AB}), where $w_{AB}(v) = w_A(R(A)) + w_B(L(B))$, and otherwise w_{AB} restricts to w_A and w_B. Then one can show, using a similar bijective argument as above, that concatenation corresponds exactly to matrix multiplication of the GDPs:

$$\mathbf{G}_{AB,w_{AB}} = \mathbf{G}_{A,w_A} \mathbf{G}_{B,w_B}.$$

To recover Theorem 5.22, we cannot just take $w_A, w_B \equiv 1$ in this formula, because then the identified vertex $R(A) = L(B)$ would have weight 2. However, we can correct for this by inserting the polarized graph $\underset{\text{LR}}{\bullet}$ whose single vertex has weight -1. That is, we take the three weighted polarized graphs $(A, 1)$, $(\underset{\text{LR}}{\bullet}, -1)$, and $(B, 1)$, whose concatenation in that order is $(AB, 1)$, and therefore

$$\mathbf{G}_{AB} = \mathbf{G}_A \mathbf{G}_{\bullet,-1} \mathbf{G}_B.$$

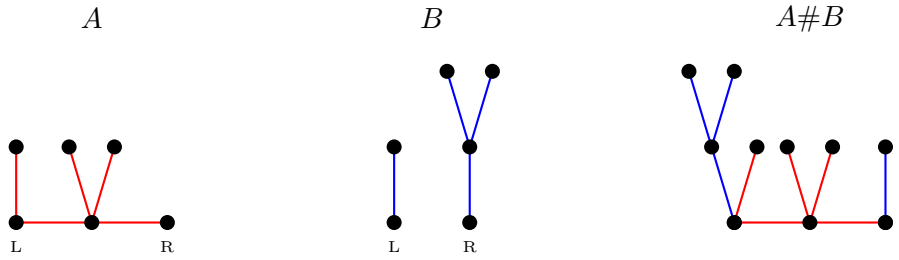

Figure 5.3: The cross-join of two polarized graphs.

Theorem 5.22 now follows because $\mathbf{G}_{\bullet,-1} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & x^{-1} \end{smallmatrix}\right]$.

Finally, we define a related operation to concatenation called *cross-join*, which turns two polarized graphs into an ordinary graph.

Definition 5.25. Let A and B be polarized graphs, at least one of which is singular. The *cross-join* of A and B, denoted $A\#B$, is the ordinary graph obtained from the disjoint union $A \sqcup B$ by identifying $R(A)$ and $L(B)$, and identifying $L(A)$ and $R(B)$.[5]

Note that unlike concatenation, cross-join is a symmetric operation, in the sense that $A\#B$ and $B\#A$ are isomorphic. We have the relation $A\#(BC) \cong (AB)\#C$ when both sides are defined.

The cross-join can also be written using concatenation. Suppose, for example, that B is singular, and write it as a tensor product of rooted graphs $B = P \otimes Q$. Then $A\#B = QAP$.

Example 5.26. Figure 5.3 depicts polarized graphs A and B, with B singular, and their cross-join $A\#B$. ◊

Using our formula for the polarized GDP of a concatenation, we can also derive an expression for the GDP of the cross-join of polarized graphs.

[5]It is possible to define $A\#B$ in this way even when both A and B are nonsingular. However, the resulting graph could have multiple edges (for example, if the poles of A were adjacent and the poles of B were adjacent) or self-loops (for example, if A were diagonal and the poles of B were adjacent) and so we would need to adjust our notion of the GDP to handle non-simple graphs. However, even if we did so in the most natural way, the below formula for the GDP of $A\#B$ would still be incorrect if both A and B were diagonal. Instead, the present definition will suffice for our purposes and lets us avoid such complications.

Theorem 5.27. Let A and B be polarized graphs, at least one of which is singular. Then

$$\mathrm{G}_{A\#B} = \mathrm{tr}(\mathbf{G}_A M \mathbf{G}_B M), \tag{5.9}$$

where $M = \left[\begin{smallmatrix} 1 & 0 \\ 0 & x^{-1} \end{smallmatrix}\right]$ as in Theorem 5.22.

Proof. Suppose that B is singular, and write $B = P \otimes Q$ for rooted graphs P and Q. Then we have

$$\begin{aligned} \mathrm{G}_{A\#B} = \mathrm{G}_{QAP} &= \mathbf{G}_Q^T M \mathbf{G}_A M \mathbf{G}_P \\ &= \mathrm{tr}(\mathbf{G}_Q^T M \mathbf{G}_A M \mathbf{G}_P) \\ &= \mathrm{tr}(\mathbf{G}_A M \mathbf{G}_P \mathbf{G}_Q^T M) \\ &= \mathrm{tr}(\mathbf{G}_A M \mathbf{G}_B M), \end{aligned}$$

where we used the fact that $\mathrm{tr}(XY) = \mathrm{tr}(YX)$ for matrices X and Y, as well as (5.5). The case in which A is singular follows by symmetry. □

In view of Theorem 5.27, we define a bilinear form on 2×2 matrices over $\mathbb{Q}(x, y, z)$ by

$$\langle f, g \rangle = \mathrm{tr}(fMgM).$$

Then (5.9) says that $\mathrm{G}_{A\#B} = \langle \mathbf{G}_A, \mathbf{G}_B \rangle$. (Note that the form $\langle \cdot, \cdot \rangle$ is symmetric.)

Before continuing, we note an easy lemma that will be useful for a proof in Chapter 6.

Lemma 5.28. The bilinear form $\langle \cdot, \cdot \rangle$ is nondegenerate. That is, if f is a 2×2 matrix such that $\langle f, g \rangle = 0$ for all g, then we must have $f = 0$.

Proof. Let $f = \left[\begin{smallmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{smallmatrix}\right]$. For $i, j \in \{1, 2\}$, let E_{ij} be the 2×2 matrix with 1 in position (i, j) and zeros everywhere else. Taking $g = M^{-1} E_{ij} M^{-1}$, one can check that

$$\langle f, g \rangle = \mathrm{tr}\,(f E_{ij}) = f_{ji}.$$

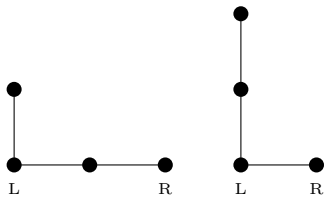

Figure 5.4: A pair of polarized graphs that are doppelgängers.

Thus all the entries of f are zero. □

5.4 Distinguishing polarized trees

We will now show that the polarized GDP distinguishes polarized trees. To do so, we will need the analogous fact about rooted trees (Corollary 5.5), although it will not quite suffice on its own.

Let A be a polarized tree. Its polarized GDP $\mathbf{G}_A$ determines $\mathbf{G}_{A_L}$ and $\mathbf{G}_{A_R}$ via (5.2) and (5.3). Then, it follows from Corollary 5.5 that $\mathbf{G}_A$ determines both A_L and A_R, up to isomorphism of rooted graphs. Unfortunately, this information is not enough to recover A, because there exist non-isomorphic polarized trees whose left-rootings are isomorphic and whose right-rootings are isomorphic. Let us give these pairs of trees a name:

Definition 5.29. Two polarized graphs A and B are called *doppelgängers* if $A \ncong B$, but $A_L \cong B_L$ and $A_R \cong B_R$ as rooted graphs.

Example 5.30. The two polarized graphs shown in Figure 5.4 are doppelgängers. ◊

Of course, doppelgängers are isomorphic as ordinary graphs. Hence, we can construct any pair of doppelgängers by taking an ordinary graph A and then choosing vertices u_1, v_1, u_2, and v_2 of A satisfying the following conditions:

(i) There exists an automorphism of A that sends u_1 to u_2.

(ii) There exists an automorphism of A that sends v_1 to v_2.

(iii) There is *no* automorphism of A that simultaneously sends u_1 to u_2 and v_1 to v_2.

Then the corresponding doppelgängers have left pole u_i and right pole v_i, for $i \in \{1, 2\}$.

Given A_L and A_R, in order to distinguish A from its possible doppelgängers, it turns out that we only need one more piece of information: the distance between the poles. We first show that $\mathbf{G}_A$ determines this distance.

Lemma 5.31. If A is a polarized tree, then $\mathbf{G}_A$ determines the length (number of edges) of the unique path in A between $L(A)$ and $R(A)$.

Proof. Let ℓ be the length of this path. Then G_A^{LR} contains a term of the form $x^{\ell+1}y^k z^\ell$ for some k, obtained by choosing the vertex subset S to be exactly the vertices along the path. Furthermore, G_A^{LR} does not contain a term of the form $x^{m+1}y^k z^m$ for any k and $m < \ell$, since such a term would correspond to a connected subgraph of A that contains both $L(A)$ and $R(A)$ but has fewer than $\ell + 1$ vertices. Hence G_A^{LR} determines ℓ, as desired. □

It remains to show that this information about a polarized tree A — namely, the isomorphism types of A_L and A_R, as well as the distance between $L(A)$ and $R(A)$ — is enough to recover A. This is a purely graph-theoretical task, not involving the GDP; our proof relies on certain facts about automorphisms of ordinary trees.

Let T be a tree. Given vertices v_1 and v_2 of T, we write $v_1 \sim v_2$ if there is an automorphism of T sending v_1 to v_2; note that $\sim$ is an equivalence relation on $V(T)$. Also, let $d(v_1, v_2)$ be the distance between v_1 and v_2, that is, the number of edges in the path in T between v_1 and v_2.

Lemma 5.32. Let T be a tree and suppose v_1 and v_2 are vertices with $v_1 \sim v_2$. Then there is an automorphism of T that swaps v_1 and v_2.

Proof. Recall that every tree has a *center*, either a single vertex or a pair of adjacent vertices, that must be fixed by any automorphism. (The center appears as the middle vertex or vertices of any longest path in the tree.) If the center of T is a pair of adjacent

vertices, we may subdivide their common edge so that the newly introduced vertex becomes the center; this does not change the automorphism group of T. Then we may consider T to be rooted at its center vertex u; let T_v denote the subtree of a vertex v.

If $v_1 \sim v_2$, then v_1 and v_2 must have the same distance from u; we induct on this distance. If v_1 and v_2 are adjacent to u, then T_{v_1} and T_{v_2} must be isomorphic as rooted trees, so there is an automorphism of T that swaps these subtrees and fixes all other vertices. This swaps v_1 and v_2, as desired.

Otherwise, let w_1 and w_2 be the parents of v_1 and v_2, respectively. Then $w_1 \sim w_2$, so by induction, there is an automorphism φ that swaps w_1 and w_2, and hence also swaps their subtrees. Let $\varphi(v_1) = v_1' \in T_{w_2}$ and $\varphi(v_2) = v_2' \in T_{w_1}$. Then the subtrees of v_1, v_1', v_2, and v_2' are all isomorphic, so there is an automorphism of T_{w_1} that swaps v_1 and v_2', as well as an automorphism of T_{w_2} that swaps v_1' and v_2. Composing these with φ, we get an automorphism of T that swaps v_1 and v_2, as desired. □

Lemma 5.33. Let T be a tree, and let u_1, u_2, v_1, and v_2 be vertices of T such that $u_1 \sim u_2$ and $v_1 \sim v_2$. If additionally $d(u_1, v_1) = d(u_2, v_2)$, then there is an automorphism of T sending $u_1 \to u_2$ and $v_1 \to v_2$.

Proof. Let φ be an automorphism of T with $\varphi(v_1) = v_2$, and let $\varphi(u_1) = u_1'$. If $u_1' = u_2$ then we are done, so assume $u_1' \neq u_2$. We construct an automorphism ψ with $\psi(u_1') = u_2$ that fixes v_2; then $\psi \circ \varphi$ will send $u_1 \to u_2$ and $v_1 \to v_2$.

Let $w_0, w_1, \ldots, w_m$ be the vertices on the unique path from $u_1' = w_0$ to $u_2 = w_m$; then T decomposes as this path, together with rooted trees $T_0, T_1, \ldots, T_m$ attached to each of the vertices on the path. Notice that

$$d(u_1', v_2) = d(u_1, v_1) = d(u_2, v_2).$$

Therefore, m must be even, and v_2 must lie in the rooted tree $T_{m/2}$.

By Lemma 5.32, there is an automorphism σ that swaps u_1' and u_2; this must also

fix $w_{m/2}$ and restrict to an automorphism τ of the rooted tree $T_{m/2}$. Extend τ to an automorphism of T by fixing all vertices outside $T_{m/2}$; then $\tau^{-1} \circ \sigma$ swaps u_1' and u_2 and it fixes v_2, as desired. □

Given our characterization of doppelgängers above, Lemma 5.31 and Lemma 5.33 immediately imply the desired theorem:

Theorem 5.34. The polarized GDP distinguishes polarized trees.

Finally, like with rooted trees, we can also extend this result to linear independence. We need a fact about the relationship between the GDPs of doppelgängers:

Proposition 5.35. If A and B are doppelgängers, then $\mathbf{G}_A - \mathbf{G}_B$ is a multiple of $\left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$.

Proof. We have $\mathbf{G}_{A_L} = \mathbf{G}_{B_L}$, so $\mathbf{G}_A\left[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right] = \mathbf{G}_B\left[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right]$. Similarly, $\mathbf{G}_{A_R} = \mathbf{G}_{B_R}$, so $\left[\begin{smallmatrix} 1 & 1 \end{smallmatrix}\right]\mathbf{G}_A = \left[\begin{smallmatrix} 1 & 1 \end{smallmatrix}\right]\mathbf{G}_B$. Hence $\mathbf{G}_A - \mathbf{G}_B$ is a 2×2 matrix whose right kernel contains $\left[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right]$ and whose left kernel contains $\left[\begin{smallmatrix} 1 & 1 \end{smallmatrix}\right]$. This implies the claim. □

Theorem 5.36. If A and B are non-isomorphic polarized trees, then $\mathbf{G}_A$ and $\mathbf{G}_B$ are linearly independent.

Proof. Suppose that $\mathbf{G}_A$ and $\mathbf{G}_B$ are linearly dependent. Then the rooted GDPs of A_L and B_L are also linearly dependent (due to (5.2)), so by Theorem 5.7, A_L and B_L are isomorphic. Similarly, A_R and B_R are isomorphic. Therefore, A and B are doppelgängers (or isomorphic, in which case we are done), so $\mathbf{G}_A - \mathbf{G}_B$ is a multiple of $\left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$ by Proposition 5.35. But we are assuming that $\mathbf{G}_A$ and $\mathbf{G}_B$ are multiples of each other: thus, if $\mathbf{G}_A - \mathbf{G}_B \neq 0$, then $\mathbf{G}_A$ and $\mathbf{G}_B$ must themselves be nonzero multiples of $\left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$. However, this is absurd because the entries of $\mathbf{G}_A$ and $\mathbf{G}_B$ are polynomials with nonnegative coefficients. Hence $\mathbf{G}_A - \mathbf{G}_B$ must be zero, so $A \cong B$ by Theorem 5.34. □

To conclude this section, we use the results we have built up to prove the converse of Proposition 5.14 for polarized forests.

Theorem 5.37. Let A be a polarized forest.

(i) If $\mathbf{G}_A$ is diagonal, then A is diagonal.

(ii) If $\mathbf{G}_A$ is symmetric, then A is symmetric.

(iii) If $\mathbf{G}_A$ is singular, then A is singular.

Proof. For (i), if $\mathbf{G}_A$ is a diagonal matrix, then $\mathrm{G}_A^L = \mathrm{G}_A^R = 0$, which is only possible if $L(A) = R(A)$, so A is diagonal. (This argument works even if A is not a forest.)

For (ii), suppose that $\mathbf{G}_A$ is a symmetric matrix. Then $\mathbf{G}_{A_L} = \mathbf{G}_{A_R}$ (interpreting both as column vectors) since $\mathbf{G}_{A_L} = \mathbf{G}_A[\begin{smallmatrix}1\\1\end{smallmatrix}]$ and $\mathbf{G}_{A_R} = \mathbf{G}_A^T[\begin{smallmatrix}1\\1\end{smallmatrix}]$. We distinguish two cases based on whether A is singular. First, suppose that $L(A)$ and $R(A)$ are in the same component C of A, and let E be the union of the other components. Then $\mathbf{G}_{A_L} = \mathrm{G}_E\mathbf{G}_{C_L}$ and $\mathbf{G}_{A_R} = \mathrm{G}_E\mathbf{G}_{C_R}$ (where C has the same poles as A), so $\mathbf{G}_{C_L} = \mathbf{G}_{C_R}$. Now C_L and C_R are rooted trees, so by Corollary 5.5, they must be isomorphic. Thus, there is an automorphism of C (as an ordinary graph) that swaps $L(A)$ and $R(A)$ (see Lemma 5.32), and this can be extended to all of A by fixing all of E. Therefore, A is symmetric.

Second, suppose that A is singular. Without loss of generality, assume that A only has two components, one containing $L(A)$ and the other containing $R(A)$. (The presence of any other components does not affect whether $\mathbf{G}_A$ and A are symmetric.) Write A as a tensor product $A = P \otimes Q$, where P and Q are rooted trees. Then we have

$$\mathbf{G}_{A_L} = \mathbf{G}_P\mathrm{G}_Q \qquad \text{and} \qquad \mathbf{G}_{A_R} = \mathrm{G}_P\mathbf{G}_Q.$$

Hence, $\mathbf{G}_P$ and $\mathbf{G}_Q$ are linearly dependent, so P and Q must be isomorphic by Theorem 5.7. Then $A = P \otimes Q$ is symmetric because an isomorphism between P and Q can be extended to an automorphism of A (as an ordinary graph) that swaps $L(A)$ and $R(A)$.

For (iii), we prove the contrapositive. Suppose that A is nonsingular. Without loss of generality, we may assume that A is a polarized tree. (If A has more than one component,

then one component must contain both poles, and the other components only serve to multiply $\mathbf{G}_A$ by a nonzero scalar.) Let $v_1, v_2, \ldots, v_m$ be the vertices on the unique path in A from $v_1 = L(A)$ to $v_m = R(A)$. Then A can be uniquely expressed as a concatenation $B_1 E B_2 E \cdots E B_m$, where each B_i is a diagonal polarized tree with $L(B_i) = R(B_i) = v_i$, and E is a "polarized edge" (the polarized tree with two vertices whose poles are different). It follows from Theorem 5.22 that

$$\mathbf{G}_A = \mathbf{G}_{B_1} M \mathbf{G}_E M \mathbf{G}_{B_2} M \mathbf{G}_E M \cdots M \mathbf{G}_{B_m}.$$

The matrices M and $\mathbf{G}_{B_i}$ are nonsingular because they are diagonal with nonzero diagonal entries, and $\mathbf{G}_E$ is nonsingular because we can explicitly calculate

$$\mathbf{G}_E = \begin{bmatrix} 1 & xy \\ xy & x^2 z \end{bmatrix}$$

so that $\det \mathbf{G}_E = x^2(z - y^2) \neq 0$. Thus $\mathbf{G}_A$ is nonsingular, as desired. □

5.5 Recurrences for the generalized degree polynomial

In the rest of this chapter and all of the next, we explore several ways in which the generalized degree polynomials of rooted and polarized graphs can be used to derive nontrivial facts about the ordinary GDP. A running theme is that such results often arise from linear dependences between polarized GDPs, which can be exploited by using concatenation.

Recall from Proposition 4.14 that there is an algebra map γ that sends the chromatic symmetric function to the generalized degree polynomial (up to factors of y). By applying the map γ to the deletion–near-contraction recurrence for the CSF (Proposition 2.7), we can obtain one possible recurrence relation for the GDP, namely

$$\mathrm{G}_T = y\mathrm{G}_{T-e} + \mathrm{G}_{T \odot e} - y(x+1)\mathrm{G}_{T/e}$$

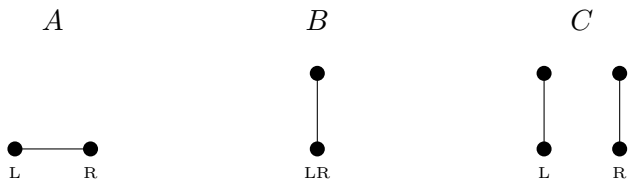

Figure 5.5: The graphs A, B, and C from the proof of Theorem 5.38.

for trees T and (non-leaf) edges e. However, this recurrence is not a particularly efficient way to compute G_T because $T \odot e$ and T/e are about as large as T itself. In this section, we use the technology of rooted and polarized GDPs to derive a pair of more efficient recurrence relations for the ordinary GDP.

Theorem 5.38. Let T be a tree and e an (non-leaf) edge of T. Let T_1 and T_2 be the two components of $T - e$, and set $T_1' = T_1 \cup \{e\}$ and $T_2' = T_2 \cup \{e\}$. Then

$$\mathrm{G}_T = \frac{y}{(xz+y)(xy+1)}\mathrm{G}_{T_1'}\mathrm{G}_{T_2'} + \frac{x(z-y^2)}{(xz+y)(xy+1)}\mathrm{G}_{T\odot e}. \tag{5.10}$$

Proof. Let $e = \{u, v\}$, where u is a vertex of T_1 and v is a vertex of T_2. Consider T_1 to be rooted at u and T_2 to be rooted at v. Then T can be written as the concatenation $T_1 A T_2$, where A is a polarized edge. Similarly, the graph $T \odot e$ can be written as $T_1 B T_2$, where B is an edge with equal poles, and $T_1' \sqcup T_2'$ can be written as $T_1 C T_2$, where C is the disjoint union of two edges with one pole in each component. (Figure 5.5 depicts A, B, and C.)

The key fact is that the polarized graphs A, B, and C have linearly dependent GDPs: namely, one can verify by direct computation that

$$\mathbf{G}_A = \frac{x(z-y^2)}{(xz+y)(xy+1)}\mathbf{G}_B + \frac{y}{(xz+y)(xy+1)}\mathbf{G}_C. \tag{5.11}$$

Multiplying on the left by $\mathbf{G}_{T_1} M$ and on the right by $M\mathbf{G}_{T_2}$, Theorem 5.22 gives

$$\mathrm{G}_T = \frac{x(z-y^2)}{(xz+y)(xy+1)}\mathrm{G}_{T\odot e} + \frac{y}{(xz+y)(xy+1)}\mathrm{G}_{T_1'\sqcup T_2'}.$$

We have $\mathrm{G}_{T_1' \sqcup T_2'} = \mathrm{G}_{T_1'} \mathrm{G}_{T_2'}$, so we are done. □

Thus, this recurrence relation comes as a result of the linear dependence between the polarized GDPs $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$. The existence of such a linear relation, with only three terms, may be surprising, since 2×2 matrices over $\mathbb{Q}(x, y, z)$ form a 4-dimensional vector space. However, later on (in Section 6.2), we will see a general method for generating such relations, of which this is the simplest one.

Theorem 5.38 allows us to compute G_T for any tree T by choosing e to be any non-leaf edge[6] and then recursively computing the GDPs of T_1', T_2', and $T \odot e$. This process terminates: even though $T \odot e$ has the same number of vertices as T, it has one fewer non-leaf edge than T. Thus, the base cases of this recurrence are the stars St_n (just like the deletion–near-contraction recurrence). We have

$$\mathrm{G}_{\mathrm{St}_n} = (xy + 1)^{n-1} + x(xz + y)^{n-1}$$

by taking cases on whether or not the vertex subset contains the central vertex.

Remark 5.39. In [4, Section 5], the authors obtained a similar recurrence relation for the half-generalized degree polynomial H_T. In fact, it can be shown that our recurrence (Theorem 5.38) specializes to theirs in the same way that G_T specializes to H_T; we leave the details of the computation to the interested reader.

Although the ordinary GDP can be effectively computed using (5.10), by changing the underlying linear dependence, we can obtain a different recurrence relation for the ordinary GDP that expresses G_T in terms of the GDPs of four smaller trees, rather than two smaller trees (T_1' and T_2') and one tree of the same size ($T \odot e$).

Theorem 5.40. Let T be a tree and e an (non-leaf) edge of T. Let T_1 and T_2 be the two

[6] If we choose e to be a leaf edge, then (without loss of generality) $T_1' \cong T \odot e \cong T$ and T_2' is the tree with two vertices, and so one can check that (5.10) reduces to $\mathrm{G}_T = \mathrm{G}_T$.

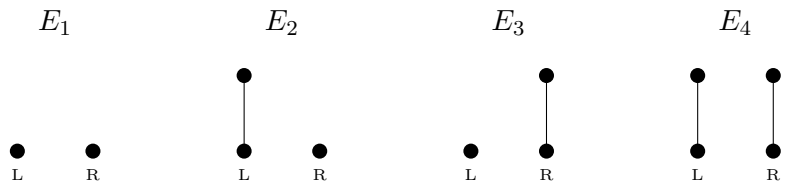

Figure 5.6: The graphs E_1, E_2, E_3, E_4 from the proof of Theorem 5.40.

components of $T - e$, and set $T_1' = T_1 \cup \{e\}$ and $T_2' = T_2 \cup \{e\}$. Then

$$\mathrm{G}_T = u_1 \mathrm{G}_{T_1} \mathrm{G}_{T_2} + u_2 \left(\mathrm{G}_{T_1'} \mathrm{G}_{T_2} + \mathrm{G}_{T_1} \mathrm{G}_{T_2'} \right) + u_3 \mathrm{G}_{T_1'} \mathrm{G}_{T_2'}, \tag{5.12}$$

where the constants u_1, u_2, u_3 are given by

$$u_1 = \frac{(z-y^2)(x^2z+2xy+1)}{(xz+y-xy-1)^2}, \quad u_2 = -\frac{(z-y^2)(x+1)}{(xz+y-xy-1)^2}, \quad u_3 = \frac{z-2y+1}{(xz+y-xy-1)^2}.$$

Proof. As in the proof of the previous recurrence, the idea is to express $\mathbf{G}_A$, where A is a polarized edge, in terms of the polarized GDPs of other polarized graphs. In this case, we use the four singular polarized graphs E_1, E_2, E_3, and E_4 depicted in Figure 5.6, for which it can be verified by direct calculation that

$$\mathbf{G}_A = u_1 \mathbf{G}_{E_1} + u_2 (\mathbf{G}_{E_2} + \mathbf{G}_{E_3}) + u_3 \mathbf{G}_{E_4}.$$

By concatenating on the left with T_1 and on the right with T_2, we obtain

$$\mathrm{G}_T = u_1 \mathrm{G}_{T_1 \sqcup T_2} + u_2 (\mathrm{G}_{T_1' \sqcup T_2} + \mathrm{G}_{T_1 \sqcup T_2'}) + u_3 \mathrm{G}_{T_1' \sqcup T_2'},$$

which yields the desired formula.

□

Unlike the linear dependence that underlies Theorem 5.38, the linear dependence be-

tween $\mathbf{G}_A$ and the $\mathbf{G}_{E_i}$ is one we would expect to exist because it involves the GDPs of five polarized graphs. (That $\mathbf{G}_{E_2}$ and $\mathbf{G}_{E_3}$ have the same coefficient u_2 is also to be expected, because $\mathbf{G}_{E_3} = \mathbf{G}_{E_2}^T$ while $\mathbf{G}_A$, $\mathbf{G}_{E_1}$, and $\mathbf{G}_{E_4}$ are symmetric.) While it has more terms than the previous relation, this one has the advantage that all the polarized graphs E_i are singular, and thus $\mathrm{G}_{T_1 E_i T_2}$ factors into smaller components. We chose the E_i so that in these factorizations, the GDPs of only four trees (namely T_1, T_2, T_1', and T_2') appear, further simplifying the computations.

Remark 5.41. When using SageMath [29] to compute the GDPs of trees for this project, we typically applied one of the above recurrence relations, together with memoization of previous results. We found that it was significantly faster to use (5.12) for this purpose than to use (5.10), especially for large-scale computations.

5.6 Divisibility results for the generalized degree polynomial

Next, we prove some nontrivial divisibility facts regarding the GDPs of ordinary, rooted, and polarized trees.[7] Our main result is the one for ordinary trees, Theorem 5.47. All our results will involve the polynomial

$$\eta = xz + y - xy - 1.$$

We first define a basic operation on rooted trees that we call *plucking*.

Definition 5.42. Let A be a rooted tree, and let v be a (non-root) vertex of A. To *pluck* A at v is to delete a non-root leaf vertex adjacent to v and then add a leaf vertex adjacent to $r(A)$.

The key calculation involves the relationship between the rooted GDPs of a rooted tree and the result of plucking it once.

[7] In this section, we work over the polynomial ring $\mathbb{Q}[x, y, z]$ in order to discuss divisibility.

Lemma 5.43. Let A be a rooted tree, let v be a (non-root) vertex of A, and let A' be the result of plucking A at v. Then η divides both entries of $\mathbf{G}_A - \mathbf{G}_{A'}$.

Proof. Let B be the result of removing a leaf vertex from A adjacent to v, considered as a polarized tree with $L(B) = v$ and $R(B) = r(A)$. Notice that A and A' can both be obtained from B by concatenation with an appropriate tree and then right-rooting. Specifically, let D be the unique diagonal polarized tree with two vertices. Then we have

$$A = (DB)_R \qquad \text{and} \qquad A' = (BD)_R.$$

We have $\mathbf{G}_D = \left[\begin{smallmatrix} xy+1 & 0 \\ 0 & x(xz+y) \end{smallmatrix}\right]$, so

$$\begin{aligned}
\mathbf{G}_A &= \begin{bmatrix} 1 & 1 \end{bmatrix} \mathbf{G}_D M \mathbf{G}_B \\
&= \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} xy+1 & 0 \\ 0 & x(xz+y) \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & x^{-1} \end{bmatrix} \begin{bmatrix} \mathrm{G}_B^1 & \mathrm{G}_B^R \\ \mathrm{G}_B^L & \mathrm{G}_B^{LR} \end{bmatrix} \\
&= \begin{bmatrix} (xy+1)\mathrm{G}_B^1 + (xz+y)\mathrm{G}_B^L & (xy+1)\mathrm{G}_B^R + (xz+y)\mathrm{G}_B^{LR} \end{bmatrix}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\mathbf{G}_{A'} &= \begin{bmatrix} 1 & 1 \end{bmatrix} \mathbf{G}_B M \mathbf{G}_D \\
&= \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \mathrm{G}_B^1 & \mathrm{G}_B^R \\ \mathrm{G}_B^L & \mathrm{G}_B^{LR} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & x^{-1} \end{bmatrix} \begin{bmatrix} xy+1 & 0 \\ 0 & x(xz+y) \end{bmatrix} \\
&= \begin{bmatrix} (xy+1)\mathrm{G}_B^1 + (xy+1)\mathrm{G}_B^L & (xz+y)\mathrm{G}_B^R + (xz+y)\mathrm{G}_B^{LR} \end{bmatrix}.
\end{aligned}$$

Subtracting these, we obtain

$$\mathbf{G}_A - \mathbf{G}_{A'} = \eta \begin{bmatrix} \mathrm{G}_B^L & -\mathrm{G}_B^R \end{bmatrix},$$

proving the claim. □

Using the above lemma, we can easily establish the following corollary that applies for any pair of rooted trees of the same order:

Corollary 5.44. If A and B are rooted trees with $|A| = |B|$, then η divides both entries of $\mathbf{G}_A - \mathbf{G}_B$.

Proof. Let $n = |A| = |B|$, and let S_n be the star graph of order n, rooted at its center vertex. Observe that by repeatedly plucking A at all its (non-root) vertices, we eventually obtain S_n. It follows from Lemma 5.43 and transitivity that both entries of $\mathbf{G}_A - \mathbf{G}_{S_n}$ are divisible by η. The same holds for B, so $\mathbf{G}_B - \mathbf{G}_{S_n}$ also has both entries divisible by η. Then the result follows since $\mathbf{G}_A - \mathbf{G}_B = (\mathbf{G}_A - \mathbf{G}_{S_n}) - (\mathbf{G}_B - \mathbf{G}_{S_n})$. □

We can also prove an analog of this result for polarized trees, given an extra hypothesis.

Theorem 5.45. If A and B are polarized trees of the same order and with the same distance between their poles, then η divides every entry of $\mathbf{G}_A - \mathbf{G}_B$.

Proof. In essence, we will prove this theorem by defining an analog of *plucking* for polarized trees and reasoning in the same way as above. However, exactly mimicking those arguments would require us to consider the GDP of a tree with *three* distinguished vertices, for which the definitions get much messier to write down. Fortunately, we have the following shortcut to cut down on notation.

Note that the vector $\mathbf{G}_{A_L} - \mathbf{G}_{B_L} = (\mathbf{G}_A - \mathbf{G}_B)\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]$, whose entries are the row sums of $\mathbf{G}_A - \mathbf{G}_B$, has all its entries divisible by η because of Corollary 5.44. For the same reason, the vector $\mathbf{G}_{A_R} - \mathbf{G}_{B_R}$, whose entries are the column sums of $\mathbf{G}_A - \mathbf{G}_B$, has all its entries divisible by η. It follows that

$$\mathbf{G}_A - \mathbf{G}_B \equiv \begin{bmatrix} t & -t \\ -t & t \end{bmatrix} \pmod{\eta}$$

for some t, where we are taking the entries modulo η. Therefore, it suffices to show that under the given hypotheses, any one entry of $\mathbf{G}_A - \mathbf{G}_B$, say $\mathrm{G}_A^1 - \mathrm{G}_B^1$, is divisible by η.

Now, let A be a polarized tree, and let $v \neq L(A)$ be a vertex of A that has an adjacent, non-pole leaf. Let C be the result of deleting that leaf and its corresponding edge from A, and let A' be the result of adding a leaf to C that is adjacent to $L(A)$. (The operation that turns A into A' is our equivalent of *plucking*.) We will show that $\mathrm{G}_A^1 - \mathrm{G}_{A'}^1$ is divisible by η. Then the desired conclusion will follow, because by repeated plucking, we can turn A into a star centered at $L(A)$ together with a path from $L(A)$ to $R(A)$, and this tree can also be obtained starting from any polarized tree with the same order and distance between poles.

Define

$$\mathrm{G}_C^1 = \sum_{L(A),R(A),v\notin S} x^{|S|}y^{d(S)}z^{e(S)} \qquad \text{and} \qquad \mathrm{G}_C^v = \sum_{\substack{L(A),R(A)\notin S\\ v\in S}} x^{|S|}y^{d(S)}z^{e(S)}.$$

(These would be two of the entries of the GDP of C with distinguished vertices $L(A)$, $R(A)$, and v.) Then a combinatorial argument gives

$$\mathrm{G}_A^1 = (xy+1)\mathrm{G}_C^1 + (xz+y)\mathrm{G}_C^v$$

and

$$\mathrm{G}_{A'}^1 = (xy+1)\mathrm{G}_C^1 + (xy+1)\mathrm{G}_C^v.$$

Hence $\mathrm{G}_A^1 - \mathrm{G}_{A'}^1 = \eta\mathrm{G}_C^v$, completing the proof. □

Corollary 5.46. If A is a polarized tree, then η divides $\mathrm{G}_A^L - \mathrm{G}_A^R$.

Proof. Recall from (5.1) that the *opposite* A^* of A satisfies $\mathbf{G}_{A^*} = \mathbf{G}_A^T$. Then

$$\mathbf{G}_A - \mathbf{G}_{A^*} = \mathbf{G}_A - \mathbf{G}_A^T = \begin{bmatrix} 0 & \mathrm{G}_A^R - \mathrm{G}_A^L \\ \mathrm{G}_A^L - \mathrm{G}_A^R & 0 \end{bmatrix}$$

Thus the result follows from applying Theorem 5.45 with $B = A^*$. □

Now that we have the desired divisibility results for rooted and polarized GDPs, we can prove a surprising analog for ordinary GDPs, as well as for determinants of (differences between) polarized GDPs.

Theorem 5.47. If T and U are trees with $|T| = |U|$, then η^2 divides $\mathrm{G}_T - \mathrm{G}_U$.

Proof. Suppose that A is a rooted tree and plucking it once gives some A'. In the proof of Lemma 5.43, we computed that

$$\mathbf{G}_A - \mathbf{G}_{A'} = \eta \begin{bmatrix} \mathrm{G}_B^L & -\mathrm{G}_B^R \end{bmatrix}$$

for some polarized tree B. Then, adding these two entries, we have $\mathrm{G}_A - \mathrm{G}_{A'} = \eta(\mathrm{G}_B^L - \mathrm{G}_B^R)$. This is divisible by η^2 due to Corollary 5.46.

In other words: for any ordinary tree T, if we arbitrarily choose a root for it, pluck once, and interpret the result as an ordinary tree T', then $\mathrm{G}_T - \mathrm{G}_{T'}$ will be divisible by η^2. The result now follows by the same argument that we used to prove Corollary 5.44. Namely, any tree T of order n can be turned into the star S_n by repeatedly applying this operation, so η^2 divides $\mathrm{G}_T - \mathrm{G}_{S_n}$ for all T. Hence η^2 also divides $\mathrm{G}_T - \mathrm{G}_U = (\mathrm{G}_T - \mathrm{G}_{S_n}) - (\mathrm{G}_U - \mathrm{G}_{S_n})$. □

Theorem 5.48. If A, B are polarized trees with $|A| = |B|$, then η^2 divides $\det(\mathbf{G}_A - \mathbf{G}_B)$.

Proof. Note that Theorem 5.45 is not applicable, since A and B may not have the same distance between poles. Instead, we use the following trickery. For any 2×2 matrix $N = \left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right]$, we have the identity

$$\det N = ad - bc = a(a + b + c + d) - (a + b)(a + c).$$

Taking N to be $\mathbf{G}_A - \mathbf{G}_B$ and reinterpreting the terms on the right-hand side, we obtain

$$\det(\mathbf{G}_A - \mathbf{G}_B) = (\mathrm{G}_A^1 - \mathrm{G}_B^1)(\mathrm{G}_A - \mathrm{G}_B) - \left(\mathrm{G}_{A_L}^1 - \mathrm{G}_{B_L}^1\right)\left(\mathrm{G}_{A_R}^1 - \mathrm{G}_{B_R}^1\right).$$

By Theorem 5.47, $\mathrm{G}_A - \mathrm{G}_B$ is divisible by η^2. By Corollary 5.44, both $\mathrm{G}^1_{A_L} - \mathrm{G}^1_{B_L}$ and $\mathrm{G}^1_{A_R} - \mathrm{G}^1_{B_R}$ are divisible by η. The result follows. □

Chapter 6

TREES WITH EQUAL GENERALIZED DEGREE POLYNOMIALS

In this final chapter, we detail our main application of the theory of rooted and polarized trees, which is to construct instances in which pairs (or larger families) of ordinary trees have the same GDP. As in Section 5.5, linear dependences between the GDPs of polarized graphs are our main tool for finding these instances. We present data of such instances for all trees with at most 23 vertices, and show that our constructions are enough to recover nearly all of them.

Furthermore, we show in certain cases that the trees in such instances, which are not distinguished by their GDPs, are still distinguished by the chromatic symmetric function. This gives some evidence that Question 1.1 has a positive answer and suggests a program for further work on this question using the GDP as an intermediary.

6.1 The reversal phenomenon

Our constructions fall under two main frameworks, called the *reversal phenomenon* and the *switching phenomenon*. We begin with the reversal phenomenon, which is the simpler of the two to describe. (Despite this, the particular constructions that arise from the reversal phenomenon are less well-understood.)

Recall from Section 5.2 that a polarized graph is *singular* if its poles lie in different components. As before, we use $\underset{\text{LR}}{\bullet}$ to denote the unique polarized graph with one vertex, whose GDP is $\mathbf{G}_{\bullet} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & x \end{smallmatrix}\right]$.

Definition 6.1. Let A and B be polarized trees, and let C be a singular polarized forest. We say that the triple $(A, B; C)$ satisfies the *reversal phenomenon* if $\mathbf{G}_A$, $\mathbf{G}_B$, $\mathbf{G}_C$, and $\mathbf{G}_{\bullet}$ are linearly dependent.

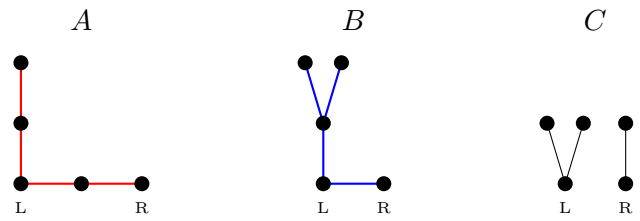

Figure 6.1: A triple of polarized graphs $(A, B; C)$ satisfying the reversal phenomenon.

Note that $(A, B; C)$ trivially satisfies the reversal phenomenon if two of A, B, and $\bullet_{\mathrm{LR}}$ are the same. However, these cases will not produce interesting results. Here is a more nontrivial example of the reversal phenomenon.

Example 6.2. The polarized graphs A, B, and C depicted in Figure 6.1 satisfy the reversal phenomenon. Specifically, by direct computation, one can check that

$$u_1\mathbf{G}_A + u_2\mathbf{G}_B + u_3\mathbf{G}_C + u_4\mathbf{G}_\bullet = 0,$$

where the constants u_1, u_2, u_3, u_4 are given by

$$\begin{aligned} u_1 &= xz + xy + y + 1, \\ u_2 &= -(xz + 1), \\ u_3 &= -y(xz + 1), \\ u_4 &= xy(z - y^2)(x^3z^2 + x^2y^2 + 2x^2yz + xy^2 + 2xy + 1). \end{aligned}$$

$\Diamond$

Later in this section, we will produce one family of triples $(A, B; C)$ satisfying the reversal phenomenon, including the triple from Example 6.2. (This will give a more enlightening explanation for the existence of this example.) However, our overall understanding of the structure of all such triples remains incomplete. For now, we show how instances of the reversal phenomenon produce pairs of ordinary trees with the same GDP, which will also

explain the use of the word "reversal."

Theorem 6.3. Suppose that $(A, B; C)$ satisfies the reversal phenomenon. Then for any positive integer m and any $E_1, E_2, \ldots, E_m \in \{A, B\}$, the ordinary trees

$$(E_1 E_2 \cdots E_m)\#C \qquad \text{and} \qquad (E_m \cdots E_2 E_1)\#C$$

have the same generalized degree polynomial.

Proof. The four matrices $\mathbf{G}_A$, $\mathbf{G}_B$, $\mathbf{G}_C$, and $\mathbf{G}_\bullet$ must satisfy a nontrivial linear relation. However, $\mathbf{G}_C$ and $\mathbf{G}_\bullet$ are linearly independent (since C is singular and hence not diagonal), so the linear relation must involve at least one of $\mathbf{G}_A$ and $\mathbf{G}_B$. Thus, we may assume without loss of generality that we can write

$$\mathbf{G}_A = b\mathbf{G}_B + c\mathbf{G}_C + d\mathbf{G}_\bullet$$

for some scalars $b, c, d \in \mathbb{Q}(x, y, z)$.

We use induction on the number of As in the list $E_1, \ldots, E_m$. If this number is zero, then $E_1 = \ldots = E_m = B$, so the statement is trivially true. Now, suppose that $E_i = A$ for some i. Define polarized trees X and Y by $X = E_1 \cdots E_{i-1}$ and $Y = E_{i+1} \cdots E_m$. (If $i = 1$, then we set $X = \underset{\text{LR}}{\bullet}$. Similarly, if $i = m$, then $Y = \underset{\text{LR}}{\bullet}$.) By Theorem 5.22, we have

$$\begin{aligned}
\mathbf{G}_{E_1 E_2 \cdots E_m} &= \mathbf{G}_{XAY} \\
&= \mathbf{G}_X M \mathbf{G}_A M \mathbf{G}_Y \\
&= \mathbf{G}_X M (b\mathbf{G}_B + c\mathbf{G}_C + d\mathbf{G}_\bullet) M \mathbf{G}_Y \\
&= b\mathbf{G}_{XBY} + c\mathbf{G}_{XCY} + d\mathbf{G}_{XY}.
\end{aligned}$$

Applying the map $f \mapsto \langle f, \mathbf{G}_C \rangle$ to both sides, Theorem 5.27 gives

$$\mathrm{G}_{(E_1 E_2 \cdots E_m)\#C} = b\mathrm{G}_{(XBY)\#C} + c\mathrm{G}_{(XCY)\#C} + d\mathrm{G}_{(XY)\#C}.$$

Because C is singular, $(XCY)\#C$ can be written as a disjoint union $(X\#C) \sqcup (Y\#C)$. Hence

$$\mathrm{G}_{(E_1E_2\cdots E_m)\#C} = b\mathrm{G}_{(XBY)\#C} + c\mathrm{G}_{X\#C}\mathrm{G}_{Y\#C} + d\mathrm{G}_{(XY)\#C}. \tag{6.1}$$

Now, suppose that we perform the above calculations again but with the order of the E_i reversed. This yields

$$\mathrm{G}_{(E_m\cdots E_2E_1)\#C} = b\mathrm{G}_{(Y'BX')\#C} + c\mathrm{G}_{Y'\#C}\mathrm{G}_{X'\#C} + d\mathrm{G}_{(Y'X')\#C}, \tag{6.2}$$

where $X' = E_{i-1}\cdots E_1$ and $Y' = E_m\cdots E_{i+1}$.

Observe that XBY, X, Y, and XY are all concatenations of As and Bs that have strictly fewer As than $XAY = E_1E_2\cdots E_m$. Therefore, the induction hypothesis applies to each of these, implying that

$$\begin{aligned}
\mathrm{G}_{(XBY)\#C} &= \mathrm{G}_{(Y'BX')\#C},\\
\mathrm{G}_{X\#C} &= \mathrm{G}_{X'\#C},\\
\mathrm{G}_{Y\#C} &= \mathrm{G}_{Y'\#C},\\
\mathrm{G}_{(XY)\#C} &= \mathrm{G}_{(Y'X')\#C}.
\end{aligned}$$

It follows from comparing (6.1) and (6.2) that $\mathrm{G}_{(E_1E_2\cdots E_m)\#C} = \mathrm{G}_{(E_m\cdots E_2E_1)\#C}$, completing the induction and the proof. □

Theorem 6.3 shows that, for every triple $(A, B; C)$ that satisfies the reversal phenomenon, we obtain an infinite family of pairs of ordinary trees with the same GDP. Let us illustrate this connection using the triple from Example 6.2.

Example 6.4. Let A, B, and C be the polarized graphs from Example 6.2, which we saw satisfy the reversal phenomenon. Then by Theorem 6.3, both of the pairs of ordinary trees shown in Figure 6.2 have the same GDP. (In fact, the trees $(AB)\#C$ and $(BA)\#C$, which

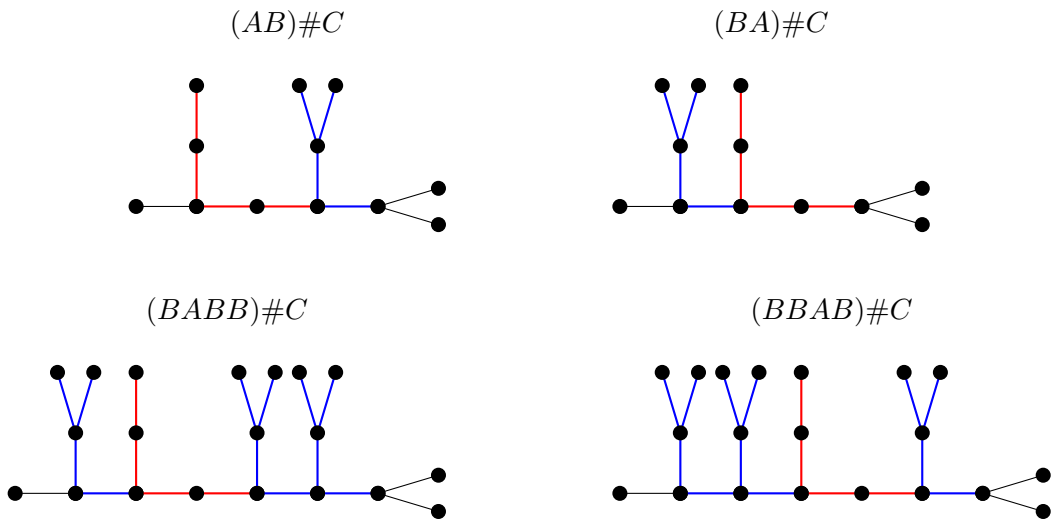

Figure 6.2: Two pairs of trees with the same GDP that arise from the reversal phenomenon.

have order 12, are the second-smallest pair of trees with the same GDP, after the trees with order 11 shown in Figure 3.1.) ◊

Next, we prove a converse of Theorem 6.3, essentially showing that a linear dependence of the above form is the *only* way in which such equalities of GDPs can arise.

Theorem 6.5. Let A and B be polarized trees, and let C be a singular polarized forest. If $\mathrm{G}_{(AB)\#C} = \mathrm{G}_{(BA)\#C}$, then $(A, B; C)$ satisfies the reversal phenomenon.

Proof. Suppose for contradiction that $\mathbf{G}_A$, $\mathbf{G}_B$, $\mathbf{G}_C$, and $\mathbf{G}_\bullet$ are linearly independent; then they span the (four-dimensional) space of 2×2 matrices over $\mathbb{Q}(x, y, z)$. Also, we may assume without loss of generality that C only has two components, so $C = P \otimes Q$ for rooted trees P and Q.

Given a 2×2 matrix f over $\mathbb{Q}(x, y, z)$, define

$$\varphi(f) = \langle f, \mathbf{G}_{AC} - \mathbf{G}_{CA} \rangle .$$

We claim that the function φ is identically zero. For any polarized graph D, we have

$$\varphi(\mathbf{G}_D) = \mathrm{G}_{D\#(AC)} - \mathrm{G}_{D\#(CA)} = \mathrm{G}_{(DA)\#C} - \mathrm{G}_{(AD)\#C}. \tag{6.3}$$

Then it is clear from (6.3) that $\varphi(\mathbf{G}_A) = \varphi(\mathbf{G}_\bullet) = 0$, as well as $\varphi(\mathbf{G}_B) = 0$ by our assumption. Furthermore, we have $(CA)\#C = C\#(AC) \cong (AC)\#C$, so $\varphi(\mathbf{G}_C) = 0$. It follows that φ must be identically zero.

But $\langle \cdot, \cdot \rangle$ is a nondegenerate bilinear form by Lemma 5.28, so we must have

$$\mathbf{G}_{AC} - \mathbf{G}_{CA} = 0.$$

This means that the singular polarized forests $AP \otimes Q$ and $P \otimes QA$ have the same GDP. However, note that if F is a singular polarized forest, then $\mathbf{G}_F$ determines the order of the component of F containing $L(F)$: indeed, that order is the unique positive integer k such that the monomial $x^k z^{k-1}$ appears in G_F^L. Thus, AP and P must have the same order, which forces $A = \underset{\mathrm{LR}}{\bullet}$, contradicting our initial assumption. □

The preceding two theorems combine to give a powerful tool for producing pairs of trees with the same generalized degree polynomial, even without knowing anything about the polarized GDP. Indeed, suppose that we have an example of trees T and U with have the same GDP. If we can write $T = (AB)\#C$ and $U = (BA)\#C$, where A, B are polarized trees and C is a singular polarized forest, then Theorem 6.3 and Theorem 6.5 show that $(E_1 \cdots E_m)\#C$ and $(E_m \cdots E_1)\#C$ have the same GDP for any m and $E_1, \ldots, E_m \in \{A, B\}$. Thus, the fact that the first pair of trees in Example 6.4 have the same GDP directly implies that the second pair does as well.

6.1.1 Examples of the reversal phenomenon

We have just seen how a triple of polarized graphs $(A, B; C)$ satisfying the reversal phenomenon gives rise to pairs of ordinary trees with the same GDP. Knowing this, the natural

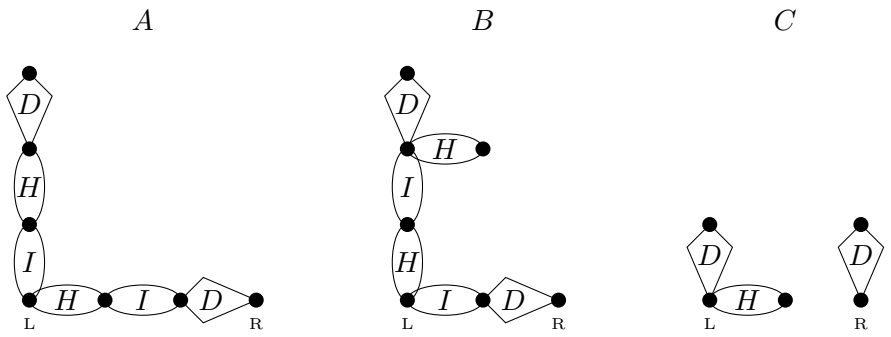

Figure 6.3: The polarized graphs from Theorem 6.6.

next step is to study these triples and attempt to characterize them. This appears to be a very difficult task, but as a partial result, we present one family of such triples. These were discovered by examining data of small trees with the same GDP; see Section 6.4 for details.

Theorem 6.6. Let D, H, and I be polarized trees, with H and I symmetric, such that HI and D commute. Let A, B, and C be polarized graphs constructed from D, H, and I as depicted[1] in Figure 6.3. Then $(A, B; C)$ satisfies the reversal phenomenon.

Note that if H is diagonal, then A and B are isomorphic, so the theorem is trivially true. Also, when D and H are polarized edges and $I = \underset{\text{LR}}{\bullet}$, we recover Example 6.2. In general, we have

$$D = (KL)^a \qquad H = (KL)^b K \qquad \text{and} \qquad I = L(KL)^c$$

where K and L are symmetric polarized trees and $a, b, c \geq 0$. (One can check that this is equivalent to the conditions on D, H, and I.)

We will prove Theorem 6.6 not by explicitly giving a linear relation between the four matrices $\mathbf{G}_A$, $\mathbf{G}_B$, $\mathbf{G}_C$, and $\mathbf{G}_\bullet$; instead, we will exhibit a nontrivial linear functional φ that annihilates them. (Since the space of 2×2 matrices over $\mathbb{Q}(x, y, z)$ is four-dimensional,

[1] We draw H and I as ovals because they are symmetric polarized graphs. Since D might not be symmetric, we draw it as a kite to specify its orientation; the "sharper" end is the right pole.

this will suffice to prove the linear dependence.) First, we need a definition.

Definition 6.7. Given a 2×2 matrix $U = \left[\begin{smallmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{smallmatrix}\right]$, the *asymmetry* of U is $\delta U = U_{21} - U_{12}$.

Thus, for a polarized tree A, we have $\delta \mathbf{G}_A = \mathrm{G}_A^L - \mathrm{G}_A^R$. We also have $\delta U = \operatorname{tr}(U\Omega)$, where $\Omega = \left[\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right]$ (the standard 2×2 symplectic matrix). Observe that $\delta \mathbf{G}_{A^*} = -\delta \mathbf{G}_A$, and $\delta \mathbf{G}_A = 0$ if and only if A is symmetric (by Proposition 5.14 and Theorem 5.37).

Here are some additional facts about the asymmetry of polarized trees.

Proposition 6.8. If A and B are doppelgängers, then $\delta \mathbf{G}_A = \delta \mathbf{G}_B$.

Proof. By Proposition 5.35, $\mathbf{G}_A - \mathbf{G}_B = t\left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$ for some t, so the result follows. □

Proposition 6.9. Let A and B be polarized trees. Then $\delta \mathbf{G}_{A^*BA} = \frac{\det \mathbf{G}_A}{x} \delta \mathbf{G}_B$.

Proof. The proof is purely computational; here is one way to do it. One can check that for any 2×2 matrix N, we have[2] $N\Omega N^T = (\det N)\Omega$. Then

$$\begin{aligned} \delta \mathbf{G}_{A^*BA} &= \delta \left((M\mathbf{G}_A)^T \mathbf{G}_B (M\mathbf{G}_A) \right) \\ &= \operatorname{tr} \left((M\mathbf{G}_A)^T \mathbf{G}_B (M\mathbf{G}_A) \Omega \right) \\ &= \operatorname{tr} \left(\mathbf{G}_B (M\mathbf{G}_A) \Omega (M\mathbf{G}_A)^T \right) \\ &= \operatorname{tr} \left(\mathbf{G}_B \det(M\mathbf{G}_A) \Omega \right) \\ &= \det(M\mathbf{G}_A) \delta \mathbf{G}_B. \end{aligned}$$

The conclusion follows because $\det M = 1/x$. □

Proposition 6.10. Suppose that polarized trees A and B commute. If $B \neq \underset{\mathrm{LR}}{\bullet}$, then $\mathbf{G}_A \in \operatorname{span}(\mathbf{G}_B, \mathbf{G}_\bullet)$.

Proof. Using the same idea as in Remark 5.23, we can translate this to an elementary fact about multiplication of 2×2 matrices. Specifically, we are given that the 2×2 matrices

[2] This corresponds to the fact that the symplectic group and special linear group coincide in dimension 2.

$\mathbf{G}_A M$ and $\mathbf{G}_B M$ commute under the usual matrix multiplication. In addition, we know that $\mathbf{G}_B M$ is not a scalar multiple of I_2 by Theorem 5.36. Then, it is a standard fact (which can be proved by analyzing the corresponding linear system) that $\mathbf{G}_A M \in \operatorname{span}(\mathbf{G}_B M, I_2)$, and this implies the desired result. □

Proposition 6.11. Let P and Q be rooted trees, and let A and B be commuting polarized trees. Then[3]

$$\delta\mathbf{G}\left[\ \right] = \frac{\det \mathbf{G}_B}{x}\,\delta\mathbf{G}\left[\ \right].$$

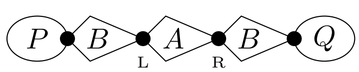

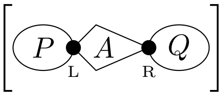

Proof. If $B = \underset{\mathrm{LR}}{\bullet}$, then this is obviously true, so assume otherwise. Then Proposition 6.10 implies that $\mathbf{G}_A \in \operatorname{span}(\mathbf{G}_B, \mathbf{G}_\bullet)$. Since the desired identity is linear in $\mathbf{G}_A$, we need only establish it for $A = B$ and $A = \underset{\mathrm{LR}}{\bullet}$. The former case can be done by direct computations like in the proof of Proposition 6.9, while the latter case is trivial as both sides are zero. □

Now we can prove our result.

Proof of Theorem 6.6. We first claim that ID is symmetric. If $D = \underset{\mathrm{LR}}{\bullet}$, then this follows because I is symmetric, so suppose otherwise. Since HI and D commute, Proposition 6.10 gives $\mathbf{G}_D \in \operatorname{span}(\mathbf{G}_{HI}, \mathbf{G}_\bullet)$, so $\mathbf{G}_{ID} \in \operatorname{span}(\mathbf{G}_{IHI}, \mathbf{G}_I)$. Since IHI and I are symmetric, we have $\delta\mathbf{G}_{ID} = 0$, so ID is symmetric as claimed.

Let φ be the linear functional such that for all polarized trees X,

$$\varphi(\mathbf{G}_X) = \det \mathbf{G}_S \cdot \delta\mathbf{G}_{UX} + \det \mathbf{G}_D \cdot \delta\mathbf{G}_{XV},$$

where $S =$ 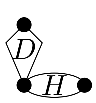

, $U =$ 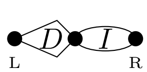

, and $V =$ 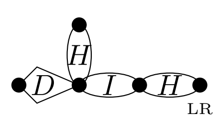

. Note that φ is not identically zero: for instance, when X is diagonal, XV is also diagonal and so $\delta\mathbf{G}_{XV} = 0$,

[3] For readability, we write $\mathbf{G}[C]$ instead of $\mathbf{G}_C$ when C is represented by a drawing.

but UX need not be symmetric. We will prove

$$\varphi(\mathbf{G}_A) = \varphi(\mathbf{G}_B) = \varphi(\mathbf{G}_C) = \varphi(\mathbf{G}_\bullet) = 0,$$

which will imply the desired linear dependence. When $X = B$ or $X = \underset{\text{LR}}{\bullet}$, both UX and XV are symmetric (note that $U = (ID)^* \cong ID$ and thus U is symmetric), so $\varphi(G_X) = 0$.

For $X = C$, we have

$$\delta\mathbf{G}_{UC} = \delta\mathbf{G}\left[\quad\right] = \frac{\det \mathbf{G}_D}{x}\,\delta\mathbf{G}\left[\quad\right]$$

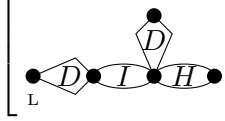

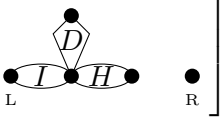

and

$$\delta\mathbf{G}_{CV} = \delta\mathbf{G}\left[\quad\right] = \frac{\det \mathbf{G}_S}{x}\,\delta\mathbf{G}\left[\quad\right],$$

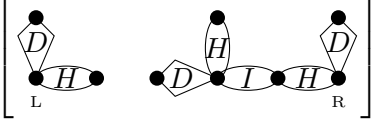

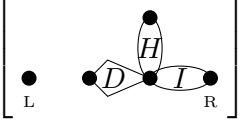

both by Proposition 6.9. Hence, $\det \mathbf{G}_S \cdot \delta\mathbf{G}_{UC} = -\det \mathbf{G}_D \cdot \delta\mathbf{G}_{CV}$, and so $\varphi(\mathbf{G}_C) = 0$.

For $X = A$, we have

$$\delta\mathbf{G}_{UA} = \delta\mathbf{G}\left[\quad\right] = \frac{\det \mathbf{G}_D}{x} \cdot \frac{\det \mathbf{G}_I}{x}\,\delta\mathbf{G}\left[\quad\right]$$

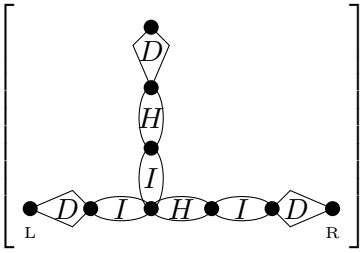

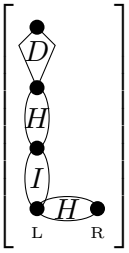

(again by Proposition 6.9). To simplify $\delta\mathbf{G}_{AV}$, we use Proposition 6.11, Proposition 6.9,

and Proposition 6.8 in that order to write

$$\delta\mathbf{G}_{AV} = \delta\mathbf{G}\left[\quad\right]$$

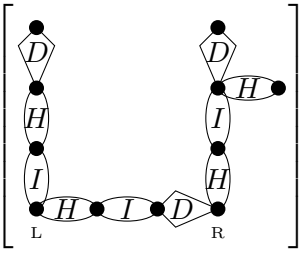

$$= \frac{\det \mathbf{G}_{HI}}{x}\,\delta\mathbf{G}\left[\quad\right]$$

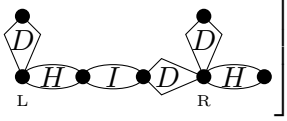

$$= \frac{\det \mathbf{G}_{HI}}{x} \cdot \frac{\det \mathbf{G}\left[\quad\right]}{x}\,\delta\mathbf{G}\left[\quad\right]$$

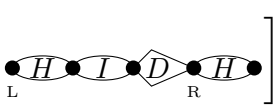

$$= \frac{\det \mathbf{G}_{I}}{x} \cdot \frac{\det \mathbf{G}_{S}}{x}\,\delta\mathbf{G}\left[\quad\right]$$

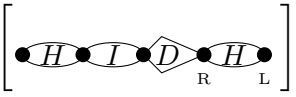

$$= \frac{\det \mathbf{G}_{I}}{x} \cdot \frac{\det \mathbf{G}_{S}}{x}\,\delta\mathbf{G}\left[\quad\right].$$

It follows that $\det \mathbf{G}_S \cdot \delta\mathbf{G}_{UA} = -\det \mathbf{G}_D \cdot \delta\mathbf{G}_{AV}$, and so $\varphi(\mathbf{G}_A) = 0$. □

6.2 The switching phenomenon

We now discuss the *switching phenomenon*, our second framework for constructing trees with equal generalized degree polynomials. Unlike the reversal phenomenon, which can only be used to find *pairs* of trees with the same GDP, the switching phenomenon can be used to construct arbitrarily large families of trees with this property.

The switching phenomenon is also defined by a linear dependence condition, a stronger version of the one defining the reversal phenomenon (Definition 6.1).

Definition 6.12. Let A and B be polarized trees, and let C be a singular polarized forest. We say that the triple $(A, B; C)$ satisfies the *switching phenomenon* if $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$ are linearly dependent.

(Throughout this section, we will implicitly assume that C only has two components

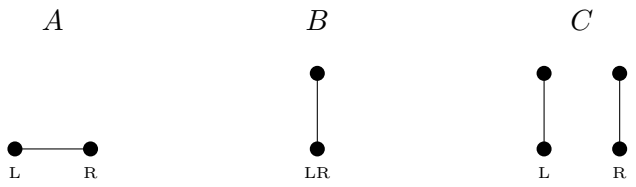

Figure 6.4: A triple $(A, B; C)$ satisfying the switching phenomenon.

and that $A \neq B$.)

Example 6.13. If A, B, and C are the polarized graphs shown in Figure 6.4, then $(A, B; C)$ satisfies the switching phenomenon. Indeed, these are the same graphs as in the proof of Theorem 5.38, and the linear relation between $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$ can be found there. $\Diamond$

Comparing Definition 6.12 with Definition 6.1, we see that any triple $(A, B; C)$ that satisfies the switching phenomenon must also satisfy the reversal phenomenon. Thus, we know from the previous section that certain pairs of trees of the form $(E_1 \cdots E_m)\#C$, for $E_1, \ldots, E_m \in \{A, B\}$, have the same GDP (namely, those related by reversing the order of the E_i). However, by using the additional strength of the switching phenomenon, we can show additional equalities of this kind. To state these equalities, we first introduce a way to associate each tree $(E_1 \cdots E_m)\#C$ with an integer composition.

Definition 6.14. Let A and B be polarized trees, and let C be a singular polarized forest. Given a composition $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$, we define a tree $T(\alpha)$ by

$$T(\alpha) = (B^{\alpha_1 - 1} A B^{\alpha_2 - 1} A \cdots A B^{\alpha_k - 1})\#C.$$

In other words, we have $T(\alpha) = (E_1 \cdots E_m)\#C$, where $m = |\alpha| - 1$ and

$$E_i = \begin{cases} A & \text{if } i \in \mathrm{Set}(\alpha) \\ B & \text{if } i \notin \mathrm{Set}(\alpha) \end{cases}$$

for $i = 1, \ldots, m$.

(Note that $T(\alpha)$ depends on A, B, C, but we suppress this in the notation for readability.)

Example 6.15. When A, B, and C are as in Figure 6.4, one can check that the tree $T(\alpha)$ is exactly the caterpillar $C(\alpha)$ as given in Definition 3.4. ◊

Suppose $\alpha = (\alpha_1, \ldots, \alpha_k)$ is a composition and $T(\alpha) = (E_1 \cdots E_m)\#C$ for $E_i \in \{A, B\}$. Then the reverse composition $\alpha^* = (\alpha_k, \ldots, \alpha_1)$ satisfies

$$T(\alpha^*) = (E_m \cdots E_1)\#C.$$

Thus, Theorem 6.3 can be rephrased in the following way: if $(A, B; C)$ satisfies the reversal phenomenon and the compositions α and β are reverses of each other, then $\mathrm{G}_{T(\alpha)} = \mathrm{G}_{T(\beta)}$. We now seek to prove an analog of this fact for the switching phenomenon, where the hypothesis "α and β are reverses of each other" is replaced by a weaker condition: that α and β are *related by switching*. Let us define this condition. (Our discussion closely follows the characterization of compositions with the same *ribbon Schur function* [9]. It extends the construction given in [4, Section 6] and is also structurally similar to the construction of [2, Section 7], which deals with the chromatic symmetric functions of *weighted* graphs.)

Given compositions $\alpha = (\alpha_1, \ldots, \alpha_k)$ and $\beta = (\beta_1, \ldots, \beta_\ell)$, recall the two operations

$$\alpha\beta = (\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_\ell) \qquad \text{(concatenation)}$$

and

$$\alpha \odot \beta = (\alpha_1, \ldots, \alpha_{k-1}, \alpha_k + \beta_1, \beta_2, \ldots, \beta_\ell). \qquad \text{(near-concatenation)}$$

Write $\beta^{\odot m}$ for the m-fold near-concatenation of β with itself: $\beta^{\odot 1} = \beta$, $\beta^{\odot 2} = \beta \odot \beta$, etc.

We now define a third operation on compositions, denoted by $\circ$, as follows:

$$\alpha \circ \beta = \beta^{\odot \alpha_1} \beta^{\odot \alpha_2} \cdots \beta^{\odot \alpha_k}.$$

For example, we have

$$23 \circ 413 = (413)^{\odot 2}(413)^{\odot 3} = (41713)(4171713) = 417134171713.$$

We leave as an exercise the verification of the following algebraic properties of $\circ$:

Proposition 6.16. If α, β, and γ are compositions, then

(i) $1 \circ \alpha = \alpha \circ 1 = \alpha$ (identity);

(ii) $(\alpha \circ \beta) \circ \gamma = \alpha \circ (\beta \circ \gamma)$ (associativity);

(iii) $(\alpha\beta) \circ \gamma = (\alpha \circ \gamma)(\beta \circ \gamma)$

(iv) $(\alpha \odot \beta) \circ \gamma = (\alpha \circ \gamma) \odot (\beta \circ \gamma)$; and

(v) $(\alpha \circ \beta)^* = \alpha^* \circ \beta^*$.

A *factorization* of a composition α is an equality $\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(m)}$ for compositions $\alpha^{(1)}, \ldots, \alpha^{(m)}$. (Because $\circ$ is associative, this expression can be interpreted unambiguously.) We say that the factorization is *nontrivial* if (i) $m \geq 2$; (ii) none of the $\alpha^{(i)}$ are the composition 1; (iii) no two consecutive $\alpha^{(i)}$ have length 1; and (iv) no two consecutive $\alpha^{(i)}$ have all parts equal to 1. Furthermore, we say that a composition α is *irreducible* if it has no nontrivial factorization. It turns out that every composition α has a unique factorization into irreducibles [9, Theorem 3.6]. We call this the *irreducible factorization* of α.

Definition 6.17. Let α and β be compositions. We say that α and β are *related by switching*, written $\alpha \sim_s \beta$, if the irreducible factorizations

$$\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(m)} \qquad \text{and} \qquad \beta = \beta^{(1)} \circ \cdots \circ \beta^{(m)}$$

have the same length, and for each i, either $\alpha^{(i)} = \beta^{(i)}$ or $\alpha^{(i)} = \left(\beta^{(i)}\right)^*$.

In particular, note that α and α^* are always related by switching: if $\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(m)}$ is the irreducible factorization of α, then Proposition 6.16(v) implies that[4]

$$\alpha^* = \left(\alpha^{(1)}\right)^* \circ \cdots \circ \left(\alpha^{(m)}\right)^*.$$

In addition, by unique factorization, $\sim_s$ is an equivalence relation on compositions.

We can now state our main result on the switching phenomenon.

Theorem 6.18. Suppose $(A, B; C)$ satisfies the switching phenomenon. If α and β are compositions with $\alpha \sim_s \beta$, then $\mathrm{G}_{T(\alpha)} = \mathrm{G}_{T(\beta)}$.

Example 6.19. Let $(A, B; C)$ be the polarized graphs from Figure 6.4. We take $\alpha = 12132$ and $\beta = 13212$, which are related by switching because $12132 = 12 \circ 12$ and $13212 = 21 \circ 12$. Then Theorem 6.18 tells us that the trees

$$T(12132) = (ABAABBAB)\#C \qquad \text{and} \qquad T(13212) = (ABBABAAB)\#C$$

have the same GDP. These are also the caterpillars $C(12132)$ and $C(13212)$, as noted in Example 6.15. In fact, these are the smallest pair of trees with the same GDP, as first presented in Figure 3.1. We have illustrated this construction in Figure 6.5.

In general, using the same triple $(A, B; C)$, we see that $\mathrm{G}_{C(\alpha)} = \mathrm{G}_{C(\beta)}$ if $\alpha \sim_s \beta$. $\Diamond$

Note that there are arbitrarily large families of compositions that are all related by switching. Indeed, if $\gamma^{(1)}, \ldots, \gamma^{(k)}$ are any compositions, then all compositions of the form $\delta^{(1)} \circ \cdots \circ \delta^{(k)}$, where each $\delta^{(i)}$ is either $\gamma^{(i)}$ or its reverse, are related by switching. If we take each $\gamma^{(i)}$ to be irreducible and non-palindromic, then all 2^k compositions of this form are guaranteed to be distinct by unique factorization. Taken together with the triple $(A, B; C)$

[4]To show that $\alpha \sim_s \beta$, it suffices to find any two factorizations $\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(k)}$ and $\beta = \beta^{(1)} \circ \cdots \circ \beta^{(k)}$ of the same length — not necessarily irreducible — such that $\alpha^{(i)}$ is either $\beta^{(i)}$ or its reverse for all i. The irreducible factorizations of α and β can then be obtained by further factoring the $\alpha^{(i)}$ and $\beta^{(i)}$, and Proposition 6.16(v) guarantees that Definition 6.17 will be satisfied.

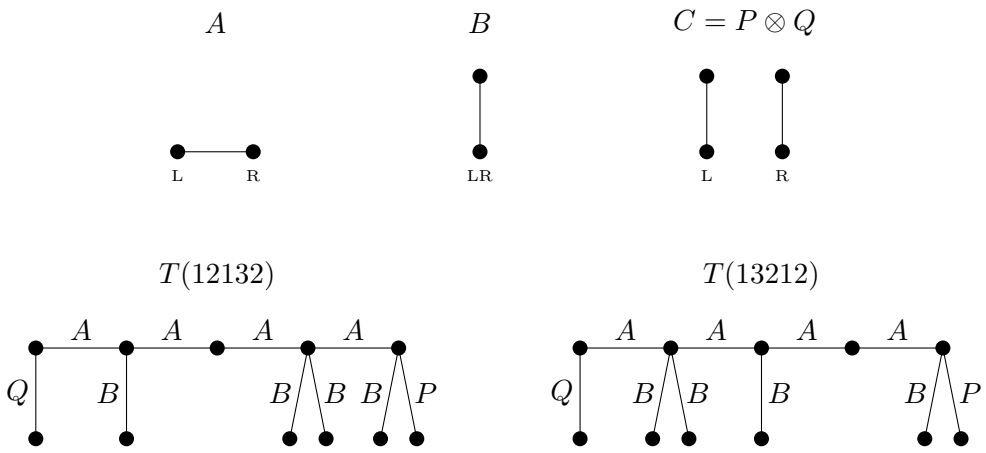

Figure 6.5: Using the switching phenomenon to recover the smallest pair of trees with the same GDP.

of Figure 6.4, this produces a family of 2^{k-1} distinct caterpillars with the same GDP (since $C(\alpha)$ and $C(\alpha^*)$ are isomorphic).

We will prove Theorem 6.18 via a sequence of lemmas. For brevity, let us write $\alpha \approx \beta$ to mean $\mathrm{G}_{T(\alpha)} = \mathrm{G}_{T(\beta)}$. Then our goal is to show that

$$\alpha \sim_s \beta \qquad \text{implies} \qquad \alpha \approx \beta.$$

By the reversal phenomenon (Theorem 6.3), we already know that $\alpha \approx \alpha^*$ for all α.

Lemma 6.20. Suppose $(A, B; C)$ satisfies the switching phenomenon. Then there exist $\kappa, \lambda \in \mathbb{Q}(x, y, z)$ such that, for any compositions α and β,

$$\mathrm{G}_{T(\alpha\beta)} = \kappa \mathrm{G}_{T(\alpha \odot \beta)} + \lambda \mathrm{G}_{T(\alpha)} \mathrm{G}_{T(\beta)}.$$

Hence, $\mathrm{G}_{T(\alpha\beta)}$ is determined by $\mathrm{G}_{T(\alpha\odot\beta)}$ and the unordered pair $\{\mathrm{G}_{T(\alpha)}, \mathrm{G}_{T(\beta)}\}$.

Proof. We are given that $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$ are linearly dependent. Neither of $\mathbf{G}_A$ and $\mathbf{G}_B$ (which are nonsingular matrices) can be a scalar multiple of $\mathbf{G}_C$ (which is singular), so

there must exist a linear relation of the form

$$\mathbf{G}_A = \kappa \mathbf{G}_B + \lambda \mathbf{G}_C \tag{6.4}$$

for some $\kappa, \lambda \in \mathbb{Q}(x, y, z)$. Now, fix compositions $\alpha = (\alpha_1, \ldots, \alpha_k)$ and $\beta = (\beta_1, \ldots, \beta_\ell)$. Define polarized trees $X = B^{\alpha_1 - 1} A \cdots A B^{\alpha_k - 1}$ and $Y = B^{\beta_1 - 1} A \cdots A B^{\beta_\ell - 1}$, so that $T(\alpha) = X \# C$ and $T(\beta) = Y \# C$. Then, one can check from the definitions that

$$\begin{aligned} T(\alpha\beta) &= (XAY)\#C, \\ T(\alpha \odot \beta) &= (XBY)\#C, \\ T(\alpha) \sqcup T(\beta) &= (XCY)\#C. \end{aligned}$$

So, in (6.4), we multiply on the left by $\mathbf{G}_X M$ and on the right by $M \mathbf{G}_Y$; then we apply the map $f \mapsto \langle f, \mathbf{G}_C \rangle$ to both sides. This produces the desired equation. □

Lemma 6.21. For any compositions α and β, we have $\alpha \circ \beta \approx \alpha \circ \beta^*$.

Proof. We proceed by induction on $\ell(\alpha)$, with β fixed. If $\ell(\alpha) = 1$, then $\alpha^* = \alpha$, so we have

$$\alpha \circ \beta = (\alpha^* \circ \beta^*)^* = (\alpha \circ \beta^*)^* \approx \alpha \circ \beta^*$$

where we applied Proposition 6.16(v). If $\ell(\alpha) \geq 2$, let $\alpha = \alpha'\alpha''$ for some nonempty compositions α' and α''. Then we have by Proposition 6.16(iii) that

$$\alpha \circ \beta = (\alpha'\alpha'') \circ \beta = (\alpha' \circ \beta)(\alpha'' \circ \beta).$$

Therefore, Lemma 6.20 and Proposition 6.16(iv) imply

$$\begin{aligned} \mathrm{G}_{T(\alpha \circ \beta)} &= \kappa \mathrm{G}_{T((\alpha' \circ \beta) \odot (\alpha'' \circ \beta))} + \lambda \mathrm{G}_{T(\alpha' \circ \beta)} \mathrm{G}_{T(\alpha'' \circ \beta)} \\ &= \kappa \mathrm{G}_{T((\alpha' \odot \alpha'') \circ \beta)} + \lambda \mathrm{G}_{T(\alpha' \circ \beta)} \mathrm{G}_{T(\alpha'' \circ \beta)}. \end{aligned}$$

On the other hand, we can perform a similar calculation but with β replaced by β^*, giving

$$\mathrm{G}_{T(\alpha\circ\beta^*)} = \kappa \mathrm{G}_{T((\alpha'\odot\alpha'')\circ\beta^*)} + \lambda \mathrm{G}_{T(\alpha'\circ\beta^*)}\mathrm{G}_{T(\alpha''\circ\beta^*)}.$$

By the induction hypothesis, we have

$$(\alpha' \odot \alpha'') \circ \beta \approx (\alpha' \odot \alpha'') \circ \beta^*,$$
$$\alpha' \circ \beta \approx \alpha' \circ \beta^*,$$
$$\alpha'' \circ \beta \approx \alpha'' \circ \beta^*.$$

Therefore we get $\mathrm{G}_{T(\alpha\circ\beta)} = \mathrm{G}_{T(\alpha\circ\beta^*)}$, completing the induction and the proof. □

Lemma 6.22. For all compositions α, β, γ, we have $\alpha \circ \beta \circ \gamma \approx \alpha \circ \beta^* \circ \gamma$.

Proof. This follows from Proposition 6.16(v) and two applications of Lemma 6.21:

$$\alpha \circ \beta \circ \gamma \approx \alpha \circ \beta \circ \gamma^* = \alpha \circ (\beta^* \circ \gamma)^* \approx \alpha \circ \beta^* \circ \gamma.$$ □

Finally, we can establish our theorem on the switching phenomenon.

Proof of Theorem 6.18. We are given compositions α and β with $\alpha \sim_s \beta$. Write

$$\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(m)} \qquad \text{and} \qquad \beta = \beta^{(1)} \circ \cdots \circ \beta^{(m)}$$

such that $\alpha^{(i)}$ is either $\beta^{(i)}$ or its reverse for all i. We define compositions $\gamma^{(0)}, \ldots, \gamma^{(m)}$ by

$$\gamma^{(j)} = \beta^{(1)} \circ \cdots \circ \beta^{(j)} \circ \alpha^{(j+1)} \circ \cdots \circ \alpha^{(m)}.$$

Thus, $\gamma^{(0)} = \alpha$ and $\gamma^{(m)} = \beta$. Observe that to get from $\gamma^{(j)}$ to $\gamma^{(j+1)}$, we replace $\alpha^{(j+1)}$ with $\beta^{(j+1)}$, which is either $\alpha^{(j+1)}$ or its reverse. It then follows (using Lemma 6.22 in the

latter case[5]) that $\gamma^{(j)} \approx \gamma^{(j+1)}$ for all $0 \leq j \leq m-1$. Thus $\alpha \approx \beta$ by transitivity. □

6.2.1 Examples of the switching phenomenon

Now that we have proved the main theorem behind the switching phenomenon, it remains to understand the triples $(A, B; C)$ that satisfy the switching phenomenon, according to Definition 6.12. Here is our main result on this matter.

Theorem 6.23. Suppose that A and B are polarized trees and D is a singular polarized forest such that $AD \cong BD$ and $DA \cong DB$. Then $(A, B; ADB)$ satisfies the switching phenomenon.

(Equivalently, if $D = P \otimes Q$ for rooted trees P and Q, then the condition is that $AP \cong BP$ and $QA \cong QB$ as rooted trees.) Note that Example 6.13 is an example of this theorem, where $D = \underset{\text{L}}{\bullet}\ \underset{\text{R}}{\bullet}$.

To prove this theorem, we will need a small lemma about 2×2 matrices:

Lemma 6.24. Suppose that X and Y are 2×2 matrices (over any field) such that Y has rank 1 and $XY = YX = 0$. Then $\operatorname{tr}(X) I_2 - X$ is a multiple of Y.

Proof. Let $Y = \left[\begin{smallmatrix} a \\ b \end{smallmatrix}\right]\left[\begin{smallmatrix} c & d \end{smallmatrix}\right] = \left[\begin{smallmatrix} ac & ad \\ bc & bd \end{smallmatrix}\right]$, where the vectors $\left[\begin{smallmatrix} a \\ b \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} c \\ d \end{smallmatrix}\right]$ are nonzero. Then $XY = 0$ means that $X\left[\begin{smallmatrix} a \\ b \end{smallmatrix}\right] = 0$, while $YX = 0$ means that $\left[\begin{smallmatrix} c & d \end{smallmatrix}\right]X = 0$. This forces $X = t\left[\begin{smallmatrix} bd & -ad \\ -bc & ac \end{smallmatrix}\right]$ for some t. Hence

$$\operatorname{tr}(X) I_2 - X = t(bd + ac)\left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right] - t\left[\begin{smallmatrix} bd & -ad \\ -bc & ac \end{smallmatrix}\right] = t\left[\begin{smallmatrix} ac & ad \\ bc & bd \end{smallmatrix}\right] = tY$$

as desired. □

Proof of Theorem 6.23. In this proof, let $\tilde{\mathbf{G}}_A = \mathbf{G}_A M$ for polarized trees A. Hence the formula from Theorem 5.22 becomes simply $\tilde{\mathbf{G}}_{AB} = \tilde{\mathbf{G}}_A \tilde{\mathbf{G}}_B$. (Also see Remark 5.23.)

---

[5] To be more explicit, we use associativity to write

$$\gamma^{(j)} = (\beta^{(1)} \circ \cdots \circ \beta^{(j)}) \circ \alpha^{(j+1)} \circ (\alpha^{(j+2)} \circ \cdots \circ \alpha^{(m)})$$

and then apply Lemma 6.22 to replace the second factor, $\alpha^{(j+1)}$, by its reverse, which gives $\gamma^{(j+1)}$.

By the Cayley-Hamilton theorem (or direct calculation), we have

$$X^2 = \mathrm{tr}(X)X - \det(X)I_2$$

for all 2×2 matrices X. Taking $X = \tilde{\mathbf{G}}_A$ and multiplying on the right by $\tilde{\mathbf{G}}_B$, we get

$$\tilde{\mathbf{G}}_{AAB} = \mathrm{tr}(\tilde{\mathbf{G}}_A)\tilde{\mathbf{G}}_{AB} - \det(\tilde{\mathbf{G}}_A)\tilde{\mathbf{G}}_B. \tag{6.5}$$

Instead taking $X = \tilde{\mathbf{G}}_B$ and multiplying on the left by $\tilde{\mathbf{G}}_A$, we get

$$\tilde{\mathbf{G}}_{ABB} = \mathrm{tr}(\tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_{AB} - \det(\tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_A. \tag{6.6}$$

We now subtract (6.6) from (6.5), which gives

$$\tilde{\mathbf{G}}_{AAB} - \tilde{\mathbf{G}}_{ABB} = \mathrm{tr}(\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_{AB} - \left(\det(\tilde{\mathbf{G}}_A)\tilde{\mathbf{G}}_B - \det(\tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_A\right),$$

or when rearranged,

$$\det(\tilde{\mathbf{G}}_A)\tilde{\mathbf{G}}_B - \det(\tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_A = \tilde{\mathbf{G}}_A\left[\mathrm{tr}(\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B)I_2 - (\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B)\right]\tilde{\mathbf{G}}_B. \tag{6.7}$$

Since $AD \cong BD$, we have $\tilde{\mathbf{G}}_{AD} = \tilde{\mathbf{G}}_{BD}$, so $(\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_D = 0$. Similarly, $DA \cong DB$ implies that $\tilde{\mathbf{G}}_D(\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B) = 0$. Therefore the matrices $\tilde{\mathbf{G}}_A - \tilde{\mathbf{G}}_B$ and $\tilde{\mathbf{G}}_D$ satisfy Lemma 6.24. This means that the bracketed quantity in (6.7) is $t \cdot \tilde{\mathbf{G}}_D$ for some $t \in \mathbb{Q}(x, y, z)$. Hence,

$$\det(\tilde{\mathbf{G}}_B)\tilde{\mathbf{G}}_A - \det(\tilde{\mathbf{G}}_A)\tilde{\mathbf{G}}_B = t \cdot \tilde{\mathbf{G}}_{ADB}.$$

Multiplying on the right by M^{-1}, we obtain a nontrivial linear relation between $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_{ADB}$. □

The preceding theorem gives a simple way to show that a triple $(A, B; C)$ satisfies the

switching phenomenon. Additionally, note that the condition on $(A, B; C)$ (namely, that $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$ are linearly dependent) is homogeneous. That is, if some $(A, B; C)$ satisfies the switching phenomenon, then so do $(AE, BE; CE)$ and $(EA, EB; EC)$ for any polarized tree E. This is because any linear relation

$$\kappa \mathbf{G}_A + \lambda \mathbf{G}_B + \mu \mathbf{G}_C = 0$$

can be multiplied on the right by $M\mathbf{G}_E$ to yield

$$\kappa \mathbf{G}_{AE} + \lambda \mathbf{G}_{BE} + \mu \mathbf{G}_{CE} = 0,$$

or on the left by $\mathbf{G}_E M$ to yield

$$\kappa \mathbf{G}_{EA} + \lambda \mathbf{G}_{EB} + \mu \mathbf{G}_{EC} = 0.$$

We do not know whether combining these two results (Theorem 6.23, as well as the ability to concatenate on the left and/or right by polarized trees) is enough to produce all nontrivial triples $(A, B; C)$ that satisfy the switching phenomenon.

For concreteness, we identify two common ways in which the conditions of Theorem 6.23 are met.

Corollary 6.25. Let H be a symmetric polarized tree and let P, Q be rooted trees. Let A, B, and C be the three polarized graphs depicted in Figure 6.6.[6] Then $(A, B; C)$ satisfies the switching phenomenon.

Proof. If $D = P \otimes Q$, then one can check that $AD \cong BD$, $DA \cong DB$, and $C = ADB$. □

Corollary 6.26. If A and B are doppelgängers, then $(A, B; C)$ satisfies the switching phenomenon where $C = A_L \sqcup A_R$.

[6] Here is a formal description of these graphs: $A = \hat{P}H\hat{Q}$ (recall that $\hat{P}$ denotes the polarization of P); B is the polarization of $(\hat{P}\hat{Q}H)_R$; and $C = (\hat{P}H\hat{P}Q) \otimes (P\hat{Q}H\hat{Q})$. We believe that the pictoral representation is more enlightening.

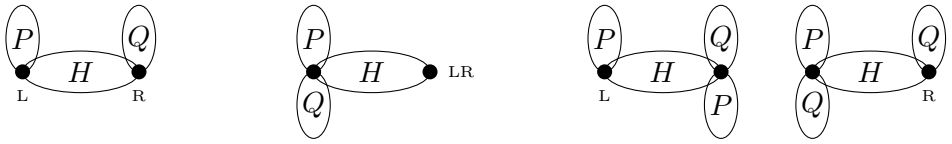

Figure 6.6: The graphs A, B, and C from Corollary 6.25.

Proof. Let D be the two-vertex polarized graph with no edges and distinct poles. Then we have $AD \cong BD$, $DA \cong DB$, and $C = ADB$. □

To summarize our discussion in this section, we present a recipe for using the switching phenomenon to generate a family of ordinary trees with the same GDP.

1. Find an example of polarized trees A, B and a singular polarized forest C such that $\mathbf{G}_A$, $\mathbf{G}_B$, and $\mathbf{G}_C$ are linearly dependent (Definition 6.12). Our Theorem 6.23, and more specifically Corollary 6.25 and Corollary 6.26, are possible sources of such an example. In addition, any example can be concatenated freely on the left and/or right by polarized trees (see the discussion after the proof of Theorem 6.23).

2. Choose a pair or larger family of compositions that are related by switching (Definition 6.17). This can be done by choosing any compositions $\gamma^{(1)}, \ldots, \gamma^{(k)}$ and then computing the expression $\delta^{(1)} \circ \cdots \circ \delta^{(k)}$ over all choices of $\delta^{(i)} \in \{\gamma^{(i)}, (\gamma^{(i)})^*\}$ for each i. (See also the discussion following Example 6.19.)

3. For each composition α chosen in the previous step, form the tree $T(\alpha)$ with respect to the chosen triple $(A, B; C)$ from the first step (Definition 6.14).

4. Then the resulting trees will all have the same GDP (Theorem 6.18).

6.3 Distinguishing trees with the CSF

Before concluding, we return to Question 1.1. In this chapter, we have seen a variety of constructions for families of trees with the same GDP. Since the GDP is determined by the CSF, studying these families could be fruitful in searching for an example of trees with the same CSF. However, we can show that certain applications of our constructions cannot produce such an example, which provides some evidence that Stanley's question has a positive answer. Let us give one instance of this for both the reversal phenomenon and the switching phenomenon.

6.3.1 Distinguishing trees from the reversal phenomenon

First, we revisit our construction for the reversal phenomenon from Section 6.1.1.

Theorem 6.27. Define the polarized graphs D, H, I, A, B, and C as in Theorem 6.6. Assume that D and H are not diagonal and $\deg R(D) = \deg L(H) = \deg R(H) = 1$. Then for $m \geq 2$, all trees of the form $(E_1 \cdots E_m)\#C$ for $E_1, \ldots, E_m \in \{A, B\}$ have distinct chromatic symmetric functions.

To prove this theorem, we will need some additional definitions and results.

Definition 6.28. Let T be a tree. A *shape* of T is a proper subtree $S \subsetneq T$ that has exactly one boundary edge: $d(S) = 1$.

For instance, a shape of order 1 is simply a leaf. (The term "shape" comes from [22].) Let $\sigma_a(T)$ denote the number of shapes of T with order a, and let $\sigma_{a,b}(T)$ denote the number of pairs of shapes (S_1, S_2) of T such that $|S_1| = a$, $|S_2| = b$, and $S_1 \subseteq S_2$. Observe that $\sigma_a(T) = [x^a y^1 z^{a-1}]\mathrm{G}_T(x, y, z)$, so the numbers $\sigma_a(T)$ are determined by G_T. On the other hand, the numbers $\sigma_{a,b}(T)$ are not determined by G_T in general, but certain linear combinations of them are determined by $\mathbf{X}_T$:

Lemma 6.29. Let $|T| = n$. If a and b are positive integers such that $a \leq b$ and $a+b \leq n-1$, then $\mathbf{X}_T$ determines $\sigma_{a,a+b}(T) + \sigma_{b,a+b}(T) - \sigma_{a,b}(T)$.

Proof. We use the higher-order generalized degree polynomials, defined in Section 4.5, which we proved are determined by the CSF. Specifically, consider the coefficient of the monomial $x_1^a x_2^b\, y^2 z_1^{a-1} z_2^{b-1}$ in $\mathrm{G}_T^{(2)}$. This is the number of pairs of disjoint subtrees (S_1, S_2) of T with order a and b, respectively, such that $d(S_1, S_2) = 2$. Then either S_1 and S_2 are disjoint shapes, or $S_1 \sqcup S_2$ is a shape of order $a + b$.

Since $a \leq b$ and $a + b \leq n - 1$, given any two shapes of order a and b, either they are disjoint or the shape of order a is contained in the shape of order b. Therefore, if S_1 and S_2 are disjoint shapes, then the number of pairs (S_1, S_2) equals $\sigma_a(T)\sigma_b(T) - \sigma_{a,b}(T)$. In the second case, if $S_1 \sqcup S_2$ is a shape of order $a + b$, then either S_1 is a shape of order a contained in $S_1 \sqcup S_2$ or S_2 is a shape of order b contained in $S_1 \sqcup S_2$. Therefore, the number of pairs (S_1, S_2) in this case is $\sigma_{a,a+b}(T) + \sigma_{b,a+b}(T)$. Thus,

$$[x_1^a x_2^b\, y^2 z_1^{a-1} z_2^{b-1}]\mathrm{G}_T^{(2)} = \sigma_{a,a+b} + \sigma_{b,a+b} + \sigma_a\sigma_b - \sigma_{a,b}.$$

Since G_T determines σ_a and σ_b (and $\mathbf{X}_T$ determines G_T and $\mathrm{G}_T^{(2)}$), the claim follows. □

We record an easy corollary of this lemma, which we will use in the proof of Theorem 6.6.

Corollary 6.30. Let T be a tree of order n and $2 \leq k \leq n - 2$. Suppose that either $\sigma_{k-1}(T) = 0$ or $\sigma_{k+1}(T) = 0$. Then $\mathbf{X}_T$ determines $\sigma_{1,k}(T)$.

Proof. If $\sigma_{k-1}(T) = 0$, then $\sigma_{1,k-1}(T) = \sigma_{k-1,k}(T) = 0$ as well, so taking $(a, b) = (1, k - 1)$ in Lemma 6.29 gives the result. If $\sigma_{k+1}(T) = 0$, then we can take $(a, b) = (1, k)$ instead. □

Proof of Theorem 6.27. Let $C = C_1 \otimes C_2$, and fix a tree

$$T = (E_1 \cdots E_m)\#C = C_2 E_1 \cdots E_m C_1$$

with $E_1, \ldots, E_m \in \{A, B\}$. (Figure 6.7 illustrates the case $m = 4$ and $T = (ABAA)\#C$.) We show how to recover the E_i from $\mathbf{X}_T$ by analyzing the shapes of T.

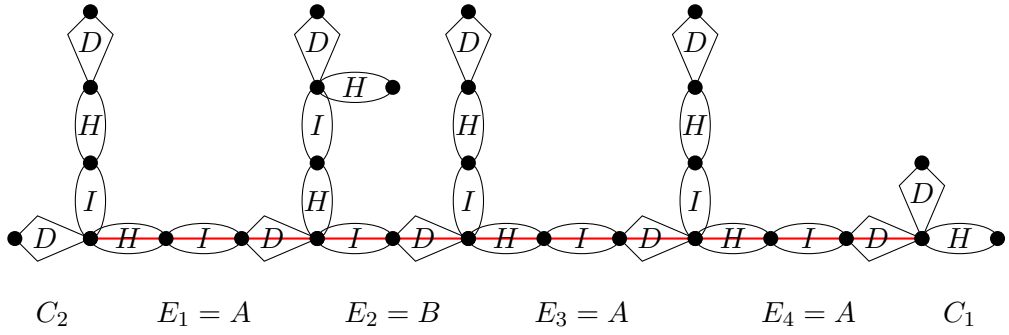

Figure 6.7: An example from the proof of Theorem 6.27. The path W is the thick red line.

It will be more convenient in this proof to consider graphs by number of edges, rather than number of vertices, since the number of edges is additive with respect to concatenation. Recall that we use $e(S)$ to denote the number of edges of S; let $d = e(D)$, $h = e(H)$, and $i = e(I)$. (By our assumptions, $d > 0$ and $h > 0$, but i can be 0.) Then

$$e(T) = (2d + 2h + 2i)m + (2d + h).$$

Let W be the unique path in T between the roots of C_2 and C_1. The shapes of T can be classified into three groups:

(i) The *outer shapes*, whose boundary edge does not lie on W;

(ii) The *left shapes*, whose boundary edge lies on W and which contain $r(C_2)$; and

(iii) The *right shapes*, whose boundary edge lies on W and which contain $r(C_1)$.

If S is an outer shape, then either $e(S) < d + 2h + i$ or $e(S) \geq e(T) - (d + 2h + i)$. Instead, if S is a left shape, then by inspection we have

$$e(S) \in \{2d + h + i, \ldots, 3d + 2h + 2i - 1\} \pmod{2d + 2h + 2i}.$$

(By this, we mean that $e(S)$ is congruent modulo $2d+2h+2i$ to one of the numbers in the set.) Analogously, if S is a right shape, then

$$e(S) \in \{d+h, \ldots, 2d+2h+i-1\} \pmod{2d+2h+2i}.$$

We now use this information to isolate shapes of T with specific orders. Fix an integer k with $1 \leq k \leq m-1$, and consider shapes S with

$$e(S) = (2d+2h+2i)k + (d+h).$$

Our observations show that there are no outer shapes or left shapes satisfying this condition. However, there is a unique right shape satisfying it, namely $S = E_{m-k+1} \cdots E_m C_1$. (Our assumptions on D guarantee that this is indeed a shape.) In addition, one can check that there are no shapes S with one fewer edge: $e(S) = (2d+2h+2i)k + (d+h-1)$. Thus, by Corollary 6.30, we can recover from $\mathbf{X}_T$ the number of leaves of T inside the shape $E_{m-k+1} \cdots E_m C_1$. Since k was arbitrary, this holds for all $1 \leq k \leq m-1$. Furthermore, our assumptions on H (namely, that it is not diagonal and has degree-1 poles) implies that A and B contribute different numbers of leaves to this count. Hence, letting k vary, this information determines whether each of $E_2, \ldots, E_m$ is A or B.

We can apply a similar argument to the left shapes. Fix an integer k with $1 \leq k \leq m-1$ and consider shapes S with

$$e(S) = (2d+2h+2i)k + (d-1).$$

There are no outer shapes or right shapes satisfying this condition, but there is a unique left shape, namely $C_2 E_1 \cdots E_k$ with the edge incident to $R(E_k)$ removed. (Again, we need the assumptions on D to know that this is a shape.) Furthermore, there are no shapes with one more edge: $e(S) = (2d+2h+2i)k + d$. Thus, by Corollary 6.30, $\mathbf{X}_T$ determines the number of leaves of T inside this shape. Letting k vary, this determines whether each of

$E_1, \ldots, E_{m-1}$ is A or B.

We have thus shown that $E_1, \ldots, E_m$ can be recovered from $\mathbf{X}_T$, so we are done. □

6.3.2 Distinguishing trees from the switching phenomenon

Next, we consider families of trees produced using the switching phenomenon. Certain such families consist of caterpillars, obtained by taking $(A, B; C)$ as in Example 6.19. Specifically, we noted that if α and β are compositions with $\alpha \sim_s \beta$, then $C(\alpha)$ and $C(\beta)$ have the same GDP. We now show that these two caterpillars do not have the same chromatic symmetric function (unless they are isomorphic).

Fix a composition $\alpha = (\alpha_1, \ldots, \alpha_\ell)$ and let $v_1, \ldots, v_\ell$ be the spine vertices of $C(\alpha)$, with $\deg v_i = \alpha_i + 1$. We ignore the shapes of $C(\alpha)$ of order 1 (which are leaves) and $|C(\alpha)| - 1$ (which are complements of leaves). Every other shape S contains either v_1 or v_ℓ, but not both. We call S a *left shape* if it contains v_1 and a *right shape* if it contains v_ℓ. Then $C(\alpha)$ has $\ell - 1$ left shapes, which are nested in each other and whose orders are one more than the elements of $\mathrm{Set}(\alpha)$. Similarly, $C(\alpha)$ has $\ell - 1$ right shapes, which are nested and whose orders are one more than the elements of $\mathrm{Set}(\alpha^*)$.

To establish the claimed result, we rely on the following lemma, which contains the meat of the proof.

Lemma 6.31. Fix a non-palindrome composition γ. Then $\mathbf{X}_{C(\beta \circ \gamma)}$ determines β for all compositions β.

Proof. Let $T = C(\beta \circ \gamma)$, and write $m = |\beta|$ and $n = |\gamma|$. (Note that $\mathbf{X}_T$ determines m because $|T| = |\beta \circ \gamma| + 2 = mn + 2$ and we know n.)

Let $s = \min(\mathrm{Set}(\gamma) \setminus \mathrm{Set}(\gamma^*))$, which exists because $\mathrm{Set}(\gamma)$ and $\mathrm{Set}(\gamma^*)$ are unequal but have the same cardinality. If $s = 1$, then we can replace γ with γ^*, since the statements of the lemma for γ and γ^* are equivalent. Thus, we may assume that $2 \leq s \leq n - 1$.

Using the definition of $\beta \circ \gamma$, one can check that

$$\mathrm{Set}(\beta \circ \gamma) = \{nb + c : 0 \leq b < m,\, c \in \mathrm{Set}(\gamma)\} \sqcup \{nb : b \in \mathrm{Set}(\beta)\}$$

and similarly

$$\mathrm{Set}(\beta^* \circ \gamma^*) = \{nb + c : 0 \leq b < m,\, c \in \mathrm{Set}(\gamma^*)\} \sqcup \{nb : b \in \mathrm{Set}(\beta^*)\}.$$

Thus, if $b \in [m-1]$, then the number of shapes of T of order $1 + nb$ (which is determined by $\mathbf{X}_T$) equals the number of true statements among "$b \in \mathrm{Set}(\beta)$" and "$b \in \mathrm{Set}(\beta^*)$". We now fix b such that exactly one of these statements is true and then determine which; this will be enough to determine β.

By Lemma 6.29, $\mathbf{X}_T$ determines the number

$$\kappa = \sigma_{s,\,1+nb+s} + \sigma_{1+nb,\,1+nb+s} - \sigma_{s,\,1+nb}.$$

Since $s \in \mathrm{Set}(\gamma) \setminus \mathrm{Set}(\gamma^*)$, it follows that $1+nb+s$ is the order of a left shape but not a right shape. Meanwhile, $1 + nb$ is either the order of a left shape (if $b \in \mathrm{Set}(\beta)$) or a right shape (if $b \in \mathrm{Set}(\beta^*)$), but not both. In addition, we know whether s is the order of a left shape or right shape (or both), because this depends on whether $s-1 \in \mathrm{Set}(\gamma)$ or $s-1 \in \mathrm{Set}(\gamma^*)$, respectively. Note that $s-1 \notin \mathrm{Set}(\gamma) \setminus \mathrm{Set}(\gamma^*)$ by minimality of s, so s cannot be the order of a left shape but not a right shape. This leaves three possibilities for s:

- Suppose that s is the order of both a left shape and a right shape. Then we have $\sigma_{s,\,1+nb+s} = \sigma_{s,\,1+nb} = 1$, so $\kappa = \sigma_{1+nb,\,1+nb+s}$. This equals 1 if $1 + nb$ is a left shape and 0 if $1 + nb$ is a right shape, so we can determine which is true.

- Suppose s is the order of neither a left shape nor a right shape. Then we have $\sigma_{s,\,1+nb+s} = \sigma_{s,\,1+nb} = 0$, so again $\kappa = \sigma_{1+nb,\,1+nb+s}$, and the logic from the previous case applies.

- Suppose s is the order of a right shape but not a left shape. Then $\sigma_{s,\,1+nb+s} = 0$, so $\kappa = \sigma_{1+nb,\,1+nb+s} - \sigma_{s,\,1+nb}$. If $1 + nb$ is the order of a left shape, then this equals $1 - 0 = 1$, and if $1 + nb$ is the order of a right shape, then this equals $0 - 1 = -1$. Thus, we can again determine which is true.

It follows that we can recover $\mathrm{Set}(\beta)$ and thus β. □

We can now prove the desired result:

Theorem 6.32. If $\alpha \sim_s \beta$ and $\beta \notin \{\alpha, \alpha^*\}$, then $\mathbf{X}_{C(\alpha)} \neq \mathbf{X}_{C(\beta)}$.

Proof. Let $\alpha = \alpha^{(1)} \circ \cdots \circ \alpha^{(k)}$ and $\beta = \beta^{(1)} \circ \cdots \circ \beta^{(k)}$, where $\beta^{(j)}$ is either $\alpha^{(j)}$ or its reverse for each j. Find the largest index i for which $\alpha^{(i)}$ is not a palindrome (which exists since $\alpha \neq \beta$). If $\beta^{(i)}$ is the reverse of $\alpha^{(i)}$, then we can replace β and all the $\beta^{(j)}$ with their reverses. Thus, we may assume that $\alpha^{(i)} = \beta^{(i)}$. We define

$$\alpha' = \alpha^{(1)} \circ \cdots \circ \alpha^{(i-1)}, \qquad \beta' = \beta^{(1)} \circ \cdots \circ \beta^{(i-1)},$$
$$\alpha'' = \alpha^{(i)} \circ \cdots \circ \alpha^{(k)}, \qquad \beta'' = \beta^{(i)} \circ \cdots \circ \beta^{(k)},$$

so that $\alpha = \alpha' \circ \alpha''$ and $\beta = \beta' \circ \beta''$. By our assumptions, we have $\alpha^{(j)} = \beta^{(j)}$ for all $j \geq i$, so $\alpha'' = \beta''$. In addition, $\alpha^{(i+1)}, \ldots, \alpha^{(k)}$ are all palindromes but $\alpha^{(i)}$ is not, so α'' and β'' are not palindromes either. Finally, we have $\alpha' \neq \beta'$ since $\alpha \neq \beta$. Thus, Lemma 6.31 implies that $\mathbf{X}_{C(\alpha)} \neq \mathbf{X}_{C(\beta)}$. □

To end this section, we state a related conjecture:

Conjecture 6.33. If $\mathrm{G}_{C(\alpha)} = \mathrm{G}_{C(\beta)}$ for compositions α and β, then $\alpha \sim_s \beta$.

We know the converse of this statement to be true due to the switching phenomenon (see Example 6.19). If this conjecture could be proved, then combining it with Theorem 6.32 would prove that non-isomorphic caterpillars have distinct chromatic symmetric functions. (Recall that a proof of this result was claimed in [22].)

n	11	12	13	14	15	16	17	18	19	20	21	22	23	Total
Pairs	1	1	1	5	1	7	19	15	23	56	22	90	160	401
Switching	1		1	5	1	4	18	11	17	48	21	82	154	363
Reversal		1				3	1	2	6	8	1	7	3	32
Unexplained								2				1	3	6

Table 6.1: The number of pairs of trees with at most 23 vertices that have the same GDP.

6.4 Future directions

We conclude this chapter (and this thesis) by considering some directions for extending our work and moving ever closer to resolving Question 1.1.

In this chapter, we discussed several ways to construct ordinary trees with the same GDP. An exhaustive computer search indicates that there are 401 pairs of trees with order at most 23 that have the same GDP, and all but six of them arise from the reversal phenomenon or switching phenomenon. (The switching phenomenon implies that there are families of arbitrarily many trees with the same GDP, but up to order 23, only pairs exist.) We have collected some summary statistics in Table 6.1,[7] and the remaining six pairs of trees appear in Figure 6.8.

Furthermore, in Section 6.3, we were able to prove that some of these constructions produce trees with different CSFs. Our work therefore suggests the following two-step program for resolving Stanley's question, which this thesis makes some progress towards executing:

(i) Characterize all pairs of trees with the same generalized degree polynomial.

(ii) For each pair, decide whether the trees have the same chromatic symmetric function.

[7]There is some overlap between the constructions produced by the reversal and switching phenomena. Specifically, if $(A, B; C)$ satisfies the switching phenomenon, and thus also the reversal phenomenon, then the fact that $\mathrm{G}_{T(\alpha)} = \mathrm{G}_{T(\alpha^*)}$ can be viewed as an instance of either phenomenon. For the purposes of Table 6.1, we considered this an instance of the switching phenomenon. That is, the triples $(A, B; C)$ that produce the 32 examples of the reversal phenomenon do *not* satisfy the switching phenomenon.

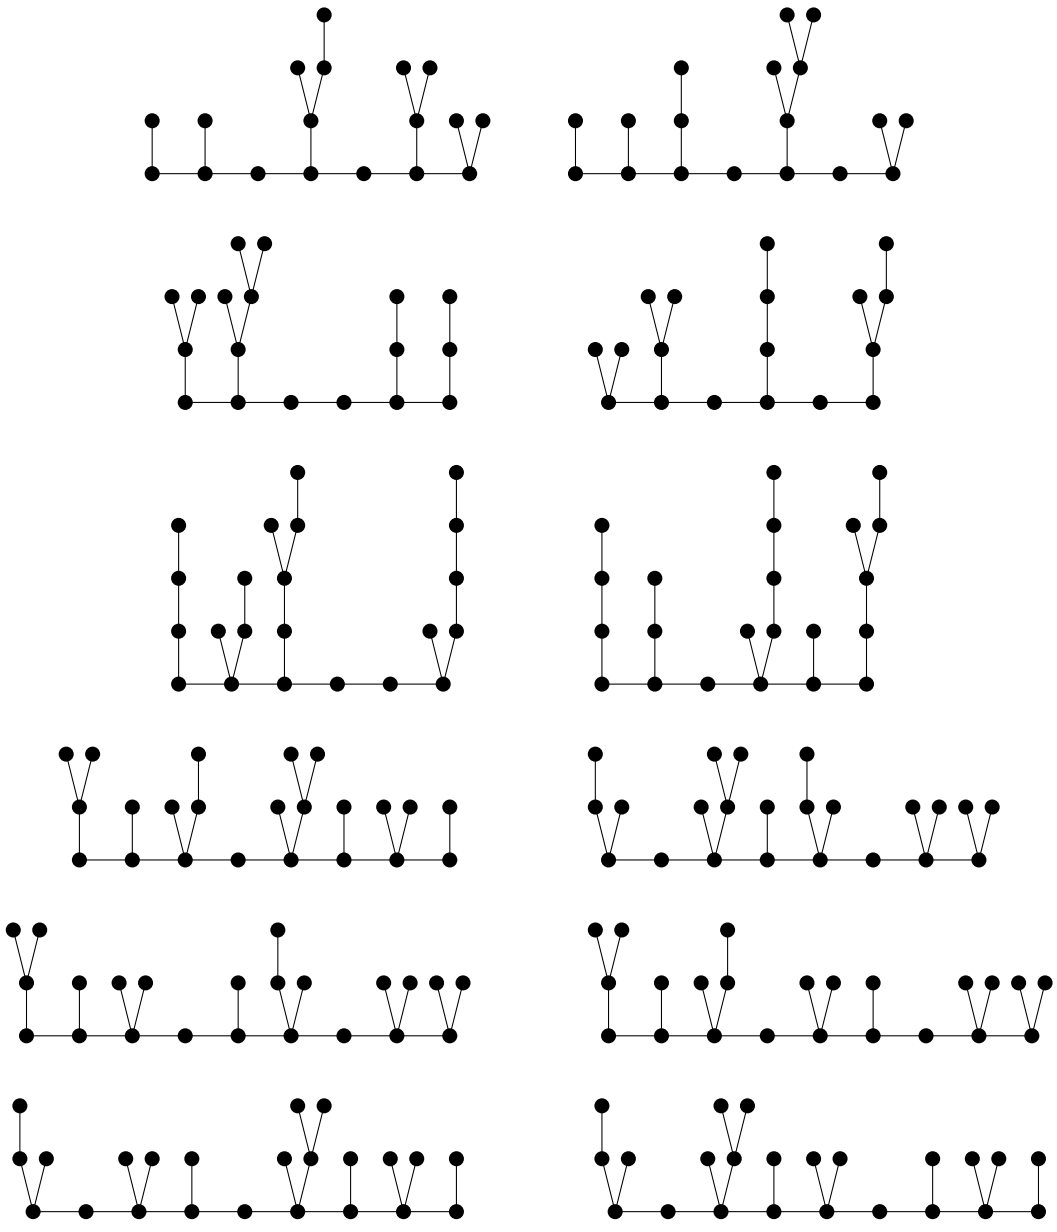

Figure 6.8: The six pairs of trees of order 23 with the same GDP that cannot be explained using the reversal phenomenon or switching phenomenon.

While we admit that this program is an ambitious one, we believe it is a promising direction in which to take the study of Stanley's question. Indeed, in our view, part of the difficulty in studying Stanley's question directly is the chromatic symmetric function is "too complex": in particular, we cannot examine concrete examples of trees with the same CSF to look for patterns. (Of course, this is an obvious statement, since Stanley's question asks whether any examples exist at all!) On the other hand, our results suggest that the generalized degree polynomial may lie in a certain "sweet spot" with regards to this problem:

(i) First, the GDP is simple enough that many "small" pairs of trees with the same GDP exist (and many can be well understood, as we have seen in this chapter). Furthermore, it is much easier to compute the GDP than the CSF, in part because of the recurrence relations of Section 5.5, so the GDP may be more amenable to computerized search methods.

(ii) On the other hand, the GDP is still a complex enough invariant that it determines a substantial amount of information about a tree (as we saw in Chapter 4). Because of this, we believe that there ought to exist an articulable explanation for the existence of most, if not all, pairs of trees with the same GDP (hopefully including the six pairs of Figure 6.8).

It is also reasonable to ask whether replacing the GDP with a similar but different invariant might result in a simpler or more elegant theory. For instance, one might consider building up the analog of this theory for the half-generalized degree polynomial or subtree polynomial (considered in Section 4.1). We know from Figure 4.2 that these are strictly weaker tree invariants than the GDP, but they still determine, for instance, the degree and path sequences of the tree [25].

In any case, we hope we have demonstrated that the generalized degree polynomial, as well as all its specializations, generalizations, and variations, are also interesting objects in

their own right. That they are related to the celebrated chromatic symmetric function is simply a cherry on top.

BIBLIOGRAPHY

[1] Marcelo Aguiar, Nantel Bergeron, and Frank Sottile. "Combinatorial Hopf algebras and generalized Dehn-Sommerville relations". In: *Compositio Mathematica* 142.01 (Jan. 2006), pp. 1–30.

[2] Farid Aliniaeifard, Victor Wang, and Stephanie van Willigenburg. "Extended chromatic symmetric functions and equality of ribbon Schur functions". In: *Advances in Applied Mathematics* 128 (2021), p. 102189. ISSN: 0196-8858.

[3] José Aliste-Prieto and José Zamora. "Proper caterpillars are distinguished by their chromatic symmetric function". In: *Discrete Mathematics* 315-316 (2014), pp. 158–164. ISSN: 0012-365X.

[4] José Aliste-Prieto et al. "Chromatic symmetric functions and polynomial invariants of trees". In: *Bulletin of the London Mathematical Society* 56.11 (2024), pp. 3452–3476.

[5] José Aliste-Prieto et al. "Marked Graphs and the Chromatic Symmetric Function". In: *SIAM Journal on Discrete Mathematics* 37.3 (2023), pp. 1881–1919.

[6] Cristina Ballantine et al. "On quasisymmetric power sums". In: *Journal of Combinatorial Theory, Series A* 175 (2020), p. 105273. ISSN: 0097-3165.

[7] Chris Berg et al. "A Lift of the Schur and Hall–Littlewood Bases to Non-commutative Symmetric Functions". In: *Canadian Journal of Mathematics* 66.3 (2014), pp. 525–565.

[8] Alexander Berkovich and Frank G. Garvan. "Some Observations on Dyson's New Symmetries of Partitions". In: *Journal of Combinatorial Theory, Series A* 100.1 (2002), pp. 61–93. ISSN: 0097-3165.

[9] Louis J. Billera, Hugh Thomas, and Stephanie van Willigenburg. "Decomposable compositions, symmetric quasisymmetric functions and equality of ribbon Schur functions". In: *Advances in Mathematics* 204.1 (2006), pp. 204–240. ISSN: 0001-8708.

[10] Soojin Cho and Stephanie van Willigenburg. "Chromatic Bases for Symmetric Functions". In: *Electron. J. Comb.* 23 (2015), p. 1.

[11] Logan Crew. "A note on distinguishing trees with the chromatic symmetric function". In: *Discrete Mathematics* 345.2 (2022), p. 112682. ISSN: 0012-365X.

[12] Logan Crew and Sophie Spirkl. "A deletion–contraction relation for the chromatic symmetric function". In: *European Journal of Combinatorics* 89 (2020), p. 103143. ISSN: 0195-6698.

[13] David Eisenstat and Gary Gordon. "Non-isomorphic caterpillars with identical subtree data". In: *Discrete Mathematics* 306.8 (2006), pp. 827–830. ISSN: 0012-365X.

[14] H. Gray Funkhouser. "A Short Account of the History of Symmetric Functions of Roots of Equations". In: *The American Mathematical Monthly* 37.7 (1930), pp. 357–365. ISSN: 00029890, 19300972.

[15] Arunkumar Ganesan et al. "Proper q-caterpillars are distinguished by their Chromatic Symmetric Functions". In: *Discrete Mathematics* 347.11 (2024), pp. 114–162. ISSN: 0012-365X.

[16] I.M. Gelfand et al. "Noncommutative Symmetrical Functions". In: *Advances in Mathematics* 112.2 (1995), pp. 218–348. ISSN: 0001-8708.

[17] Michael Gonzalez, Rosa Orellana, and Mario Tomba. "The chromatic symmetric function in the star-basis". In: *Discrete Mathematics* 349.2 (2026), p. 114691. ISSN: 0012-365X.

[18] Darij Grinberg and Victor Reiner. *Hopf Algebras in Combinatorics*. 2020. arXiv: 1409.8356 [math.CO].

[19] Tatsuyuki Hikita. *A proof of the Stanley–Stembridge conjecture*. 2025. arXiv: 2410.12758 [math.CO].

[20] Brian Hopkins and James Sellers. "On Blecher and Knopfmacher's fixed points for integer partitions". In: *Discrete Mathematics* 347 (May 2024), p. 113938.

[21] Ricky Ini Liu and Michael Tang. *Generalized degree polynomials of trees*. 2024. arXiv: 2411.18972. To appear in *Transactions of the AMS*.

[22] Martin Loebl and Jean-Sébastien Sereni. "Isomorphism of Weighted Trees and Stanley's Isomorphism Conjecture for Caterpillars". In: *Annales de l'Institut Henri Poincaré D* 6 (Jan. 2018).

[23] Nicholas A. Loehr and Gregory S. Warrington. *A rooted variant of Stanley's chromatic symmetric function*. 2023. arXiv: 2206.05392 [math.CO].

[24] Ian G. Macdonald. *Symmetric Functions and Hall Polynomials*. 2nd. Oxford University Press, 1995.

[25] Jeremy L. Martin, Matthew Morin, and Jennifer D. Wagner. "On distinguishing trees by their chromatic symmetric functions". In: *Journal of Combinatorial Theory, Series A* 115.2 (2008), pp. 237–253. ISSN: 0097-3165.

[26] Rosa Orellana and Geoffrey Scott. "Graphs with equal chromatic symmetric functions". In: *Discrete Mathematics* 320 (2014), pp. 1–14. ISSN: 0012-365X.

[27] Rosa Orellana and Foster Tom. *Linear relations on star coefficients of the chromatic symmetric function*. 2026. arXiv: 2601.13390 [math.CO].

[28] Mercedes H. Rosas. "MacMahon Symmetric Functions, the Partition Lattice, and Young Subgroups". In: *Journal of Combinatorial Theory, Series A* 96.2 (2001), pp. 326–340. ISSN: 0097-3165.

[29] W. A. Stein et al. *Sage Mathematics Software (Version 10.6)*. Accessible online at `http://www.sagemath.org`. The Sage Development Team. 2025.

[30] John Shareshian and Michelle L. Wachs. “Chromatic quasisymmetric functions”. In: *Advances in Mathematics* 295 (2016), pp. 497–551. ISSN: 0001-8708.

[31] Richard P. Stanley. “A Symmetric Function Generalization of the Chromatic Polynomial of a Graph”. In: *Advances in Mathematics* 111.1 (1995), pp. 166–194. ISSN: 0001-8708.

[32] Richard P. Stanley. “Graph colorings and related symmetric functions: ideas and applications. A description of results, interesting applications, & notable open problems”. In: *Discrete Mathematics* 193.1 (1998), pp. 267–286. ISSN: 0012-365X.

[33] Richard P. Stanley and Sergey Fomin. *Enumerative Combinatorics: Volume 2*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999.

[34] Richard P. Stanley and John R. Stembridge. “On immanants of Jacobi-Trudi matrices and permutations with restricted position”. In: *Journal of Combinatorial Theory, Series A* 62.2 (1993), pp. 261–279. ISSN: 0097-3165.

[35] Yuzhenni Wang, Xingxing Yu, and Xiao-Dong Zhang. “A class of trees determined by their chromatic symmetric functions”. In: *Discrete Mathematics* 347.9 (2024), p. 114096. ISSN: 0012-365X.