VERTEX-WEIGHTED GENERALIZATIONS OF CHROMATIC SYMMETRIC FUNCTIONS

Logan Crew

A DISSERTATION

in

Mathematics

Presented to the Faculties of the University of Pennsylvania in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2020

Co-Supervisor of Dissertation   Co-Supervisor of Dissertation

_____   _____

James Haglund,   Greta Panova, Associate
Professor of Mathematics   Professor of Mathematics

Graduate Group Chairperson

_____

Julia Hartmann, Professor of Mathematics

Dissertation Committee:
David Harbater, Professor of Mathematics
James Haglund, Professor of Mathematics
Greta Panova, Associate Professor of Mathematics at USC
Dennis DeTurck, Professor of Mathematics

Acknowledgments

In preparation.

ABSTRACT

VERTEX-WEIGHTED GENERALIZATIONS OF CHROMATIC SYMMETRIC

FUNCTIONS

Logan Crew

James Haglund, Greta Panova

Discovered by Richard Stanley in the early 1990s, the chromatic symmetric function X_G of a graph G enumerates for each integer partition λ of $|V(G)|$ the number of proper colorings of G that partition $V(G)$ into stable sets of sizes equal to the parts of λ. Thus, X_G is a refinement of the well-known chromatic polynomial χ_G, and its coefficients in different symmetric function bases provide further information on the structure of G than χ_G. However, X_G loses some of the utility of χ_G because it fails to admit a natural edge deletion-contraction relation. To address this shortcoming we introduce *vertex-weighted graphs* (G, w) consisting of a graph G and a weight function $w : V(G) \to \mathbb{N}$. Then X_G extends in a natural way to a new function $X_{(G,w)}$ on vertex-weighted graphs. We demonstrate that $X_{(G,w)}$ satisfies a deletion-contraction relation akin to that of the chromatic polynomial, and use this relation to both derive new properties of the chromatic symmetric function and prove previously known properties in an original way. In the case of prior results, the new proofs are typically simpler and more intuitive than the original proofs, and are more closely related to analogous proofs of properties of the chromatic polyno-

mial. We then demonstrate how the deletion-contraction relation can be used as a new tool to research open questions involving X_G. We also explore a similar extension of the bad-coloring chromatic symmetric function XB_G to vertex-weighted graphs, and we consider applications of these new functions to graph isomorphism and symmetric function bases.

Contents

Chapter 1

Introduction

Between numerous practical applications in scheduling, programming, and data analysis, as well as major results such as the Four-Color Theorem, it is no surprise that the study of proper graph colorings took off in the twentieth century, and the area continues to be among the most active in discrete mathematics. An early attempt to bring an algebraic perspective to the field was the discovery of the chromatic polynomial by Birkhoff [3], defined as the unique function χ_G satisfying that for any positive integer n, $\chi_G(n)$ is the number of proper n-colorings of G. Birkhoff showed that χ_G satisfies the deletion-contraction relation

$$\chi_G(x) = \chi_{G \backslash e}(x) - \chi_{G/e}(x)$$

where $G \backslash e$ denotes the graph G with the edge e removed, and G/e denotes the graph G with the endpoints of e identified. This surprisingly simple relation allows for recursive computation of χ_G, and provides a framework for discovering and proving

properties of χ_G using an inductive argument on the edges of a graph. For example, via an inductive argument combined with the fact that $\chi_{\overline{K_k}}(n) = n^k$ where $\overline{K_k}$ is a graph with k vertices and no edges, it is easy to verify that $\chi_G(n)$ is always a polynomial in n (see Section 2.4 for more details).

A further discovery in this direction was that of the chromatic symmetric function X_G of a graph G by Stanley in the early 1990s [38], defined as a power series (over $\mathbb{R}$) in countably many variables satisfying

$$X_G(x_1, x_2, \dots) = \sum_{\kappa} \prod_{v \in V(G)} x_{\kappa(v)}$$

where the sum ranges over all proper colorings κ of G. This function admits χ_G as a specialization, since $X_G(\underbrace{1, 1, \dots, 1}_{\text{n 1s}}, 0, 0, \dots) = \chi_G(n)$, but it allows additionally the usage of tools from algebra and symmetric function theory to study colorings of graphs. In particular, X_G can be expressed in a number of different bases of the space of symmetric functions in $x_1, x_2, \dots$, and the coefficients of X_G in different bases often give different information about the graph in a compact way. Some examples of this are given and extended in Chapter 3.

In this thesis we further extend X_G. The motivation for our construction is the observation that X_G does not admit a simple deletion-contraction relation like as χ_G does. This is because every monomial of X_G has degree $|V(G)|$, so trying to formulate such a relation will fail when considering edge contraction, which reduces the number of vertices. Instead, the best-known edge recurrence relation discovered

for X_G itself is a triangular relation discovered by Orellana and Scott in 2014 [32], and its generalization to all cycles in 2018 by Dahlberg and van Willigenburg [11].

To provide X_G with a deletion-contraction relation in a natural way, we extend the definition to include pairs (G, w) consisting of a graph G and a vertex-weight function $w : V(G) \to \mathbb{N}$. Then the extended function

$$X_{(G,w)}(x_1, x_2, \dots) = \sum_{\kappa} \prod_{v \in V(G)} x_{\kappa(v)}^{w(v)}$$

admits a generalization of the classic deletion-contraction relation of the chromatic polynomial of the form

$$X_{(G,w)} = X_{(G \backslash e, w)} - X_{(G/e, w/e)}$$

where w/e indicates that when we contract an edge e, the weight of the contacted vertex is the sum of the weights of the endpoints of e.

This approach builds upon previous extensions of the chromatic symmetric function that admit deletion-contraction relations. An early example was the chromatic symmetric function in noncommuting variables $Y_G(x_1, x_2, \dots)$, introduced by Gebhard and Sagan in 1999 [17], which satisfies $Y_G = Y_{G \backslash e} - Y_{G/e} \uparrow$, where $\uparrow$ is an operation they define called induction. This induction operation takes a variable x_i representing the color of a vertex and duplicates it, which is similar to our contraction term $X_{(G/e, w/e)}$ that adds vertex weights. More recently, the pointed chromatic symmetric function $X_{G,v}(t, x_1, x_2, \dots)$, rooted at a vertex v, was introduced by Pawlowski in 2018 [34] and satisfies $X_{G,v} = X_{G \backslash e, v} - t X_{G/e, v}$ when e is

an edge incident to v. There are many other results that the function $X_{(G,w)}$ and its properties build upon; several such examples will given throughout the thesis as appropriate.

The thesis will be organized as follows. In Chapter 2 we provide background on integer partitions, symmetric functions, and graph theory that will be used throughout this paper. We also provide some basic results of χ_G and discuss recent research on X_G to provide further context for our new results. Chapter 3 is the largest section and will discuss vertex-weighted generalizations of both the chromatic symmetric function X_G and its bad-coloring version XB_G as defined by Stanley in [41]. Finally, in Chapter 4, we discuss problems and applications of the newly defined $X_{(G,w)}$.

Some of the material for this chapter is taken from [7] (joint with Sophie Spirkl). In all future chapters, references to the author's preprints which contribute the material will be noted at the beginning of the chapter.

Chapter 2

Background

In this chapter, we give an overview of concepts and terminology that will be used throughout this work. Much of this background also appears in [7].

2.1 Integer Partitions

An *integer partition* (or just *partition*) is a tuple $\lambda = (\lambda_1, ..., \lambda_k)$ of positive integers such that $\lambda_1 \geq ... \geq \lambda_k$. The integers λ_i are the *parts* of λ. If $\sum_{i=1}^{k} \lambda_i = n$, we say that λ is a partition of n, and we write $\lambda \vdash n$, or $|\lambda| = n$. The number of parts k is the *length* of λ, and is denoted by $l(\lambda)$. The number of parts equal to i in λ is given by $r_i(\lambda)$. The *complement* (or conjugate) of λ is the partition λ' defined by $r_i(\lambda') = \lambda_i - \lambda_{i+1}$ for $i \leq l(\lambda)$.

Pictorially, we may associate to λ a diagram of boxes called its *Young diagram*, where the beginnings of the rows are aligned on the left, and the i^{th} row from the

top has λ_i boxes. Then λ' is described simply as the partition whose Young diagram is the reflection of the diagram for λ across its main diagonal. For example, the following is the Young diagram of $(4, 4, 1, 1)$ and its complement $(4, 2, 2, 2)$:

Given two partitions $\lambda = (\lambda_1, ..., \lambda_k)$ and $\mu = (\mu_1, ..., \mu_l)$ of n, we define a *puzzle of μ into λ* as an ordered tuple of partitions $(\mu^1, ..., \mu^k)$ such that

- For all $1 \leq i \leq k$ we have $\mu^i \vdash \lambda_i$.

- The disjoint union of the parts of the μ^i is μ.

If there exists a puzzle of μ into λ, we say that μ is a *refinement* of λ. Note that refinement induces a partial order on the set of partitions of n for any fixed n.

Given λ and μ partitions of n, we say that λ *dominates* μ if for all $1 \leq l \leq \max(l(\lambda), l(\mu))$ (defining $\lambda_i = 0$ if $i > l(\lambda)$ and likewise for μ) we have

$$\sum_{i=1}^{l} \mu_i \leq \sum_{i=1}^{l} \lambda_i$$

and we write $\mu \leq \lambda$. Dominance also induces a partial order on partitions of n for any fixed n. Both the partial order induced by refinement and the partial order induced by dominance are extended by the total order given by listing the partitions of n in reverse lexicographic order.

2.2 Symmetric Functions

A function $f(x_1, x_2, \dots) \in \mathbb{R}[[x_1, x_2, \dots]]^1$ is *symmetric* if $f(x_1, x_2, \dots) = f(x_{\sigma(1)}, x_{\sigma(2)}, \dots)$ for every permutation σ of the positive integers $\mathbb{N}$. The *algebra of symmetric functions* Λ is the subalgebra of $\mathbb{R}[[x_1, x_2, \dots]]$ consisting of those symmetric functions f that are of bounded degree (that is, there exists a positive integer n such that every monomial of f has degree $\leq n$). Furthermore, Λ is a graded algebra, with natural grading

$$\Lambda = \bigoplus_{k=0}^{\infty} \Lambda^d$$

where Λ^d consists of symmetric functions that are homogeneous of degree d [27, 40].

Partitions arise in considering Λ because they naturally index its bases. One way to see this is to consider the simplest symmetric functions, the *monomial symmetric functions*, defined for an integer partition λ as

$$m_\lambda = \sum_{\sigma \in S_\mathbb{N}} x_{\sigma(1)}^{\lambda_1} x_{\sigma(2)}^{\lambda_2} \cdots x_{\sigma(k)}^{\lambda_k}.$$

It is easy to verify that for each d, every function $f \in \Lambda^d$ may be written as a unique linear combination of elements of $\{m_\lambda : \lambda \vdash d\}$, so the monomial symmetric functions of degree d form a basis of Λ^d as a vector space, and thus the dimension of Λ^d is equal to the number of partitions of d (and Λ is infinite-dimensional).

---

[1]The choice of $\mathbb{R}$ as the coefficient ring is not particularly important for this work. Choosing $\mathbb{C}$ or even $\mathbb{Q}$ would work just as well, and in fact most (although not quite all) of the elementary results in this section also work over $\mathbb{Z}$, provided that "vector space" is replaced by "module", etc.

There are five commonly used symmetric function bases, all indexed by integer partitions λ (and discussed in further detail that will be given here in [27, 40]). The monomial symmetric functions m_λ are one. A second basis consists of the *elementary symmetric functions*, defined by the equations

$$e_n = \sum_{i_1 < \cdots < i_n} x_{i_1} \ldots x_{i_n}, \quad e_\lambda = e_{\lambda_1} e_{\lambda_2} \ldots e_{\lambda_k}.$$

These functions arise when considering polynomials (considered in one variable x that is distinct from the x_i). If $(x - r_1) \ldots (x - r_k)$ is a monic polynomial of degree k, then its expands as

$$\sum_{i=0}^{k} (-1)^{k-i} e_i(r_1, \ldots, r_k, 0, 0, \ldots) x^i.$$

Let $\lambda \vdash d$, and let M be the matrix with rows and columns indexed by partitions of d such that

$$e_\lambda = \sum_{\mu \vdash n} M_{\lambda\mu} m_\mu.$$

Then it is easy to verify that $M_{\lambda\mu} \geq 0$, and $M_{\lambda\mu} > 0$ if and only if $\mu \leq \lambda'$ [40]. Furthermore, clearly $M_{\lambda\lambda'} = 1$, so after an appropriate rearrangement of its rows and columns M is an upper triangular matrix with 1s on the main diagonal and is thus invertible (even over $\mathbb{Z}$), verifying that $\{e_\lambda : \lambda \vdash d\}$ is indeed a basis of Λ^d.

A third basis is the *complete homogeneous symmetric functions*, defined by the equations

$$h_n = \sum_{i_1 \leq \cdots \leq i_n} x_{i_1} \ldots x_{i_n}, \quad h_\lambda = h_{\lambda_1} h_{\lambda_2} \ldots h_{\lambda_k}.$$

It may be shown that for all $n \geq 1$ the identity

$$\sum_{i=0}^{n}(-1)^i h_i e_{n-i} = 0$$

holds, and as a corollary that the endomorphism $\omega : \Lambda \to \Lambda$ defined by $\omega(e_\lambda) = h_\lambda$ extended linearly is an involution (that is, it also satisfies $\omega(h_\lambda) = e_\lambda$) [40]. Thus we see that $\{h_\lambda : \lambda \vdash d\}$ is a basis of Λ^d.

A fourth basis is the *power-sum symmetric functions*, defined by the equations

$$p_n = \sum_{i=1}^{\infty} x_i^n, \quad p_\lambda = p_{\lambda_1} p_{\lambda_2} \cdots p_{\lambda_k}.$$

It is straightforward to verify that if $\lambda \vdash d$ and $R_{\lambda\mu}$ is the matrix indexed by partitions of d such that $p_\lambda = \sum_{\mu \vdash d} R_{\lambda\mu} m_\mu$ then

$$R_{\lambda\mu} = \sum_{\text{puzzles } \lambda \to \mu} \frac{\prod_i r_i(\lambda)!}{\prod_{i,j} r_i(\lambda^j)!} \tag{2.1}$$

This verifies that the set $\{p_\lambda : \lambda \vdash d\}$ is a basis of Λ^d since $R_{\lambda\mu} = 0$ if λ is not a refinement of μ, and $R_{\lambda\lambda} > 0$ [2](in fact it is true that $R_{\lambda\mu} > 0$ if and only if $\lambda \leq \mu$ [40]). We also note that the action of the symemtric function involution ω on the p-basis is $\omega(p_\lambda) = (-1)^{|\lambda|-l(\lambda)} p_\lambda$ [40].

The last commonly used basis of symmetric functions is the basis of *Schur functions* s_λ. These are defined using the Young diagram of the corresponding partition. A *filling* of a Young diagram is an assignment of a positive integer to

---

[2]However, $R_{\lambda\lambda}$ is not equal to 1 for all λ, so the matrix R is *not* invertible over $\mathbb{Z}$, and accordingly the set $\{p_\lambda : \lambda \vdash d\}$ is not generally a $\mathbb{Z}$-basis for Λ^d.

each box. A filled Young diagram is called a *Young tableau.* A Young tableau is called *semi-standard* if the positive integers are:

- Weakly increasing ($\leq$) along rows, and

- Strictly increasing ($<$) down columns.[3]

For example, the following is a semi-standard Young tableau for $(4, 2, 2, 1)$:

$$
\begin{array}{|c|c|c|c|}
\hline
1 & 1 & 2 & 4 \\
\hline
2 & 3 \\
\cline{1-2}
4 & 4 \\
\cline{1-2}
6 \\
\cline{1-1}
\end{array}
$$

Let $SSYT(\lambda)$ denote the set of semi-standard Young tableau of shape λ. If $r_i(T)$ is the number of occurences of the number i in T, define the *content of T* as the partition $1^{r_1(T)}2^{r_2(T)}\ldots$ given by part multiplicities. Let $K_{\lambda\mu}$ be the number of $T \in SSYT(\lambda)$ with content μ. Then for $\lambda \vdash d$ the Schur functions s_λ may be defined by

$$
s_\lambda = \sum_{\mu \vdash d} K_{\lambda\mu} m_\mu.
$$

It may be shown that $K_{\lambda\mu} > 0$ if and only if $\mu \leq \lambda$ [13]. Since also clearly $K_{\lambda\lambda} = 1$ for all λ, the set $\{s_\lambda : \lambda \vdash d\}$ is a basis for Λ^d. It is easy to verify

---

[3]As one might guess, a Young tableau is called *standard* if the integers are strictly increasing along both rows and columns, and typically in this context it is also required that the entries are the numbers 1 through n each exactly once. Standard Young tableaux are important in defining the modules corresponding to irreducible representations of the symmetric group, but this topic is outside the scope of this work. For more information see [36].

that $s_n = h_n$ and $s_{1^n} = e_n$, so the Schur functions may be viewed as "interpolating" between these. There are numerous other equivalent ways to define Schur functions, e.g. using the identity

$$s_\lambda = \det\left[(h_{\lambda_i + j - i})|_{i,j=1}^{l(\lambda)}\right]$$

We also note that the symmetric function involution ω satisfies $\omega(s_\lambda) = s_{\lambda'}$. For more exposition see [27, 40].

The importance of the Schur functions lies in their connection to the representation theory of the symmetric group. Given a group G, a *class function* f on G is any function $f : G \to \mathbb{C}$ such that whenever $g, h \in G$ are conjugate, $f(g) = f(h)$. The set of class functions of G forms a $\mathbb{C}$-vector space, and its dimension is equal to the number of conjugacy classes of G, since for example one of its bases consists of those functions that take value 1 on some fixed conjugacy class and 0 on the others.

In the specific case of S_n, every element $w \in S_n$ has an associated partition of n called its *cycle type*, given by the sizes of the cycles in the cycle decomposition of w. It is well-known that two elements of S_n are conjugate if and only if they have the same cycle type, so there are as many conjugacy classes of S_n as there are partitions of n [40]. It is also well-known that there is a natural way irreducible characters of S_n may be indexed as $\{\chi^\lambda : \lambda \vdash n\}$, such that these irreducible characters are an orthonormal basis of the set F of class functions of S_n with respect to the inner product

$$\langle f, g \rangle = \frac{1}{n!} \sum_{w \in S_n} f(w)\overline{g(w)}$$

11

The ring of symmetric functions Λ also has an inner product, defined by the relation $\langle s_\lambda, s_\mu \rangle = \delta_{\lambda\mu}$ extended linearly. It may be shown that the map $ch : F \to \Lambda$ defined by extending linearly the relation $ch(\chi^\lambda) = s_\lambda$ is an isometry and bijective graded ring homomorphism of F with Λ, and it is also a graded ring isomorphism if the coefficient rings of Λ and F are identical [40]. This map is called the *Frobenius characteristic*, and it sheds light on the role of the Schur functions in providing a direct link between symmetric function theory and the representation theory of S_n.

Given a symmetric function f and a basis b of Λ, we say that f is b-*positive* if when we write f in the basis b, all coefficients are nonnegative. A common theme of research in symmetric function theory arises when a certain family of symmetric functions seems to be positive with respect to some basis other than the m-basis.

If a family of functions is s-positive, the Frobenius characteristic implies that these functions are characters of some S_n-module, and thus the coefficients may be coming from some deep algebraic connection. If a family is e-positive or h-positive, then the function represents a positive algebraic combination of sign representations or trivial representations (respectively) of permutation groups. Additionally, if $\lambda \vdash n$, and $f_\lambda \in F$ is the class function of S_n such that $f_\lambda(w) = 1$ if w has cycle type λ and 0 otherwise, then the Frobenius characteristic also satisfies

$$ch(f_\lambda) = \frac{1}{n!} \left(\prod_i i^{r_i(\lambda)} r_i(\lambda)! \right) p_\lambda$$

and thus p-positive symmetric function families represent a positive linear combination of such canonical characters [40].

2.3 Graphs

A *graph $G = (V, E)$* consists of a *vertex set V* and an *edge multiset E* where the elements of E are pairs of (not necessarily distinct) elements of V. Graphs are typically depicted with points representing vertices, and line segments (or occasionally curves) between two vertices representing edges. An edge $e \in E$ that contains the same vertex twice is called a *loop*. If there are two or more edges that each contain the same two vertices, they are called *multi-edges*. A *simple graph* is a graph $G = (V, E)$ in which E does not contain loops or multi-edges (thus, $E \subseteq \binom{V}{2}$). If $\{v_1, v_2\}$ is an edge (or nonedge), we will write it as $v_1 v_2 = v_2 v_1$. The vertices v_1 and v_2 are the *endpoints* of the edge $v_1 v_2$. We will use $V(G)$ and $E(G)$ to denote the vertex set and edge multiset of a graph G respectively.

Two graphs G and H are said to be *isomorphic* if there exists a map $f : V(G) \to V(H)$ such that for all $v_1, v_2 \in V(G)$ (not necessarily distinct), the number of edges $v_1 v_2$ in $E(G)$ is the same as the number of edges $f(v_1) f(v_2)$ in $E(H)$.

The *degree of a vertex v* in a graph G, denoted $d(v)$, is the number of edges of G having v as an endpoint (where loops are counted twice). The *degree sequence of a graph G* is the tuple $(d_1, d_2, \ldots, d_{|V(G)|})$ where $d_1 \geq \cdots \geq d_{|V(G)|}$ are the degrees of the vertices of G.

The *complement* of a simple graph $G = (V, E)$ is denoted $\overline{G}$, and is defined as $\overline{G} = (V, \binom{V}{2} \setminus E)$, so in $\overline{G}$ every edge of G is replaced by a nonedge, and every nonedge is replaced by an edge.

A *subgraph* of a graph $G = (V, E)$ is a graph $G' = (V', E')$ where $V' \subseteq V$ and $E' \subseteq E|_{V'}$, where $E|_{V'}$ is the set of edges of G with both endpoints in V'. An *induced subgraph* of G is a graph $G' = (V', E|_{V'})$ with $V' \subseteq V$, and we will write $G|_{V'}$ to mean the induced subgraph of G with vertex set V'. Given graphs G and H, we say that G is *H-free* if G has no induced subgraph isomorphic to H. A *stable set* of G is a subset $V' \subseteq V$ such that $E|_{V'} = \emptyset$. A *clique* of G is a subset $V' \subseteq V$ such that for every pair of distinct vertices v_1 and v_2 of V', $v_1 v_2 \in E(G)$.

A *path* in a graph G is a sequence of edges $v_1 v_2$, $v_2 v_3$, ..., $v_{k-1} v_k$ such that $v_i \neq v_j$ for all $i \neq j$. The vertices v_1 and v_k are the *endpoints* of the path. A *cycle* in a graph is a sequence of edges $v_1 v_2$, $v_2 v_3$,..., $v_k v_1$ such that $v_i \neq v_j$ for all $i \neq j$. Note that in a simple graph every cycle must have at least 3 edges, although in a nonsimple graph there may be cycles of size 1 (a loop) or 2 (multi-edges).

A graph G is *connected* if for every pair of distinct vertices v_1 and v_2 of G there is a path in G with v_1 and v_2 as its endpoints. The *connected components* of G are the maximal induced subgraphs of G which are connected.

A *tree* is a connected graph that does not contain any cycles. Every tree with n vertices must have exactly $n - 1$ edges, since if we start with n vertices and no edges and add edges one at a time, each new edge decreases the number of connected components by one unless it creates a cycle. Thus, to go from n connected components to 1 without adding cycles uses exactly $n - 1$ edges. A vertex of degree 1 in a tree is called a *leaf.*

The *complete graph* K_n on n vertices is the unique simple graph having all possible edges, that is, $E(K_n) = \binom{V}{2}$.

Given a graph G, there are two commonly used operations that produce new graphs. One is *deletion*: given an edge $e \in E(G)$, the graph of G *with e deleted* is the graph $G' = (V(G), E(G)\backslash\{e\})$, and is denoted $G\backslash e$. Likewise, if S is a multiset of edges, we use $G\backslash S$ to denote the graph $(V(G), E(G)\backslash S)$.

The other operation is the *contraction of an edge* $e = v_1v_2$, denoted G/e. If $v_1 = v_2$ (e is a loop), we define $G/e = G\backslash e$. Otherwise, we create a new vertex v^*, and define G/e as the graph G' with $V(G') = (V(G)\backslash\{v_1, v_2\}) \cup v^*$, and $E(G') = (E(G)\backslash E(v_1, v_2)) \cup E(v^*)$, where $E(v_1, v_2)$ is the set of edges with at least one of v_1 or v_2 as an endpoint, and $E(v^*)$ consists of each edge in $E(v_1, v_2)\backslash e$ with the endpoint v_1 and/or v_2 replaced with the new vertex v^*. Note that this is an operation on a (possibly nonsimple) graph that identifies two vertices while keeping and/or creating multi-edges and loops.

There is also a different version of edge contraction that is defined only on simple graphs. In the case that G is a simple graph, we define the *simple contraction* $G \nmid e$ to be the same as G/e except that after performing the contraction operation, we delete any loops and all but a single copy of each multi-edge so that the result is again a simple graph.

2.4   Coloring Graphs

Let $G = (V(G), E(G))$ be a (not necessarily simple) graph. A map $\kappa : V(G) \to \mathbb{Z}^+$ is called a *coloring* of G. It is called an *n-coloring* of G if the image is contained in $\{1, 2, \ldots, n\}$. A coloring is called *proper* if $\kappa(v_1) \neq \kappa(v_2)$ for all v_1, v_2 such that there exists an edge $e = v_1 v_2$ in G.

Let $\chi_G(n) : \mathbb{Z}^+ \to \mathbb{Z}^+$ be a function on G defined on positive integers by letting $\chi_G(n)$ be the number of proper n-colorings of G. Then the following *deletion-contraction relation* holds:

Lemma 1. *For a graph G, and an edge $e \in E(G)$,*

$$\chi_G(n) = \chi_{G \backslash e}(n) - \chi_{G/e}(n)$$

Proof. First, note that if e is a loop, then the statement follows immediately since $G/e = G \backslash e$ and $\chi_G(n) = 0$ by definition. If e is one of multiple edges connecting the same two vertices, then the statement also follows immediately since then $\chi_G(n) = \chi_{G \backslash e}(n)$, and $\chi_{G/e}(n) = 0$ since the contraction forms a loop.

Thus we may assume that e is the only edge connecting two distinct vertices, which we will call v_1 and v_2.

We show the rearranged formula

$$\chi_G(n) + \chi_{G/e}(n) = \chi_{G \backslash e}(n)$$

It suffices to show a one-to-one correspondence of proper n-colorings κ of $G\backslash e$, and proper n-colorings of exactly one of G and G/e. We split into cases based on κ. If $\kappa(v_1) \neq \kappa(v_2)$, then the same κ is a proper n-coloring of G. If $\kappa(v_1) = \kappa(v_2) = m$, then the same κ is a proper n-coloring of G/e if we assign $\kappa(v^*) = m$ for the vertex v^* formed by contracting v_1 and v_2. Conversely, every proper n-coloring of G corresponds to a proper n-coloring of $G\backslash e$ with $\kappa(v_1) \neq \kappa(v_2)$, and every proper n-coloring of G/e corresponds to a proper n-coloring of $G\backslash e$ with $\kappa(v_1) = \kappa(v_2)$, and this concludes the proof.

$\square$

When G is a graph with k vertices and no edges it is easy to compute that $\chi_G(n) = n^k$, so we may use the deletion-contraction relation and induction on the number of edges to conclude that $\chi_G(n)$ is always a polynomial in n for any graph G. Hence χ_G is called the *chromatic polynomial of G*. Furthermore, we may extend χ_G to have domain and range $\mathbb{R}$ by defining $\chi_G(x)$ to be the unique real-valued polynomial satisfying that $\chi_G(n)$ is the number of proper n-colorings of G for every $n \in \mathbb{Z}^+$, and this is how we will view χ_G from now on.

As was described in Section 1, the *chromatic symmetric function X_G of G* is defined as

$$X_G(x_1, x_2, ...) = \sum_{\kappa} \prod_{v \in V(G)} x_{\kappa(v)}$$

where the sum runs over all proper colorings κ of G. Clearly X_G is symmetric since for any summand $\prod_{v \in V(G)} x_{\kappa(v)}$, we get all of its permutations by simply permuting the colors of κ. Note that if G contains a loop then $X_G = 0$, and X_G is unchanged by replacing any multi-edges by a single edge.

A thorough overview of X_G is given in [38], including interpretations of coefficients of X_G when expressed in some of the symmetric function bases given in Section 2.2; we postpone introducing these properties to the next chapter, where we will also prove generalizations of them. We also postpone discussion of recent developments in the literature until the relevant properties are discussed.

Chapter 3

Deletion-Contraction Relation

In this chapter we introduce *vertex-weighted graphs* (G, w), and extend the chromatic symmetric function X_G to include such graphs. We will demonstrate the utility of this new function $X_{(G,w)}$ by proving new properties of and rederiving known results for X_G by proving them for the more general $X_{(G,w)}$. In the case of previously known results these proofs are substantially different in nature from the original ones, as they depend primarily on simple enumerative techniques and induction using deletion-contraction.

We now outline how this chapter is organized. Section 3.1 defines vertex-weighted graphs (G, w) and the extended chromatic symmetric function $X_{(G,w)}$. We demonstrate that $X_{(G,w)}$ satisfies a deletion-contraction relation akin to that of the chromatic polynomial, and we use this relation to prove a number of generalizations of properties of X_G. In Section 3.2 we prove an interpretation of the

coefficients of $X_{(G,w)}$ in the e-basis involving acyclic orientations and maps defined on the sinks of such orientations. This interpretation generalizes a well-known result of Stanley [38], and thus also proves Stanley's result in a new way. In Section 3.3 we discuss possible applications of $X_{(G,w)}$ to research. In particular, we define an extension of $X_{(G,w)}$ that generalizes the chromatic quasisymmetric function of Shareshian and Wachs [42] and show that it satisfies a deletion-contraction relation, and we show that $X_{(G,w)}$ is neither e-positive nor e-negative when w is not the weight function that gives every vertex weight 1. We also discuss potential applications of $X_{(G,w)}$ to s-positivity, partition systems, the umbral chromatic polynomial, the path-cycle symmetric function for digraphs, functions on double posets, and tree isomorphism. Finally, in Section 3.4, we also generalize Stanley's *bad-coloring extension* XB_G of the chromatic symmetric function to vertex-weighted graphs, demonstrate a deletion-contraction relation in this setting, and show that the generalized $XB_{(G,w)}$ is closely related to the W-polynomial [31] and the (r,q)-chromatic function [24].

Material in Sections 3.1, 3.2, and 3.3 is taken from the author's paper [7], written with Sophie Spirkl. Material in Section 3.4 will be featured in [9] with Sophie Spirkl, currently in preparation.

3.1 Extending X_G to Vertex-Weighted Graphs

Define a *vertex-weighted graph* (G, w) to be a graph G together with a vertex-weight function $w : V(G) \to \mathbb{Z}^+$. The *weight of a vertex $v \in V(G)$* is $w(v)$. Using the notation

$$x_\kappa(G, w) = \prod_{v \in V(G)} x_{\kappa(v)}^{w(v)}$$

we generalize the chromatic symmetric function to vertex-weighted graphs as

$$X_{(G,w)} = \sum_\kappa x_\kappa(G, w) \tag{3.1}$$

where the sum is taken over all proper colorings κ of G. We use this nonstandard notation as it will be convenient to refer explicitly to individual summands of $X_{(G,w)}$ in proofs. Note that the usual chromatic symmetric function X_G is equivalent to $X_{(G,w)}$ where w is the function assigning weight 1 to each vertex.

Given a vertex-weighted graph (G, w) and $A \subseteq V(G)$, define the *total weight of A*, denoted $w(A)$, to be $\sum_{v \in A} w(v)$. Define the *total weight of G* to be the total weight of $V(G)$. Throughout this thesis, when G is clear we will generally use n to denote the number of vertices of G, and d to denote the total weight of G.

Let $\lambda = (\lambda_1, \dots, \lambda_k)$ be a partition. Define $St_\lambda(G, w)$ to be the set of (unordered) partitions of $V(G)$ into $k = l(\lambda)$ stable sets whose total weights are $\lambda_1, \dots, \lambda_k$. We begin by establishing a simple formula for expanding $X_{(G,w)}$ in the monomial basis:

Lemma 2. *If (G, w) is a vertex-weighted graph with n vertices and total weight d, then*

$$X_{(G,w)} = \sum_{\lambda \vdash d} |St_\lambda(G, w)| \left(\prod_{i=1}^{d} r_i(\lambda)! \right) m_\lambda \qquad (3.2)$$

where we recall that $r_i(\lambda)$ is the number of parts of λ equal to i.

Proof. The proof is a simple modification of the proof of ([38], Theorem 2.4). Since $X_{(G,w)}$ is symmetric, it suffices to show that the coefficient of $x_1^{\lambda_1} \ldots x_k^{\lambda_k}$ is correct. For every element of $St_\lambda(G, w)$, label the stable sets $L_1, \ldots, L_k$ in some order such that $|L_i| = \lambda_i$. Then there are $\left(\prod_{i=1}^{d} r_i(\lambda)! \right)$ corresponding proper colorings κ of (G, w) such that $\forall i \exists j$ with $\kappa^{-1}(j) = L_i$ and also $x_\kappa(G, w) = x_1^{\lambda_1} \ldots x_k^{\lambda_k}$, since one such coloring is $\kappa(L_i) = i$ for $1 \leq i \leq k$, and we may also permute the colors among those L_i that have the same cardinality. Since also clearly every proper κ with $x_\kappa(G, w) = x_1^{\lambda_1} \ldots x_k^{\lambda_k}$ has a corresponding element of $St_\lambda(G, w)$ for which κ is monochromatic on each part, the terms of $X_{(G,w)}$ are in one-to-one correspondence with those of the right-hand side of (3.2), so the lemma is proved.

$\square$

As an example, for a partition $\lambda \vdash d$ with $l(\lambda) = n$, we define the graph $K^\lambda = (K_{l(\lambda)}, w)$ where $w(v_i) = \lambda_i$ for some ordering $v_1, v_2, \ldots, v_n$ of the vertices. Since the only stable sets of $K_{l(\lambda)}$ are single vertices, every coloring of $K_{l(\lambda)}$ colors every vertex with a distinct color, and so only monomials of m_λ appear. Each monomial

of m_λ will occur once for each permutation of the colors of vertices with the same weights, and so we have

$$X_{K^\lambda} = \left(\prod_{i=1}^{\infty} r_i(\lambda)! \right) m_\lambda.$$

Analogously, we define $\overline{K^\lambda} = (\overline{K_{l(\lambda)}}, w)$ where $w(v_i) = \lambda_i$ for some ordering of the vertices. Note that the chromatic symmetric function is multiplicative across disjoint unions, since we may color each of the connected components independently. Since $X_{K^n} = p_n$, it follows that

$$X_{\overline{K^\lambda}} = p_\lambda.$$

Note that in the case of unweighted graphs, there is no G such that X_G is equal to a nonzero multiple of m_λ or p_λ except in the case that $\lambda = 1^n$ [5].[4]

3.1.1 A Deletion-Contraction Relation

One of the primary motivations for extending the chromatic symmetric function to vertex-weighted graphs is the existence of a deletion-contraction relation in this setting. Given a vertex-weighted graph (G, w), and an edge $e = v_1 v_2$ of G, let w/e be the modified weight function on G/e such that $w/e = w$ if e is a loop, and

---

[4]It is natural to ask which bases are "representable" as chromatic symmetric functions in this way. In addition to the m- and p-bases, we obtain $X_G = (\prod \lambda_i!) e_\lambda$ even in the unweighted case by taking G to be a disjoint union of cliques of sizes equal to the parts of λ. Furthermore, it is easy to show by combining the p-positivity of (3.12) with the results of [5] that there are no vertex-weighted graphs (G, w) such that $X_{(G,w)}$ is a nonzero multiple of h_λ or s_λ except in the case that $\lambda = 1^n$.

otherwise $(w/e)(v) = w(v)$ if $v \neq v_1, v_2$, and for the vertex v^* of G/e formed by the contraction, $(w/e)(v^*) = w(v_1) + w(v_2)$. Note that the same definition of w/e may be applied to the simple contraction $G \dagger e$, so we use the same notation.

Lemma 3. *Let (G, w) be a vertex-weighted graph, and let $e \in E(G)$ be any edge. Then*

$$X_{(G,w)} = X_{(G \backslash e, w)} - X_{(G/e, w/e)} \tag{3.3}$$

and if G is a simple graph,

$$X_{(G,w)} = X_{(G \backslash e, w)} - X_{(G \dagger e, w/e)}.$$

Proof. First, we note that for a simple graph G, $X_{(G \dagger e, w/e)} = X_{(G/e, w/e)}$. This is because the only case in which $G \dagger e$ is different from G/e is in the case that some vertex v' had edges to both endpoints of e, for then G/e would have a multi-edge where $G \dagger e$ has a single edge. But by Lemma 1 multi-edges may be reduced to a single edge without affecting the chromatic symmetric function, establishing the claim. Thus, it suffices to prove (3.3).

We rewrite (3.3) in the form

$$X_{(G \backslash e, w)} = X_{(G,w)} + X_{(G/e, w/e)}. \tag{3.4}$$

The statement is immediate if e is a loop, so we may assume $e = v_1 v_2$ connects distinct vertices, and we also let v^* be the contracted vertex in G/e. It suffices to show a one-to-one correspondence between terms of $X_{(G \backslash e, w)}$ and terms of $X_{(G,w)}$ or $X_{(G/e, w/e)}$. We consider two cases for each term $x_\kappa(G \backslash e, w)$ (as defined in Section 2)

24

occurring in the left-hand side of (3.4) based on the proper coloring κ. If $\kappa(v_1) = \kappa(v_2)$, then $x_\kappa(G\backslash e, w) = x_{\kappa_e}(G/e, w/e)$, where κ_e is the proper coloring of G/e such that $\kappa_e(v^*) = \kappa(v_1)$, and for all other vertices v, $\kappa_e(v) = \kappa(v)$. If $\kappa(v_1) \neq \kappa(v_2)$, then $x_\kappa(G\backslash e, w) = x_\kappa(G, w)$. This correspondence is injective, since changing the color of any vertex in $G\backslash e$ changes the corresponding proper coloring of either G or G/e. This correspondence is also surjective, since given a proper coloring of G or G/e, we can recover a proper coloring of $G\backslash e$ that is its preimage under this map by removing e or uncontracting v^*, respectively.

$\square$

Note that it is also possible to write this relation in the "vertex uncontraction" form

$$X_{(G/e, w/e)} = X_{(G\backslash e, w)} - X_{(G,w)}. \tag{3.5}$$

This form has increased flexibility, because if we are given $(G/e, w/e)$, we may make two choices in uncontracting: first, if the vertex being uncontracted has weight greater than 2, we may choose how to distribute the weights to the two new vertices in G, and second, for edges that were incident to the contracted vertex, we may choose how those edges are incident to the newly created vertices in G. Thus, whenever this uncontraction form is used on a graph $(G/e, w/e)$ throughout this paper, we will specify the graph (G, w).

One advantage of having a deletion-contraction relation is that to prove a property on graphs, we can pass to an appropriate property on vertex-weighted graphs,

and either use the deletion-contraction property directly, or an inductive approach by showing that the property holds on graphs with no edges, and applying induction to the number of edges using deletion-contraction.

To illustrate the power of this approach, we extend known properties of the chromatic symmetric function on unweighted graphs to the set of vertex-weighted graphs. In doing so we provide new, alternate proofs of these properties in the unweighted case.

3.1.2 p-Basis Expansion Formula

Given a vertex-weighted graph (G, w), and $S \subseteq E(G)$, we define $\lambda(G, w, S)$ to be the partition whose parts are the total weights of the connected components of (G', w), where $G' = (V(G), S)$.

Lemma 4.

$$X_{(G,w)} = \sum_{S \subseteq E(G)} (-1)^{|S|} p_{\lambda(G,w,S)} \tag{3.6}$$

Proof. This could be proved by adapting the proof of ([38], Theorem 2.5), but we give a different proof using deletion-contraction. In our vertex-weighted graph (G, w), let $e_1, e_2, \ldots, e_m$ be an ordering of the edges. We expand $X_{(G,w)}$ in the following manner: First, in step 1 we apply deletion-contraction to e_1, and get

$$X_{(G,w)} = X_{(G \setminus e_1, w)} - X_{(G/e_1, w/e_1)}.$$

26

Then, in step 2, we apply deletion-contraction to both of $(G\backslash e_1, e)$ and $(G/e_1, w/e_1)$ using edge e_2, and obtain an equation with four terms. Continuing in this manner, in step i we apply deletion-contraction to all 2^{i-1} summands created in the previous step, until after step m we have an equation of the form

$$X_{(G,w)} = \sum_{S \subseteq E(G)} (-1)^{|S|} X_{(G(S),w(S))}$$

where $(G(S), w(S))$ is the graph resulting from contracting the edges in S and deleting the edges in $E(G)\backslash S$, using our given ordering. This graph has no edges, and each vertex corresponds to a connected component of the graph $G' = (V(G), S)$, since the vertices have been formed by the contraction of exactly those edges in S. Furthermore, the weights of these vertices are the total weights of the connected components of (G', w), since the weight of a vertex in $(G(S), w(S))$ is the sum of the weights of all the vertices in the corresponding component of (G', w). We recall that if $\overline{K^\lambda}$ is the graph with no edges and vertices of weights $\lambda_1 \geq \cdots \geq \lambda_k$, then $X_{\overline{K^\lambda}} = p_\lambda$. Thus $X_{(G(S),w(S))} = p_{\lambda(G,w,S)}$, and the result follows.

$\square$

3.1.3 The Effect of the Symmetric Function Involution

Given a graph G, define an *orientation of G* to be an assignment of an order (or orientation) to the endpoints of each edge $e \in E(G)$. If we orient the edge $v_1 v_2$ by placing v_1 before v_2, we write $v_1 \rightarrow v_2$, and say that v_1 is the *tail*, and v_2 the *head*. An *oriented cycle* of an orientation of G is a sequence of edges

$v_1v_2, v_2v_3, \ldots, v_{m-1}v_0, v_0v_1$ that forms a cycle in G, and such that these edges are either all oriented such that $v_i \to v_{i+1}$ for all i, or $v_{i+1} \to v_i$ for all i (with indices taken mod m in both cases). An *acyclic orientation* of G is one which contains no oriented cycle.

Recall that with respect to an edge $e = v_1v_2 \in E(G)$ that is not a loop, we define the contracted graph G/e to be G' with $V(G') = (V(G)\backslash\{v_1, v_2\}) \cup v^*$, and $E(G') = (E(G)\backslash E(v_1, v_2)) \cup E(v^*)$, where $E(v_1, v_2)$ is the set of edges with at least one of v_1 or v_2 as an endpoint, and $E(v^*)$ consists of each edge in $E(v_1, v_2)\backslash\{e\}$ with the endpoint v_1 and/or v_2 replaced with the new vertex v^*. Using this notation, we define the *contraction of the orientation γ with respect to e* to be the orientation γ_e of G/e where

- Edges of $(E(G)\backslash E(v_1, v_2))$ are oriented as they are in γ.

- If $vv_i, i \in \{1, 2\}$ is an edge of G for $v \neq v_1, v_2$, orient the corresponding edge vv^* of G/e such that v is a head of vv_i in γ if and only if v is a head of vv^* in γ_e.

- If $v_iv_j, i, j \in \{1, 2\}$ is an edge of G other than e, the corresponding edge v^*v^* is oriented trivially as $v^* \to v^*$.

Given a vertex-weighted graph (G, w) and an acyclic orientation γ of G, we define a coloring κ of G to be *weakly proper with respect to γ* if for every edge $e = v_1v_2$ oriented as $v_1 \to v_2$ by γ, we have $\kappa(v_1) \leq \kappa(v_2)$, and in this case we write

$$\overline{x}_\kappa(G, w, \gamma) = \prod_{v \in V(G)} x_{\kappa(v)}^{w(v)}.$$

For a vertex-weighted graph (G, w), we define its *weak chromatic symmetric function* as

$$\overline{X}_{(G,w)} = \sum_{(\gamma,\kappa)} \overline{x}_\kappa(G, w, \gamma) \tag{3.7}$$

where the sum ranges over all ordered pairs (γ, κ) with γ an acyclic orientation of G, and κ a weakly proper coloring of G with respect to γ.

We prove the following formula for the vertex-weighted weak chromatic symmetric function, extending the formula for unweighted graphs given in ([38], Theorem 4.2):

Theorem 5. *Let (G, w) be a vertex-weighted graph with n vertices and with total weight d. Then*

$$\overline{X}_{(G,w)} = (-1)^{d-n} \omega(X_{(G,w)}) \tag{3.8}$$

where we recall that ω is the involution on symmetric functions satisfying $\omega(p_\lambda) = (-1)^{|\lambda|-l(\lambda)} p_\lambda$.

Proof. We proceed by induction on the number of edges of G. In the base case, the graph has no edges, and vertices of weights $\lambda_1 \geq \cdots \geq \lambda_k$, say. Then $\overline{X}_{(G,w)} = X_{(G,w)} = p_\lambda$, where $\lambda = (\lambda_1, \ldots, \lambda_k)$, and since $\omega(p_\lambda) = (-1)^{|\lambda|-l(\lambda)} p_\lambda = (-1)^{d-n} p_\lambda$, the result follows.

For the inductive step, we consider (G, w) where G has $m \geq 1$ edges, and assume that (3.8) holds for graphs with $m - 1$ or fewer edges. Let $e = v_1 v_2$ be an edge of G. Then from the deletion-contraction relation (3.3), we deduce that

$$(-1)^{d-n} \omega(X_{(G,w)}) = (-1)^{d-n} \omega(X_{(G \backslash e, w)}) + (-1)^{d-n-1} \omega(X_{(G/e, w/e)}). \qquad (3.9)$$

By applying the inductive hypothesis to $(G \backslash e, w)$ and $(G/e, w/e)$, it suffices to show that

$$\overline{X}_{(G,w)} = \overline{X}_{(G \backslash e, w)} + \overline{X}_{(G/e, w/e)}. \qquad (3.10)$$

We extend the definition of $\overline{x}_\kappa(G, w, \gamma)$ to include all orientations γ and all colorings κ by defining that $\overline{x}_\kappa(G, w, \gamma) = 0$ if γ is not acyclic, or if κ is not a weakly proper coloring of G with respect to γ. Given an orientation γ and coloring κ on $(G \backslash e, w)$, we also define the following:

- γ_1 is the orientation of (G, w) with $v_1 \to v_2$ and all other edges oriented as in γ.

- γ_2 is the orientation of (G, w) with $v_2 \to v_1$ and all other edges oriented as in γ.

- If $\kappa(v_1) = \kappa(v_2)$, κ_e is the coloring of $(G/e, w/e)$ with $\kappa_e(v^*) = \kappa(v_1)$ where v^* is the vertex created by the contraction of e, and for all other vertices v, $\kappa_e(v) = \kappa(v)$.

- If $\kappa(v_1) \neq \kappa(v_2)$, then κ_e does not exist (and so $\overline{x}_{\kappa_e}(G/e, w/e, \gamma_e) = 0$).

30

Using these definitions, to show (3.10) it suffices to show the stronger statement that for every acyclic orientation γ of $G\backslash e$, and every weakly proper coloring κ of $(G\backslash e, w)$ with respect to γ, we have

$$\overline{x}_\kappa(G, w, \gamma_1) + \overline{x}_\kappa(G, w, \gamma_2) = \overline{x}_\kappa(G\backslash e, w, \gamma) + \overline{x}_{\kappa_e}(G/e, w/e, \gamma_e) \qquad (3.11)$$

since every summand of $\overline{X}_{(G,w)}$, $\overline{X}_{(G/e,w/e)}$ and $\overline{X}_{(G\backslash e,w)}$ is counted exactly once in this way. Note that each of $\overline{x}_\kappa(G, w, \gamma_1)$, $\overline{x}_\kappa(G, w, \gamma_2)$, $\overline{x}_\kappa(G\backslash e, w, \gamma)$, and $\overline{x}_{\kappa_e}(G/e, w, \gamma_e)$ is either zero or equal to $\overline{x}_\kappa(G\backslash e, w, \gamma)$, so it is enough to show that the same number of summands on both sides of (3.11) are nonzero.

We split into cases based on whether γ has a directed path between v_1 and v_2 (note that it does not contain both a path from v_1 to v_2 and one from v_2 to v_1 since then γ would contain an oriented cycle). Suppose for a contradiction that γ contains such a path; without loss of generality we may assume it is from v_1 to v_2. Then γ_2 and γ_e both contain oriented cycles in their respective graphs. However, γ_1 does not contain an oriented cycle in (G, w). Furthermore, $\kappa(v_1) \leq \kappa(v_2)$ since κ is proper with respect to γ in $(G\backslash e, w)$ and there is a directed path from v_1 to v_2, so κ is proper with respect to γ_1 in (G, w). Thus, (3.11) holds in this case.

Now assume that there is no directed path. Then all of γ_1, γ_2, and γ_e are acyclic orientations. We split into subcases based on κ. If $\kappa(v_1) = \kappa(v_2)$, then κ_e exists and is proper with respect to γ_e, and κ is proper with respect to all of γ, γ_1, and γ_2, so (3.11) holds. Otherwise, without loss of generality suppose that $\kappa(v_1) < \kappa(v_2)$.

31

Then κ_e does not exist, and κ is not proper with respect to γ_2, but κ is proper with respect to γ_1, so (3.11) also holds in this case. This concludes the proof. $\square$

As a corollary, we deduce a further result about the function $\overline{X}_{(G,w)}$ that extends the corresponding result on unweighted graphs from ([38], Theorem 2.7):

Corollary 6. *If (G, w) is a vertex-weighted graph with n vertices and total weight d, then*

$$\overline{X}_{(G,w)} = (-1)^{d-n}\omega(X_{(G,w)}) \tag{3.12}$$

is p-positive.

Proof. We proceed by induction on the number of edges. The base case is a graph with no edges, and as was noted at the beginning of the previous proof, if such a graph (G, w) has vertices of weights $\lambda_1 \geq \cdots \geq \lambda_k$ say, then $\overline{X}_{(G,w)} = p_\lambda$ where $\lambda = (\lambda_1, \ldots, \lambda_k)$, and this is p-positive.

For the inductive step, suppose that (G, w) has $m \geq 1$ edges, and suppose that we have shown that the claim holds for vertex-weighted graphs (G, w) with $m - 1$ edges. Then for any edge $e \in E(G)$, using the inductive hypothesis and the relation (3.10) shows that $\overline{X}_{(G,w)}$ is a sum of two p-positive functions, and hence it is p-positive, and this concludes the proof.

$\square$

3.1.4 A Formula on Cycles

We now prove a modular relation on cycles that was originally proved for unweighted graphs by ([11], Proposition 5):

Theorem 7. *Let (G, w) be a vertex-weighted graph containing a cycle C, and let e be a fixed edge of this cycle. Then*

$$\sum_{S \subseteq E(C) \setminus e} (-1)^{|S|} X_{(G \setminus S, w)} = 0. \tag{3.13}$$

Proof. We proceed by induction on the number of edges in the cycle. The base case of a 1-edge cycle (a loop) is immediate.

For the inductive step, we assume the claim holds for graphs with an n-edge cycle and show that it holds on graphs with an $(n + 1)$-edge cycle. Let (G, w) be a vertex-weighted graph with an $(n + 1)$-edge cycle C, let e be the edge in the statement of Theorem 7, and let $f = v_1 v_2$ be an edge of the cycle with $e \neq f$. We apply deletion-contraction to the edge f to get

$$X_{(G,w)} = X_{(G \setminus f, w)} - X_{(G/f, w/f)}.$$

Let v^* be the vertex of G/f formed by the contraction of v_1 and v_2. We now apply the inductive hypothesis to $(G/f, w/f)$, since in this graph C/f is an n-edge cycle containing the edge e. We obtain

$$X_{(G,w)} = X_{(G \setminus f, w)} - \sum_{\emptyset \neq S' \subseteq E(C/f) \setminus e} (-1)^{|S'|-1} X_{((G/f) \setminus S', w/f)}.$$

Now, in this sum, for every summand we will uncontract the graph $((G/f) \setminus S', w/f)$ to $(G \setminus S', w)$, thus obtaining

33

$$X_{(G,w)} = X_{(G\setminus f,w)}$$

$$- \sum_{\emptyset \neq S' \subseteq E(C)\setminus\{e,f\}} (-1)^{|S'|-1} X_{(G\setminus(S'\cup\{f\}),w)}$$

$$+ \sum_{\emptyset \neq S' \subseteq E(C)\setminus\{e,f\}} (-1)^{|S'|-1} X_{(G\setminus S',w)}. \qquad (3.14)$$

We claim that the right-hand side of this equation is equal to

$$\sum_{\emptyset \neq S \subseteq E(C)\setminus e} (-1)^{|S|-1} X_{(G\setminus S,w)} \qquad (3.15)$$

which is sufficient to complete the proof.

The term $X_{(G\setminus f,w)}$ of (3.14) is equal to the term of (3.15) corresponding to $S = \{f\}$. The subtracted sum

$$- \sum_{\emptyset \subsetneq S' \subseteq E(C)\setminus\{e,f\}} (-1)^{|S'|-1} X_{(G\setminus(S'\cup\{f\}),w)}$$

in (3.14) is equal to the sum of those terms of (3.15) corresponding to sets $S = \{f\} \cup S'$ with $S' \neq \emptyset$. Finally, the sum

$$\sum_{\emptyset \subsetneq S' \subseteq E(C)\setminus\{e,f\}} (-1)^{|S'|-1} X_{(G\setminus S',w)}$$

of (3.14) is equal to the sum of the terms of (3.15) corresponding to sets $S = S'$ where S' is a nonempty subset of $C\setminus\{e, f\}$.

$\square$

3.2 Acyclic Orientations

Let $a(G)$ denote the number of acyclic orientations of a graph G. In terms of deletion-contraction, for any edge $e \in E(G)$ that is not a loop, it is easy to check that

$$a(G) = a(G\backslash e) + a(G/e) \tag{3.16}$$

It can be shown, either by using (3.16) and induction, or using a chromatic polynomial version of (3.10) as in [37], that if G is a graph on n vertices, then

$$a(G) = (-1)^n \chi_G(-1) \tag{3.17}$$

Additionally, if γ is an orientation of a graph G, we call a vertex $v \in V(G)$ a *sink* of γ if v is not the tail of any edge of γ. Then (3.17) is generalized by the following theorem:

Theorem 8. *([38], Theorem 3.3)*

Let G be an unweighted graph. We write its chromatic symmetric function in the elementary symmetric function basis as

$$X_G = \sum_\lambda c_\lambda e_\lambda.$$

Then

$$a(G) = \sum_\lambda c_\lambda. \tag{3.18}$$

35

Furthermore, as a refinement, define $a_m(G)$ to be the number of acyclic orientations

of G having exactly m sinks. Then

$$a_m(G) = \sum_{l(\lambda)=m} c_\lambda. \qquad (3.19)$$

That is, $a_m(G)$ is given by the sum of those c_λ corresponding to partitions λ with

exactly m parts.

Notably, the proof method of [38] uses a novel algebraic argument that does not

generalize directly either the argument of [37] or the inductive method suggested

by (3.16).

We will prove a generalization for vertex-weighted graphs using induction and

the deletion-contraction relation. In this way, we also provide an alternate proof of

(3.19) that is a natural extension of enumerative proofs of (3.17).

We first establish some notation and terminology. For a symmetric function f,

if $f = \sum_\lambda c_\lambda e_\lambda$ is its expansion in the basis of elementary symmetric functions, we

define $\sigma(f) = \sum_\lambda c_\lambda$, and $\sigma_m(f) = \sum_{l(\lambda)=m} c_\lambda$.

For an acyclic orientation γ of a vertex-weighted graph (G, w), we define $Sink(\gamma)$

to be the set of sinks of G with respect to γ (note that as γ is acyclic, $Sink(\gamma)$

is always nonempty). Let $sink(\gamma) = |Sink(\gamma)|$. Define a *sink map S of γ* to

be a function $S : Sink(\gamma) \to 2^{\mathbb{N}}$ such that for all $v \in Sink(\gamma)$, $\emptyset \neq S(v) \subseteq$

$\{1, 2, \ldots, w(v)\}$. Given a sink map S of an acyclic orientation γ on a vertex-weighted

graph (G, w), we define its *sink weight* to be $swt(G, w, \gamma, S) = \sum_{v \in Sink(\gamma)} |S(v)|$.

When (G, w) and/or γ are clear from context we may use $swt(S)$ or $swt(\gamma, S)$ in place of $swt(G, w, \gamma, S)$ for brevity.

We now state the main theorem of this section:

Theorem 9. *Let (G, w) be a vertex-weighted graph with n vertices and total weight d. Then*

$$\sigma(X_{(G,w)}) = (-1)^{d-n} \sum_{(\gamma, S)} (-1)^{swt(S)-sink(\gamma)}$$

where the sum runs over all ordered pairs (γ, S) such that γ is an acyclic orientation of (G, w), and S is a sink map of γ. Additionally,

$$\sigma_m(X_{(G,w)}) = (-1)^{d-n} \sum_{swt(\gamma,S)=m} (-1)^{m-sink(\gamma)} \tag{3.20}$$

where the sum ranges only over those ordered pairs (γ, S) with $swt(S) = m$.

Proof. It suffices to prove (3.20). We proceed by induction on the number of edges of (G, w). The base case is a vertex-weighted graph with no edges. If such a graph has vertices of weights $\lambda_1 \geq \cdots \geq \lambda_k$, then $X_{(G,w)} = p_\lambda$ where $\lambda = (\lambda_1, \ldots, \lambda_k)$.

First, we establish the following identity for any positive integer a:

$$\sigma_m(p_a) = (-1)^{a-m} \binom{a}{m} \tag{3.21}$$

We show this for fixed a by induction on m, making use of Newton's identity ([27], Chapter 1.2):

$$p_a = (-1)^{a-1} a e_a + \sum_{i=1}^{a-1} (-1)^{a-1+i} e_{a-i} p_i.$$

37

The case $m = 1$ is clear from this. Now we assume the claim holds for $m = b-1$ and prove it for $m = b$. Using Newton's identity followed by the inductive hypothesis we have

$$\sigma_b(p_a) = \sum_{i=1}^{a-1} (-1)^{a-1+i} \sigma_{b-1}(p_i) = \sum_{i=1}^{a-1} (-1)^{a-1+i}(-1)^{i-b+1} \binom{i}{b-1}$$

$$= (-1)^{a-b} \sum_{i=1}^{a-1} \binom{i}{b-1} = (-1)^{a-b} \binom{a}{b}$$

where we have used the Hockey Stick Identity.

We now establish the base case for the induction of the main proof. Recall that (G, w) has vertices of weights $\lambda_1 \geq \cdots \geq \lambda_k$ and no edges. First we evaluate directly the left-hand side of (3.20):

$$\sigma_m(X_{(G,w)}) = \sigma_m(p_{(\lambda_1,\dots,\lambda_k)}) = \sum_{(a_1,\dots,a_k)} \sigma_{a_1}(p_{\lambda_1}) \cdots \sigma_{a_k}(p_{\lambda_k})$$

where this sum runs over all tuples $(a_1, \dots, a_k)$ of positive integers satisfying $a_i \leq \lambda_i$ and

$a_1 + \cdots + a_k = m$. Expanding using (3.21), we get

$$\sum_{(a_1,\dots,a_k)} \prod_{i=1}^{k} (-1)^{\lambda_i-a_i} \binom{\lambda_i}{a_i} = \sum_{(a_1,\dots,a_k)} (-1)^{|\lambda|-m} \prod_{i=1}^{k} \binom{\lambda_i}{a_i}. \tag{3.22}$$

Next we will simplify the right-hand side of (3.20) and show that it is equal to (3.22). In (G, w) there is only one acyclic orientation γ, the empty orientation, and all vertices are sinks in this orientation, so equivalently we are looking for all ways to choose the sink map S such that $swt(S) = m$. Then the sum simplifies to

$$(-1)^{d-n} \sum_{swt(\gamma,S)=m} (-1)^{m-sink(\gamma)} = (-1)^{|\lambda|-k}(-1)^{m-k} \sum_{(a_1,\dots,a_k)} \prod_{i=1}^{k} \binom{\lambda_i}{a_i}$$

where the sum runs over the same tuples as in (3.22). Clearly these sums are equal, and this establishes the base case.

We now show the inductive step. Let (G, w) be a vertex-weighted graph with $g \geq 1$ edges, and assume that (3.20) holds for all vertex-weighted graphs with fewer than g edges. We may assume that (G, w) has no loops, as otherwise both sides of (3.20) are 0. Let e be an edge of (G, w), with endpoints v_1 and v_2. In $(G/e, w/e)$, let v^* be the vertex arising from the contraction of v_1 and v_2. Taking the deletion-contraction relation (3.4), applying σ_m to both sides, and multiplying both sides by $(-1)^{d-n}$ we have

$$(-1)^{d-n}\sigma_m(X_{(G \backslash e, w)}) = (-1)^{d-n}\sigma_m(X_{(G, w)}) - (-1)^{d-n-1}\sigma_m(X_{(G/e, w/e)}).$$

By the inductive hypothesis

$$(-1)^{d-n}\sigma_m(X_{(G \backslash e, w)}) = \sum_{swt(G \backslash e, \gamma, S)=m} (-1)^{m-sink(\gamma)}$$

and

$$(-1)^{d-n-1}\sigma_m(X_{(G/e, w/e)}) = \sum_{swt(G/e, \gamma, S)=m} (-1)^{m-sink(\gamma)}.$$

To finish the proof it suffices to show that

$$\sum_{swt(G \backslash e, \gamma, S)=m} (-1)^{m-sink(\gamma)} = \sum_{swt(G, \gamma, S)=m} (-1)^{m-sink(\gamma)} - \sum_{swt(G/e, \gamma, S)=m} (-1)^{m-sink(\gamma)}$$

or after multiplying both sides by $(-1)^m$,

$$\sum_{swt(G \backslash e, \gamma, S)=m} (-1)^{sink(\gamma)} = \sum_{swt(G, \gamma, S)=m} (-1)^{sink(\gamma)} - \sum_{swt(G/e, \gamma, S)=m} (-1)^{sink(\gamma)}. \quad (3.23)$$

39

To prove (3.23), we will work over a larger class of maps S whose domain is the set of all vertices of a graph instead of just the sinks of a given acyclic orientation γ, and we also allow $S(v) = \emptyset$ for all vertices v, while still requiring that $S(v) \subseteq \{1, 2, \ldots, w(v)\}$. We call S γ-admissible if $S(v) \neq \emptyset$ if and only if $v \in Sink(\gamma)$. Thus we may rephrase (3.23) by allowing γ and S in the summations to range over all acyclic orientations γ and all sink maps S with $S(v) \subseteq \{1, 2, \ldots, w(v)\}$ for all v, but where we define the corresponding summand to be $(-1)^{sink(\gamma)}$ if and only if S is γ-admissible, and 0 otherwise.

We show that for every acyclic orientation γ_0 of $G \backslash e$, and every map $S_0 : V(G) \to 2^{\mathbb{N}}$ with $S_0(v) \subseteq \{1, 2, \ldots, w(v)\}$ such that $\sum_{v \in V(G)} |S_0(v)| = m$, the equation (3.23) is satisfied when summing over those γ and S where

- in $G \backslash e$, $\gamma = \gamma_0$ and $S = S_0$.

- in G, γ restricted to $G \backslash e$ is γ_0, and $S = S_0$. This yields two choices for γ depending on the orientation of the edge $v_1 v_2$. Let γ_{v_1} be the one where $v_1 \to v_2$, and γ_{v_2} the one where $v_2 \to v_1$.

- in G/e, $\gamma = \gamma_{v^*}$ is the contraction of γ_0, and $S = S'$ is defined by $S'(v) = S_0(v)$ if $v \neq v^*$, and $S'(v^*) = S(v_1) \cup \{w(v_1) + i : i \in S(v_2)\}$.

It is easy to check that every pair (γ, S) for each of $G \backslash e$, G, and G/e is derived from exactly one such (γ_0, S_0), so proving this claim will finish the proof of the theorem.

40

For ease of notation, we fix γ_0 and S_0 in what follows. Let $T(G\backslash e)$ denote the term corresponding to γ_0 and S_0 in the summation for $G\backslash e$ in (3.23), and likewise let $T(G/e)$ denote the term in the summation for G/e corresponding to γ_{v_*} and S'. Let $T(G_{v_1})$ denote the term in the summation for G corresponding to S_0 and γ_{v_1}, and likewise for $T(G_{v_2})$. Thus what we must show for every fixed γ_0 and S_0 is

$$T(G\backslash e) = T(G_{v_1}) + T(G_{v_2}) - T(G/e) \tag{3.24}$$

We proceed by cases:

Case 1: γ_0 has a directed path from v_1 to v_2 or from v_2 to v_1

Note that γ_0 cannot have both of these directed paths since γ_0 is acyclic. Without loss of generality we assume the path is from v_2 to v_1. Then $T(G/e) = 0$ because γ_{v^*} is not acyclic. Also $T(G\backslash e) = (-1)^{sink(\gamma_0)}$ if S_0 is γ_0-admissible, and 0 otherwise. The orientation γ_{v_1} is not acyclic, so $T(G_{v_1}) = 0$. However, γ_{v_2} is acyclic, and has the same set of sinks as γ_0, so also $T(G_{v_2}) = (-1)^{sink(\gamma_0)}$ if S_0 is γ_0-admissible and 0 otherwise, and this satisfies (3.24).

Note that from now on, since we may assume there is no directed path between v_1 and v_2 in γ_0, the orientation γ_{v^*} is acyclic.

Case 2: Neither v_1 nor v_2 is a sink with respect to γ_0

In this case, v_1 and v_2 are also not sinks in γ_{v_1} or γ_{v_2}, and v^* is not a sink in γ_{v^*}, so if it is not the case that $S_0(v_1) = S_0(v_2) = \emptyset$, all terms of (3.24) are equal to 0. Otherwise, all terms are equal to 1. In either case, (3.24) is satisfied.

Case 3: Exactly one of v_1 or v_2 is a sink with respect to γ_0

Without loss of generality we may assume that v_1 is a sink; the case where v_2 is a sink is analogous. Similarly to the previous case, if $S_0(v_2) \neq \emptyset$ then all terms of (3.24) are equal to 0 (note that in γ_{v^*}, vertex v^* is a sink if and only if both v_1 and v_2 are). Thus, we may assume that $S_0(v_2) = \emptyset$. We have two subcases to consider:

Case 3.1: $S_0(v_1) = \emptyset$

In this case $T(G \backslash e) = 0$ as S_0 is not γ_0-admissible. Additionally, $T(G_{v_2}) = 0$ as S_0 is not γ_{v_2}-admissible. However, S_0 is γ_{v_1}-admissible since v_1 is no longer a sink in γ_{v_1}, so $T(G_{v_1}) = (-1)^{sink(\gamma_0)-1}$. Also, as v^* is not a sink in γ_{v^*}, we have $T(G/e) = (-1)^{sink(\gamma_0)-1}$. Thus, this case satisfies (3.24).

Case 3.2: $S_0(v_1) \neq \emptyset$

Then S_0 is γ_0-admissible, so $T(G \backslash e) = (-1)^{sink(\gamma_0)}$. Also S_0 is γ_{v_2}-admissible, but it is not γ_{v_1}-admissible, so $T(G_{v_2}) = (-1)^{sink(\gamma_0)}$ and $T(G_{v_1}) = 0$. The map S' is not γ_{v^*}-admissible, since v^* is not a sink, but $S'(v^*) \neq \emptyset$ since $S_0(u) \neq \emptyset$, so $T(G/e) = 0$. This satisfies (3.24).

Case 4: Both v_1 and v_2 are sinks with respect to γ_0

If $S_0(v_1) = S_0(v_2) = \emptyset$, then all terms of (3.24) are equal to 0. Thus we may assume that at least one of these sets is nonempty. We again split into subcases.

Case 4.1: Exactly one of $S_0(v_1)$ and $S_0(v_2)$ is nonempty

Without loss of generality we may assume that $S_0(v_1) \neq \emptyset$ and $S_0(v_2) = \emptyset$; the other case is analogous. Then S_0 is not γ_0-admissible, so $T(G \backslash e) = 0$. Also, S_0 is not γ_{v_1}-admissible, so $T(G_{v_1}) = 0$. However, S_0 is γ_{v_2}-admissible since here v_2 is no longer a sink, so $T(G_{v_2}) = (-1)^{sink(\gamma_0)-1}$. In γ_{v^*}, the contracted vertex v^* is a sink and $S'(v^*)$ is nonempty, so $T(G/e) = (-1)^{sink(\gamma_0)-1}$ since the two sinks v_1 and v_2 became one sink. This satisfies (3.24).

Case 4.2: Both $S_0(v_1)$ and $S_0(v_2)$ are nonempty

In this case S_0 is γ_0-admissible, so $T(G \backslash e) = (-1)^{sink(\gamma_0)}$. However, S_0 is neither γ_{v_1}-admissible nor γ_{v_2}-admissible, so $T(G_{v_1}) = T(G_{v_2}) = 0$. In γ_{v^*}, the contracted vertex v^* is a sink, and $S'(v^*)$ is nonempty, so $T(G/e) = (-1)^{sink(\gamma_0)-1}$, since the two sinks v_1 and v_2 became one sink. This satisfies (3.24).

Thus we have shown that (3.24) holds in all cases, and this finishes the proof.

$\square$

3.3   Further Applications

Considering vertex-weighted graphs with the chromatic symmetric function provides a new perspective and new tools for approaching major unsolved problems. We mention some of these problems and possible approaches.

3.3.1   Chromatic Quasisymmetric Functions

A *quasisymmetric function* is a function $f \in \mathbb{R}[[x_1, x_2, \ldots]]$ of bounded degree such that the coefficient of any specific monomial $x_{i_1}^{a_1} \ldots x_{i_k}^{a_k}$ with $i_1 < \cdots < i_k$ is dependent only on the order of the indices $i_1, \ldots, i_k$ and not on their value. Thus, the symmetry of the coefficients of monomials does not necessarily hold for all permutations of their indices, but for any one that "slides" the indices while maintaining their order. A well-researched generalization of the chromatic symmetric function is the chromatic quasisymmetric function of Shareshian and Wachs [42], defined on vertex-labeled graphs, or equivalently on graphs equipped with an acyclic orientation. In the context of finding a deletion-contraction relation it is more natural to look at the generalization of this function to simple graphs with an arbitrary orientation, considered by Ellzey [12] and Alexandersson and Panova [1]. Given a graph G with a fixed orientation γ, for any proper coloring κ of G define the *ascent number* $asc(\kappa)$ to be the number of edges $v_1 \to v_2$ of γ such that $\kappa(v_1) < \kappa(v_2)$. Using the notation $x_\kappa(G) = \prod_{v \in V(G)} x_{\kappa(v)}$, define the *chromatic quasisymmetric function of G with respect to γ* as

$$X_{(G,\gamma)}(q, x_1, x_2, \dots) = \sum_{\kappa} x_{\kappa}(G) q^{asc(\kappa)} \tag{3.25}$$

where the sum runs over all proper colorings κ of G.

It is natural to try to extend our definition on vertex-weighted graphs to work in this setting. Ideally, such an extension would equip the chromatic quasisymmetric function with a deletion-contraction relation. However, a first attempt

$$X_{(G,w,\gamma)}(q, x_1, x_2, \dots) = \sum_{\kappa} x_{\kappa}(G, w) q^{asc(\kappa)} \tag{3.26}$$

fails to provide a deletion-contraction relation. To see this, consider a vertex-weighted graph (G, w) with edge $e = v_1 v_2$, and assume we are considering an orientation of G in which $v_1 \to v_2$. In the following table, we determine how the power of q in a proper coloring κ of $G \backslash e$ relates to the power of q of corresponding proper colorings of G/e or G (thus all numbers $asc(\kappa)$ are relative to $G \backslash e$):

$\kappa(v_1)$ vs $\kappa(v_2)$	$G \backslash e$	G/e	G
$=$	$q^{asc(\kappa)}$	$q^{asc(\kappa)}$	N/A
$>$	$q^{asc(\kappa)}$	N/A	$q^{asc(\kappa)}$
$<$	$q^{asc(\kappa)}$	N/A	$q^{asc(\kappa)+1}$

An explicit example illustrates the problem implied by the asymmetry of this table. Consider the three-vertex path P_3 with orientation γ_0 that has exactly one sink, and exactly one source (a vertex that is not the head of any oriented edge), and with all vertex weights equal to 1. Then if we consider powers of q in the specialization $x_1 = x_2 = x_3 = 1$, $x_i = 0$ for $i \geq 4$, we have

45

$$X_{(P_3,1^3,\gamma_0)}(q,1,1,1,0,0,\dots) = q^2 + 10q + 1.$$

If we take the edge e of this path that has the source as an endpoint, then using the same specialization we have

$$X_{(P_3 \backslash e, 1^3, \gamma_0)} = 9q + 9$$

and

$$X_{(P_3/e,1^3/e,\gamma_0/e)} = 3q + 3.$$

Note that $q + 1$ divides the second and third of these functions but not the first, so we cannot have a simple deletion-contraction relation of the form $X_{(G,w,\gamma)} = (f(q))^a X_{(G\backslash e,w,\gamma\backslash e)} \pm (g(q))^b X_{(G/e,w/e,\gamma/e)}$ where f and g are polynomials in q.

There is a way to provide a deletion-contraction analog if we expand upon how the relation may look. Given a vertex-weighted graph (G, w) with fixed edge $e = v_1 v_2$, and an orientation γ, define $\gamma_{\overleftarrow{e}}$ to be the same orientation as γ except with the order of the head and tail of edge e reversed. Then the following relation holds:

Lemma 10. *Let (G, w) be a vertex-weighted graph, and let e be an edge of G. Let γ be an orientation of G. Then*

$$X_{(G,w,\gamma)} + X_{(G,w,\gamma_{\overleftarrow{e}})} = (1+q)(X_{(G\backslash e,w,\gamma)} - X_{(G/e,w/e,\gamma/e)}). \tag{3.27}$$

Proof. We first rearrange (3.27) into

$$(1+q)X_{(G\backslash e,w,\gamma)} = X_{(G,w,\gamma)} + X_{(G,w,\gamma_{\overleftarrow{e}})} + (1+q)X_{(G/e,w/e,\gamma/e)}. \tag{3.28}$$

The result is clear if e is a loop, so let $e = v_1 v_2$ have distinct endpoints. We show that (3.28) holds by showing a one-to-one correspondence between the terms on the left-hand side, and sets of terms on the right-hand side. Consider a proper coloring κ of $G\backslash e$, and let $asc(\kappa)$ be the ascent number of κ with respect to γ. We split into cases based on the colors κ gives to v_1 and v_2.

If $\kappa(v_1) = \kappa(v_2)$, then the term $(1+q)x_\kappa(G\backslash e, w)q^{asc(\kappa)}$ is equal to the term $(1+q)x_{\kappa^e}(G/e, w/e)q^{asc(\kappa)}$ of $(1+q)X_{(G/e, w/e, \gamma/e)}$, where κ^e is the contraction of the coloring κ with respect to e. Furthermore, there is no corresponding term of either $X_{(G,w,\gamma)}$ or $X_{(G,w,\gamma_{\overleftarrow{e}})}$ since κ is not a proper coloring of G.

If $\kappa(v_1) \neq \kappa(v_2)$, then we get a term of $(1+q)x_\kappa(G\backslash e, w)q^{asc(\kappa)}$ on the left-hand side of (3.28) as before. There is no corresponding term of $X_{(G/e, w/e, \gamma/e)}$ since the coloring κ does not contract to one on G/e. We do get corresponding terms of both $X_{(G,w,\gamma)}$ and $X_{(G,w,\gamma_{\overleftarrow{e}})}$, equal to $x_\kappa(G, w)q^{asc(\kappa)}$ and $x_\kappa(G, w)q^{asc(\kappa)+1}$ in some order depending on the orientation of e. Furthermore, these terms satisfy

$$(1+q)x_\kappa(G\backslash e, w)q^{asc(\kappa)} = x_\kappa(G, w)q^{asc(\kappa)} + x_\kappa(G, w)q^{asc(\kappa)+1}$$

since $x_\kappa(G\backslash e, w) = x_\kappa(G, w)$.

Thus, there is a bijective correspondence of terms from the left-hand side of (3.28) to a unique set of terms on the right-hand side, and this concludes the proof.

$\square$

3.3.2 e- and s-positivity

Another possible application of the deletion-contraction method is expanding the chromatic symmetric functions of certain families of graphs to prove that the coefficients are nonnegative in a fixed basis.

For example, since the sum of all coefficients of fixed length in the e-basis is nonnegative, it is natural to consider whether X_G might generally be e-positive. This is not the case, and there are many small graphs that give counterexamples. The smallest such graph is the *claw*, a graph isomorphic to $(\{a, b, c, d\}, \{ab, ac, ad\})$. Many other graphs that contain the claw as an induced subgraph are also not e-positive. Moreover, it is not even sufficient to be claw-free, as the *net* (a graph isomorphic to $(\{a, b, c, d, f, g\}, \{ab, ac, bc, ad, bf, cg\})$) is a claw-free graph that is not e-positive (see Figure 3.1).

There are some ways around these obstacles. One approach is to attempt to continue forbidding counterexamples as induced subgraphs to reduce to a family that is e-positive. Some research in this direction has been conducted in [20], in which the authors classify for almost all four-vertex graphs H which ones satisfy that $\{claw, H\}$-free graphs are e-positive, and [15], which conjectures that a graph G is e-positive and has all induced subgraphs e-positive if and only if G is $\{claw, net\}$-free.

More commonly, research in positivity has focused on reducing to an e-positive family by considering only graphs with certain nice properties. Currently the most

Figure 3.1: The claw (top) and the net (bottom), small graphs that are not e-positive.

commonly researched family is that of incomparability graphs of partially ordered sets [4, 10, 11, 15, 19, 38]. A *partially ordered set* (or *poset*) is pair $(P, \preceq)$ where P is a set, and $\preceq$ is a binary relation on P that is a *partial order* on P, meaning that the following are satisfied:

- $p \preceq p$ for all $p \in P$ (reflexivity)

- For any distinct $p, q \in P$ at most one of $p \preceq q$ and $q \preceq p$ holds (antisymmetry)

- For any $p, q, r \in P$, if $p \preceq q$ and $q \preceq r$ then $p \preceq r$ (transitivity)

Distinct elements $p, q \in P$ are *comparable* if either $p \preceq q$ or $q \preceq p$; otherwise p and q are *incomparable*. Two posets $(P, \preceq_P)$ and $(Q, \preceq_Q)$ are *isomorphic* if there

exists a map $f : P \to Q$ such that for every $p_1, p_2 \in P$, we have $p_1 \preceq_P p_2$ if and only if $f(p_1) \preceq_Q f(p_2)$, and analogously for $p_2 \preceq_P p_1$. The *incomparability graph* of the poset $(P, \preceq)$ is the simple graph with vertex set P and edge set $\{pq :$ p and q are incomparable$\}$. It is denoted by $Inc(P, \preceq)$, or just $Inc(P)$ if there is no ambiguity in the choice of relation $\preceq$. Any $R \subseteq P$ defines an *induced subposet* of P by simply considering $(R, \preceq)$.

Research on positivity focuses on incomparability graphs of $(3+1)$-free posets, meaning those $Inc(P)$ where P does not contain an induced subposet isomorphic to $(\{a, b, c, d\}, \preceq)$, where $a \preceq b, b \preceq c, a \preceq c$, and d is incomparable to all of a, b, c. This specific choice of forbidden subposet is made since the corresponding incomparability graphs are then claw-free. In particular, one of the most important open problems involving the chromatic symmetric function is the Stanley-Stembridge conjecture:

Conjecture 11. *([38]) Every incomparability graph of a $(3+1)$-free poset is e-positive.*[5]

Thus, we consider *e*-positivity of vertex-weighted graphs in an attempt to approach the Stanley-Stembridge conjecture, perhaps by generalizing it to a class of graphs on which deletion-contraction may be applied.

---

[5]Notably, [19] shows that this conjecture is equivalent to the statement that unit interval graphs are *e*-positive. This is a seemingly stronger statement, since unit interval graphs are precisely the incomparability graphs of posets that are simultaneously $(3+1)-$ and $(2+2)-$free.

In the case of a vertex-weighted graph (G, w) with n vertices and total weight d, it is easy to see from equations (3.6) and (3.12) that in the e-basis expansion the coefficient of e_d is $(-1)^{d-n}$, so the natural extension is to ask whether $(-1)^{d-n} X_{(G,w)}$ is e-positive.

This question of e-positivity can be answered for all vertex-weighted graphs with nontrivial vertex weights. We define a *connected partition* of a vertex-weighted graph (G, w) to be a partition $P_1 \sqcup \cdots \sqcup P_m = V(G)$ of the vertex set such that for each i, the subgraph of G induced by restricting to P_i is connected. We define the *type* of a connected partition to be the integer partition whose parts are $w(P_1), \ldots, w(P_m)$. The following lemma may be proved by a straightforward generalization of the proof of ([45], Proposition 1.3.3):

Lemma 12. *If (G, w) is a vertex-weighted graph with n vertices and total weight d such that $(-1)^{d-n} X_{(G,w)}$ is e-positive, and (G, w) has a connected partition of type $\lambda \vdash d$, then it also has a connected partition of type μ for every partition μ that is a refinement of λ.*

This yields:

Corollary 13. *Let (G, w) be a vertex-weighted graph. If there is a vertex $v \in V(G)$ such that $w(v) > 1$, then $(-1)^{d-n} X_{(G,w)}$ is not e-positive.*

Proof. Let (G, w) have vertex weights $w_1 \geq \cdots \geq w_n$. Then (G, w) has a connected partition of type $(w_1, \ldots, w_n)$, but it does not have one of type 1^d, so by the previous lemma, $(-1)^{d-n} X_{(G,w)}$ is not e-positive.

□

Although this answers the natural weighted analogue of the Stanley-Stembridge Conjecture, there is more work to be done here. The e-basis may not be optimal for considering vertex-weighted graphs; perhaps there is a choice of basis better suited for positivity questions in this setting. Alternatively, perhaps we may still modify the conjecture to apply in this setting; for example, one possible approach would be to lower-bound the e-basis coefficients of certain vertex-weighted graphs of n vertices and total weight d by a function of the "excess weight" $d - n$ in such a way that this lower bound is 0 when $d - n = 0$.

In addition to e-positivity, a result of Gasharov provided the first major step into studying s-positivity of the chromatic symmetric function. Given a poset $(P, \preceq)$, and a partition $\lambda \vdash |P|$, define a *P-tableau of shape λ* to be a filling of the Young diagram of shape λ with elements of P, each occurring exactly once, satisfying

- The entries are strictly increasing along rows (if q is immediately to the right of p, then $p \preceq q$).

- The entries are weakly increasing down columns (if q is immediately below p, then $q \npreceq p$).

Thus, a P-tableau of shape λ is a generalization of the notion of a semi-standard Young tableau (with rows and columns switched). Gasharov [16] showed:

52

Theorem 14. *Let P be a $(3+1)$-free poset, and for each $\lambda \vdash |P|$, let f_λ^P be the number of P-tableaux of shape λ. Then*

$$X_{Inc(P)} = \sum_{\lambda \vdash |P|} f_\lambda^P s_\lambda.$$

In particular, the incomparability graph of a $(3+1)$-free poset is s-positive.

This is a strictly weaker result than the Stanley-Stembridge conjecture since every e_λ is s-positive, but it provides strong evidence in support of the conjecture. It also admits a different natural conjecture, that perhaps all claw-free graphs are s-positive, even those that are not incomparability graphs, and this conjecture has separately received further attention [33, 34]. Some further results have even determined the value of specific coefficients in the s-basis expansion of X_G for any G. For example, Kaliszewski [23] showed

Theorem 15. *Let G be a graph with n vertices, and let $a_k(G)$ be the number of acyclic orientations of G with exactly k sinks. Then when expanding X_G in the basis of Schur functions, the coefficient of $s_{(m,1^{n-m})}$ is*

$$\sum_{k=1}^{n} \binom{k-1}{m-1} a_k(G)$$

Very recently, David and Monica Wang [44] have shown a general formula for any Schur function coefficient of a chromatic symmetric function as a weighted, signed sum of special rim-hook tabloids. Although there are not yet have results concerning the s-positivity of vertex-weighted graphs, we believe that the use of a deletion-contraction relation could prove useful in this context as well.

3.3.3 Other Possible Applications

There are many further possible applications of the vertex-weighted graph construction:

- Partition systems as described by Lenart and Ray [25]. Given a set S and a partition σ of S, they define a partition system P to be a set of subsets of S such that P contains the empty set and all blocks of σ, and all other elements of P are unions of the blocks of σ (henceforth the blocks of σ are called atoms of P, and the set of atoms is denoted $At(P)$ so we may drop mention of σ). One example of a partition system is the *independence complex* $I(G)$ of a graph G, where the set is $V(G)$, and $I(G) = \{A \subseteq V(G) : \forall v_1, v_2 \in A, v_1 v_2 \notin E(G)\}$. In particular, $I(G)$ is an abstract simplicial complex, meaning that for every $A \subseteq V(G)$, $A \in I(G) \to B \in I(G) \forall B \subseteq A$.

  For a partition system P on a set S, one may define

  $$X_P = \sum_{\kappa} \prod_{s \in S} x_{\kappa(s)}$$

  where the sum ranges over all $\kappa : S \to \mathbb{Z}^+$ such that for each positive integer i, the set $\{s \in S : \kappa(s) = i\}$ is in P. Then clearly $X_G = X_{I(G)}$. Lenart and Ray showed that for any nonempty $U \in P$ that is not an atom

  $$X_P = X_{P\backslash\backslash U} + X_{P/U}$$

  where $P\backslash\backslash U = P\backslash\{W : U \subseteq W\}$ and $P/U = \{W \in P : U \subseteq W$ or $U \cap W = \emptyset\}$. In the case $P = I(G)$ where G is a simple graph, the atoms are all vertices

of G, and the possible choices of U are the independent sets of G. Taking the particular choice of $U = v_1 v_2$ where $v_1 v_2$ is a nonedge of G yields

$$X_{G \backslash e} = X_G + X_{P/v_1 v_2}.$$

This does not provide a direct deletion-contraction relation for X_G because in $P/v_1 v_2$ not all atoms are singletons, and thus the partition system does not correspond to the independence complex of a graph. Our $X_{(G,w)}$ provides the necessary construction to generalize this result to graphs, and it may be possible to relate other results on X_P to $X_{(G,w)}$.

- Lenart and Ray also define an *umbral chromatic polynomial* $\chi^\varphi(P; x)$ [25, 35] for any partition system P, defined as a polynomial over the algebra $\mathbb{Z}[\varphi_1, \varphi_2, \dots]$, where the φ_i are indeterminates. They show that for graphs G, knowing $\chi^\varphi(I(G); x)$ is equivalent to knowing X_G, and that $\chi^\varphi(P; x)$ satisfies an analog of deletion-contraction ([25], Proposition 5.7). As with X_P, this relation does not correspond to one on X_G because it includes terms involving partition systems P that do not correspond to graphs; however, the introduction of $X_{(G,w)}$ provides an intermediate step, and it may be of interest to determine how the vertex-weighted graph construction fits into this setting, especially in the context of the addition-contraction tree given in [35].

- The path-cycle symmetric function on digraphs defined by Chow [6]. In particular, the most appropriate formulation to generalize is likely

$$\Theta_D(x_1, x_2, \ldots, y_1, y_2, \ldots) = \sum_{(S,\kappa)} \prod_{v_1 \text{ in a path}} x_{\kappa(v_1)} \prod_{v_2 \text{ in a cycle}} y_{\kappa(v_2)}$$

where the sum ranges over path-cycle covers S and colorings κ of D such that paths and cycles are monochromatic, and the paths receive distinct colors. Perhaps there is something akin to a natural deletion-contraction relation in this context.

- The double poset construction of Grinberg [18]. He defines a double poset $(E, \preceq_1, \preceq_2)$ using two partial orders on the same base set, and given a weight function $w : E \to \mathbb{Z}^+$, he defines the function

$$\Gamma(E, w) = \sum_{\pi : E \to \mathbb{Z}^+} \prod_{e \in E} x_{\pi(e)}^{w(e)}$$

where the sum runs over all π that are E-partitions, a notion that generalizes (P, ω)-partitions (see [39]). Results on $\Gamma(E, w)$ are in some cases directly applicable to $X_{(G,w)}$. For example, as an alternate proof of (3.8), one could fix an orientation γ and pass to quasisymmetric functions. Then, upon swapping out the symmetric function involution for the related antipode on quasisymmetric functions, (3.8) reduces to a special case of ([18], Theorem 4.2). There are likely other relations between these functions waiting to be discovered.

3.3.4 A Brief Note On Weighted Trees

As a final note we would be remiss not to mention one of the other major open problems of the chromatic symmetric function, the tree isomorphism conjecture, inspired by a question of Stanley [38]:

Conjecture 16. *If G and H are trees, and $X_G = X_H$, then G and H are isomorphic.*

This is a popularly-researched conjecture [2, 21, 26, 28, 29] and has been shown to be true for trees with up to 29 vertices [22]. A natural question is whether it is possible that a stronger statement holds, that the chromatic symmetric function distinguishes vertex-weighted trees. We define a *w-isomorphism* of vertex-weighted graphs (G, w) and (G', w') to be a map $f : V(G) \to V(G')$ that is a graph isomorphism of G and G' that also satisfies that for any $v \in V(G)$, we have $w(v) = w'(f(v))$. Then the graphs (G, w) and (G', w') are called *w-isomorphic.* Thus the corresponding question is whether there exist non-w-isomorphic vertex-weighted trees (T, w) and (T', w') with $X_{(T,w)} = X_{(T',w')}$. Such pairs of trees do exist, like the following example from [26] shown in Figure 3.2. We compare

(a) The five-vertex path with vertex weights $1, 2, 1, 3, 2$ in that order, and

(b) The five-vertex path with vertex weights $1, 3, 2, 1, 2$ in that order.

It is seen easily that these vertex-weighted trees are not w-isomorphic. To see that nonetheless they have the same chromatic symmetric function, we apply the addition form (3.4) of the deletion-contraction rule to the non-edge represented by

Figure 3.2: Weighted trees with the same chromatic symmetric function

the dashed line. Then the chromatic symmetric function of both (a) and (b) is the same as that of a five-vertex cycle with vertex weights $1, 2, 1, 3, 2$ cyclically, added to that of a four-vertex cycle with vertex weights $3, 3, 2, 1$ cyclically.

However, in this example the two underlying unweighted trees are isomorphic. We do not know of an example of two vertex-weighted trees (T, w) and (T', w') with T and T' nonisomorphic as unweighted trees, but with $X_{(T,w)} = X_{(T',w')}$.

3.4 The Weighted Version of the Bad-Coloring Extension of X_G

In this section, we extend the definition of the vertex-weighted chromatic symmetric function to include *all* colorings of a graph, not just the proper ones. To this end, for a given (not necessarily proper) coloring κ of G, define

$$x_\kappa(G, w, t) = x_\kappa(G, w)(1 + t)^{c_\kappa(G)}$$

where $c_\kappa(G)$ is the number of edges in G that are monochromatic with respect to κ. We then define the *bad coloring symmetric function of a vertex-weighted graph* (G, w) to be the following analogue of the bad coloring function given in [41]:

$$XB_{(G,w)}(t, x_1, x_2, \dots) = \sum_\kappa x_\kappa(G, w, t) \tag{3.29}$$

where the sum is over all colorings κ of G (not just the proper ones). We use the convention $0^0 = 1$, so that when $t = -1$ it follows that

$$XB_{(G,w)}(-1, x_1, x_2, \dots) = X_{(G,w)}(x_1, x_2, \dots)$$

On vertex-weighted graphs, $XB_{(G,w)}$ allows the following deletion-contraction relation that generalizes Lemma 3:

Lemma 17. *Let (G, w) be a vertex-weighted graph. For all $e \in E(G)$,*

$$XB_{(G,w)} = XB_{(G\backslash e, w)} + tXB_{(G/e, w/e)} \tag{3.30}$$

Proof. First, note that when $t = -1$, the deletion-contraction relation (3.30) reduces to the one on $X_{(G,w)}$ that we have already proved. Thus from now on we may assume $t \neq -1$.

The case when e is a loop follows immediately from the defintion of XB, so we may assume that e is not a loop. Let v_1 and v_2 be the endpoints of e. We start with the right-hand side of (3.30) and expand using the definition (3.29) of XB:

$$XB_{(G\backslash e, w)} + tXB_{(G/e, w/e)} =$$

$$\sum_{\kappa:V(G\backslash e)\to\mathbb{Z}^+} x_\kappa(G\backslash e, w, t) + t \sum_{\kappa:V(G/e)\to\mathbb{Z}^+} x_\kappa(G/e, w/e, t)$$

Note that all colorings of G are also colorings of $G\backslash e$, and vice versa. We split the κ in the first summand based on $\kappa(v_1)$ and $\kappa(v_2)$. In those κ where $\kappa(v_1) \neq \kappa(v_2)$, we have $c_\kappa(G\backslash e) = c_\kappa(G)$, so $x_\kappa(G\backslash e, w, t) = x_\kappa(G, w, t)$. In all κ with $\kappa(v_1) = \kappa(v_2)$, we have $c_\kappa(G\backslash e) = c_\kappa(G) - 1$ because of the missing edge e, so for these κ, we have $x_\kappa(G\backslash e, w, t) = (1 + t)^{-1} x_\kappa(G, w, t)$.

For the second summand, note that every κ of G/e corresponds naturally to a κ in G with $\kappa(v_1) = \kappa(v_2)$, and vice-versa (we will use the same κ to denote both of these colorings in a slight abuse of notation). For these κ we will have $c_\kappa(G/e) = c_\kappa(G) - 1$ since we are missing the contracted edge e, and thus for each such κ we will have $x_\kappa(G/e, w/e, t) = (1 + t)^{-1} x_\kappa(G, w, t)$. Putting everything together, we have

$$
\begin{aligned}
XB_{(G\backslash e, w)} + tXB_{(G/e, w/e)} &= \sum_{\kappa:V(G\backslash e)\to\mathbb{Z}^+} x_\kappa(G\backslash e, w, t) + t \sum_{\kappa:V(G/e)\to\mathbb{Z}^+} x_\kappa(G/e, w/e, t) \\
&= \sum_{\substack{\kappa:V(G)\to\mathbb{Z}^+ \\ \kappa(v_1)\neq\kappa(v_2)}} x_\kappa(G, w, t) + \sum_{\substack{\kappa:V(G)\to\mathbb{Z}^+ \\ \kappa(v_1)=\kappa(v_2)}} (1+t)^{-1} x_\kappa(G, w, t) \\
&\quad + t \sum_{\substack{\kappa:V(G)\to\mathbb{Z}^+ \\ \kappa(v_1)=\kappa(v_2)}} (1+t)^{-1} x_\kappa(G, w, t) \\
&= \sum_{\kappa:V(G)\to\mathbb{Z}^+} x_\kappa(G, w, t) \\
&= XB_{(G, w)}
\end{aligned}
$$

as desired.

$\square$

As a consequence of this relation, we can modify the proof of (3.6) by simply replacing (-1)s with ts to give the following p-basis expansion formula, an analogue of the one in [41]:

Corollary 18.

$$XB_{(G,w)}(t, x_1, x_2, \dots) = \sum_{S \subseteq E(G)} t^{|S|} p_{\lambda(G,w,S)} \tag{3.31}$$

Furthermore, this recurrence relation together with $XB_{(G,w)} = p_{(w_1,\dots,w_k)}$ when (G, w) is a graph with no edges and vertices of weights $w_1 \geq \cdots \geq w_k$ can be taken as an alternate definition of $XB_{(G,w)}$. There is a similar recursive definition of the W-polynomial from invariant theory [31]. This (nonsymmetric!) function $W_{(G,w)}(y, x_1, x_2, \dots)$ on vertex-weighted graphs is defined by the following relations:

- When (G, w) is a graph with no edges and vertices of weights $w_1 \geq \cdots \geq w_k$, we have $W_{(G,w)} = x_{w_1} \dots x_{w_k}$.

- When $e \in E(G)$ is a loop, $W_{(G,w)} = y W_{(G \backslash e, w)}$.

- When $e \in E(G)$ is not a loop, $W_{(G,w)} = W_{(G \backslash e, w)} + W_{(G/e, w/e)}$.

Noble and Welsh show in [31] that using the above relations, one can prove by induction on the number of edges that the W-polynomial satisfies

$$W_{(G,w)}(y, x_1, x_2, \dots) = \sum_{S \subseteq E(G)} x_{c_1} \dots x_{c_k} (y - 1)^{|S|+k-|V(G)|} \tag{3.32}$$

61

where $c_1, \ldots, c_k$ are the total weights of the connected components of (G, w) induced by edges in S.

Two functions on vertex-weighted graphs (G, w) are said to be *equivalent* if given one, we can entirely recover the other, without necessarily knowing the graph (G, w) or any of its properties.

Lemma 19. *The functions $XB_{(G,w)}(t, x_1, x_2, \ldots)$ and $W_{(G,w)}(y, x_1, x_2, \ldots)$ are equivalent.*

Proof. We will in fact prove a stronger statement, that given $W_{(G,w)}$, we may recover the p-basis expansion of $XB_{(G,w)}$ by writing

$$XB_{(G,w)} = t^{|V(G)|} W_{(G,w)} \left(t + 1, \frac{p_1}{t}, \frac{p_2}{t}, \ldots, \frac{p_k}{t}, \ldots \right) \tag{3.33}$$

and conversely, given the p-basis expansion of $XB_{(G,w)}$, we may recover $W_{(G,w)}$ by dividing by $t^{|V(G)|}$, setting $t = y - 1$, and replacing each p_k with tx_k. This stronger statement may be proven as a simple vertex-weighted generalization of ([31], Theorem 6.2) by showing that these substitutions take (3.31) to (3.32) and vice-versa. We provide a different proof by showing that this substitution works not just for these equations, but for the base cases and inductive steps of the recursive definitions for $XB_{(G,w)}$ and $W_{(G,w)}$. In this sense these functions are not only equivalent, but essentially the same up to a change of variables.

The base cases for both functions are vertex-weighted graphs with no edges. Let (G, w) be a vertex-weighted graph with no edges and vertices of weights $w_1 \geq \cdots \geq$

62

w_k. Then $XB_{(G,w)} = p_{w_1} \ldots p_{w_k}$, $W_{(G,w)} = x_{w_1} \ldots x_{w_k}$, and we now verify that the substitution works. Going from W to XB we have:

$$x_{w_1} \ldots x_{w_k} \to t^k \left(\frac{p_{w_1}}{t}\right) \ldots \left(\frac{p_{w_k}}{t}\right) = p_{(w_1,\ldots,w_k)}$$

and the converse is analogous.

For the inductive step, suppose that we have demonstrated that this substitution is valid for graphs of $\leq m$ edges for some m. Let (G,w) be a vertex-weighted graph with $m+1$ edges and let e be any edge of G. Starting with the W-polynomial and using deletion-contraction we have two cases. First, if e is a loop, then $W_{(G,w)} = yW_{(G\backslash e,w)}$. Then applying our substitution we may derive $(t+1)XB_{(G\backslash e,w)} = XB_{(G,w)}$, and the converse is analogous.

If e is not a loop then $W_{(G,w)} = W_{(G\backslash e,w)} + W_{(G/e,w/e)}$. We must apply some care here as $G\backslash e$ and G/e have a different number of vertices. We make the substitution $x_i = \frac{p_i}{t}$, $y = t+1$, and multiply by $t^{|V(G)|}$. Then by the inductive hypothesis the resulting function is $XB_{(G\backslash e,w)} + tW_{(G/e,w/e)} = XB_{(G,w)}$ as desired, and again the converse process of recovering W from XB is analogous.

$\square$

As a brief note, this proof shows a formal of equivalence of the two functions as formal objects, but obviously there is a problem when $t = 0$. This is easily rectified by considering the $t = 0$ case separately, in which case it reduces essentially to the base case of the above argument.

The fact that these functions are not just equivalent but in some sense identical provides both functions with alternate perspectives for research. The function XB has the advantage that its deletion-contraction relation does not depend on whether e is a loop, which provides some simplification for inductive proofs. On the other hand, W of course is a sum of finitely many monomials, which may be easier to work with for elementary enumerative arguments.

The function $XB_{(G,w)}$ is also related to the weighted (r, q)-chromatic function of [24]. For a vertex-weighted graph (G, w) with n vertices, this function is defined as

$$M_{(G,w)}(r, q) = \sum_{S \subseteq E} (-1)^{|S|} \prod_{C \in C(S)} \sum_{i=0}^{n-1} r^{w(C)q^i}$$

where $C(S)$ is the set of connected components of G induced by S, and $w(C)$ is the total weight of component C.

This function has a natural extension with an additional parameter in the form

$$B_{(G,w)}(r, q, t) = \sum_{S \subseteq E} t^{|S|} \prod_{C \in C(S)} \sum_{i=0}^{n-1} r^{w(C)q^i}$$

Using the arguments from ([24], Section 3) and adjusting them to the vertex-weighted case it is easy to show

Lemma 20. $M_{(G,w)}(r, q)$ *is equivalent to* $X_{(G,w)}(x_1, x_2, \dots)$, *and* $B_{(G,w)}(r, q, t)$ *is equivalent to* $XB_{(G,w)}(t, x_1, x_2, \dots)$.

Chapter 4

Applications and Related

Questions

In this chapter we consider numerous problems related to chromatic symmetric functions and how to integrate and use the vertex-weighted extensions in these contexts. In Section 4.1 we answer in the negative a question posed to the author by Loebl as to whether X_G can distinguish nonisomorphic *split graphs* by giving a way to construct many pairs of counterexamples. This leads naturally into Section 4.2, in which we discuss other methods of constructing pairs of graphs with equal chromatic symmetric function, and provide a new conjecture for a nontrivial class of graphs which are distinguished by X_G with supporting qualitative evidence and numerical data. We also provide the first examples in the literature of pairs of graphs with equal XB_G. In Section 4.3 we define a new graph invariant, the *gen-*

eralized degree sequence, which arises from considering vertex-weighted chromatic symmetric functions ranging over different vertex weight functions w of the same base graph G. We prove some elementary properties of the generalized degree sequence, both in the case of all graphs and when restricted to trees. Finally, in Section 4.4 we consider a new basis of Λ defined as chromatic symmetric functions of certain graphs, determine some ways in which this basis relates to the five main bases, and conjecture a formula for relating it to the monomial basis. Although results in this section do not directly use any of the vertex-weighted material, it was largely inspired by it.

All material in this section is joint work with Sophie Spirkl. Sections 4.1 and 4.2 will be presented in [9] (currently in preparation). The author is not presently planning a specific manuscript using the material in Sections 4.3 and 4.4.

4.1 Chromatic Symmetric Functions on Split Graphs

A *bipartite graph* is a graph that has a proper 2-coloring, i.e. a graph whose vertices may be partitioned into two nonempty stable sets. A *split graph* is a graph that arises from taking a simple bipartite graph G with $V(G)$ partitioned into nonempty stable sets S_1 and S_2, and switching all of the nonedges in either (but not both) of $G|_{S_1}$ and $G|_{S_2}$ to edges. Thus, the vertices of a split graph may be partitioned (not

necessarily uniquely) into a stable set and a clique. The class of split graphs can also be characterized by the property that they contain no induced subgraph isomorphic to a five-vertex cycle, a four-vertex cycle, or the complement of a four-vertex cycle [14].

There is a natural way noted by Loebl and Sereni [26] to associate to any (possibly non-simple) graph a corresponding simple split graph: given $G = (V, E)$, suppose that $V(G) = \{v_1, ..., v_n\}$ and $E(G) = \{e_1, ..., e_m\}$. The corresponding split graph H has vertex set $V(H) = \{t_1, t_2, ..., t_n, t_{n+1}, ..., t_{n+m}\}$, and edge set $E(H) = \{t_i t_j : 1 \le i < j \le n\} \cup \{t_i t_{n+j}, t_{i'} t_{n+j} : e_j = v_i v_{i'} \text{ in } G\}$. In other words, H is formed by taking the vertices of G, making them into a clique, and then adding a "hat" corresponding to each edge of G. We will denote $H = sp(G)$. Using the above notation, we say that vertex t_{n+j} of $sp(G)$ is the *splitting vertex* of the edge $e_j = v_i v_{i'}$ in G. The construction is illustrated in Figure 4.1.


Figure 4.1: An example of the split graph construction

For any two nonisomorphic graphs, the corresponding split graphs are clearly nonisomorphic, so distinguishing split graphs in some appropriate sense is equivalent to distinguishing all graphs. This motivates considering which functions may

distinguish split graphs; in [26] the authors conjecture that the U-polynomial (or equivalently the bad-coloring chromatic symmetric function XB) does.

Loebl further asked (in personal communication) whether the stronger statement holds that the chromatic symmetric function distinguishes split graphs. Unfortunately it does not, and in particular, the following lemma allows for the construction of many pairs of split graphs that have identical chromatic symmetric functions.

This construction will make use of graph automorphisms. An *automorphism* of a graph G is an isomorphism f of G with itself, and likewise a *w-automorphism* of a vertex-weighted graph (G, w) is a map f that is a w-isomorphism of (G, w) with itself.

Additionally, for $v_1, v_2 \in V(G)$, if $v_1 v_2 \notin E(G)$, we use the shorthand $G \cup v_1 v_2$ to mean the graph $(V(G), E(G) \cup v_1 v_2)$. For brevity if $v \in V(G)$ we also use v to refer to the corresponding vertex of $sp(G)$.

Lemma 21. *Let G be an unweighted graph. Suppose G has (not necessarily distinct) vertices u, u', v, v' such that:*

- *There is no edge of G between u and v, or between u' and v'.*

- *There is some automorphism of G that maps u to u', and some (possibly different) automorphism of G that maps v to v'.*

Then $X_{sp(G \cup uv)} = X_{sp(G \cup u'v')}$.

Proof. Suppose we have G as stated. In $sp(G \cup uv)$, let x be the splitting vertex of uv, and likewise in $sp(G \cup u'v')$ let x' be the splitting vertex of $u'v'$. Applying deletion-contraction to the edge $ux \in sp(G \cup uv)$, we let u^* be the vertex formed by contraction (now with weight 2), and we have

$$X_{sp(G \cup uv)} = X_{sp(G \cup uv) \setminus ux} - X_{sp(G \cup uv) \restriction ux} \tag{4.1}$$

where we are applying simple contraction, so we reduce the multi-edge formed between u^* and v to a single edge. Likewise, applying deletion-contraction to the edge $u'x' \in sp(G \cup u'v')$, we let $u^\#$ be the vertex formed by contraction (now with weight 2), and we have

$$X_{sp(G \cup u'v')} = X_{sp(G \cup u'v') \setminus u'x'} - X_{sp(G \cup u'v') \restriction u'x'} \tag{4.2}$$

where again we are applying simple contraction, so we reduce the multi-edge formed between $u^\#$ and v to a single edge.

Note that if $f : V(G) \to V(G)$ is an automorphism of G, we may extend it to an automorphism of $sp(G)$ by defining that for $z \in V(sp(G)) \setminus V(G)$, if z is the splitting vertex of ab, $f(z)$ is the splitting vertex of $f(a)f(b)$.

Let G_x denote $sp(G \cup uv) \setminus ux$, and let $G_{x'}$ denote $sp(G \cup u'v') \setminus u'x'$. Then the graph G_x has a vertex x that isn't in $G_{x'}$, the graph $G_{x'}$ has a vertex x' that isn't in G_x, and otherwise these graphs have the same vertex set. Similarly, in G_x there is an edge vx that isn't in $G_{x'}$, and in $G_{x'}$ there is an edge $v'x'$ that isn't in G_x, and otherwise these graphs have the same edge set. By hypothesis there is an

69

automorphism f of G with $f(v) = v'$, which may be extended to an automorphism of $sp(G)$ as described above. It is then easy to verify that if we extend f once more to a function $f : V(G_x) \to V(G_{x'})$ by defining $f(x) = x'$, f is an isomorphism of G_x with $G_{x'}$. Since these are unweighted isomorphic graphs, clearly they have the same chromatic symmetric function.

Let G_{u^*} denote $sp(G \cup uv) \nmid ux$ and let $G_{U\#}$ denote $sp(G \cup u'v') \nmid u'x'$. Then the vertex-weighted graph G_{u^*} has a vertex u^* of weight 2 that isn't in $G_{U\#}$, the vertex-weighted graph $G_{U\#}$ has a vertex $u^\#$ of weight 2 that isn't in G_{u^*}, and aside from these exceptions these graphs have the same vertex set, and furthermore all other vertices in both graphs have weight 1. By hypothesis there is an automorphism of G taking u to u' that extends to an automorphism of $sp(G)$. By taking the same f as a function $f : V(G_{u^*}) \to V(G_{U\#})$ with $f(u^*) = u^\#$ (instead of $f(u) = u'$), this f is a w-isomorphism of $sp(G \cup uv) \nmid ux$ and $sp(G \cup u'v') \nmid u'x'$. Since these graphs are w-isomorphic, clearly they have the same chromatic symmetric function.

Thus, comparing (4.1) with (4.2) we see that in the right-hand sides the first terms are identical and the second terms are identical, from which we may conclude that $X_{sp(G \cup uv)} = X_{sp(G \cup u'v')}$.

$\square$

In particular, when $G \cup uv$ is not isomorphic to $G \cup u'v'$, this provides examples of when the chromatic symmetric function fails to distinguish split graphs arising from normal graphs. One way to generate such examples easily is by taking an arbitrary

noncomplete connected graph G, and choosing any nonedge ab in G. Construct G^2 as the disjoint union of graphs G and G^*, where G^* is isomorphic to G. Let $f : V(G) \rightarrow V(G^*)$ be an isomorphism of G and G^*. In the statement of Lemma 21, let $u = u' = a$, $v = b$, and $v' = f(b)$. Then clearly these choices for u, u', v, v' satisfy the lemma, and the two graphs $G^2 \cup uv$ and $G^2 \cup u'v'$ are nonisomorphic, since the latter is connected, and the former is not.

We can also use Lemma 21 to produce two nonisomorphic graphs, both connected, with the same chromatic symmetric function of their split graphs, as shown in Figure 4.2.


Figure 4.2: An unweighted graph G such that $X_{sp(G \cup uv)} = X_{sp(G \cup u'v')}$ with G connected

Note that $G \cup uv$ is not isomorphic to $G \cup u'v'$ because for example, $G \cup uv$ contains a three-vertex cycle, and $G \cup u'v'$ does not.

However, it is worth noting that these examples *are* distinguished by the bad coloring function XB. In fact, Lemma 21 does not generalize to XB since $\nmid$ does not

extend to a simple deletion-contraction relation on XB. If we instead use normal contraction $/$ on the edge ux, we get a multiedge that we must keep between u and v, and likewise for u' and v'. Thus, to generalize Lemma 21 as proven we would need an automorphism of G that takes u to u' and v to v' simultaneously; but then clearly $G \cup uv$ and $G \cup u'v'$ are isomorphic! So this wouldn't provide any counterexamples, and thus the question of whether XB (and thus also the U-polynomial) may be able to distinguish split graphs is still open.

4.2 Graphs with Equal Chromatic Symmetric Function

In the previous section we established a way to generate pairs (S_1, S_2) of simple split graphs such that $X_{S_1} = X_{S_2}$. In this section we consider additional ways to construct pairs (G_1, G_2) of simple graphs with equal chromatic symmetric function, and by combining these constructions with data for small graphs derived via computer program, we conjecture sufficient conditions for a class F of simple graphs to be distinguished by the chromatic symmetric function.

4.2.1 Construction Theorems

In much of the recent literature on the chromatic symmetric function, examples of graphs with equal chromatic symmetric function have been generated using the following result of Orellana and Scott ([32], Theorem 4.2):

Theorem 22. *Let G be a simple graph with distinct vertices v_1, v_2, v_3, v_4 such that*

- *$v_1v_2, v_2v_3, v_3v_4 \in E(G)$, and $v_1v_3, v_1v_4, v_2v_4 \notin E(G)$.*

- *There is an automorphism f of $G \backslash v_2v_3$ such that $f(\{v_1, v_3\}) = \{v_2, v_4\}$ and $f(\{v_2, v_4\}) = \{v_1, v_3\}$.*

Then the graphs $G \cup v_1v_3$ and $G \cup v_2v_4$ have equal chromatic symmetric function.

In addition to Lemma 21 and Theorem 22, we demonstrate one more method for constructing graphs with equal chromatic symmetric function. This method is inspired by the case $u = u'$ of Lemma 21, but can be used in slightly more general contexts and is more akin to Theorem 22.

Given a simple graph G and a vertex $v \in V(G)$, define the *open neighborhood of v* to be $N(v) = \{u \in V(G) : uv \in E(G)\}$ (note that $v \notin N(v)$).

Lemma 23. *Let G be a simple graph, and let v_1, v_2, v_3 be distinct vertices of G satisfying*

- *$v_1v_2 \in E(G)$, $v_1v_3, v_2v_3 \notin E(G)$.*

- *$N(v_3) \subseteq N(v_1) \cap N(v_2)$.*

- *There is an automorphism f of $G \backslash v_3$ such that $f(v_1) = v_2$ and $f(v_2) = v_1$.*

  Then the graphs $G \cup v_1 v_3$ and $G \cup v_2 v_3$ have equal chromatic symmetric functions.

Proof. For clarity of notation we let $e_1 = v_1 v_3$ and $e_2 = v_2 v_3$ be nonedges of G. We apply simple deletion-contraction to each of $G \cup e_1$ and $G \cup e_2$. For the first graph, we let u_1 denote the vertex formed by contraction, and we see that

$$X_{G \cup e_1} = X_G - X_{(G \nmid e_1, w/e_1)}$$

where w is the weight function on $V(G)$ assigning every vertex weight 1. For the second graph we find that

$$X_{G \cup e_2} = X_G - X_{(G \nmid e_2, w/e_2)}.$$

Thus it suffices to show that

$$X_{(G \nmid e_1, w/e_1)} = X_{(G \nmid e_2, w/e_2)}.$$

In fact, we show the stronger statement that $(G \nmid e_1, w/e_1)$ and $(G \nmid e_2, w/e_2)$ are w-isomorphic. Let f be the automorphism of $G \backslash v_3$ that swaps v_1 and v_2 that exists by assumption. We use this to define the map $g : V(G \nmid e_1) \to V(G \nmid e_2)$ by $g(v) = f(v)$ if $v \neq u_1, v_2$, $g(u_1) = u_2$, and $g(v_2) = v_1$ (recall that $v_1 \in V(G)$ was contracted into u_1). Clearly this g is a w-isomorphism if it is an isomorphism. All edge and nonedge relations between vertices of $G \nmid e_1$ other than u_1 and v_2 are preserved in $G \nmid e_2$ by g since they were preserved by f, so it suffices to look at just edges and nonedges involving u_1 and v_2.

74

Clearly $v_2 \in N(u_1)$ and also $g(v_2) = v_1 \in N(u_2) = N(g(u_1))$. Now, let $v \in V(G \restriction e_1)$ be any vertex other than u_1 or v_2. We first consider its edge relation to v_2. By the definition of f we have that $v \in N(v_2)$ if and only if $f(v) \in N(f(v_2)) = N(v_1)$, so also $v \in N(v_2)$ if and only if $g(v) = f(v) \in N(v_1) = N(g(v_2))$. It remains to consider the edge relation between v and u_1.

First, note that the effect of simple contraction means that $v \in N(u_1)$ in $G \restriction e_1$ if and only if $v \in N(v_1) \cup N(v_3)$ in G, and likewise $g(v) \in N(g(u_1))$ if and only if $g(v) = f(v) \in N(v_2) \cup N(v_3)$. Since by assumption $N(v_3) \subseteq N(v_1) \cap N(v_2)$, this simplifies to $v \in N(u_1)$ in $G \restriction e_1$ if and only if $v \in N(v_1)$ in G, and $g(v) \in N(g(u_1))$ in $G \restriction e_2$ if and only if $g(v) = f(v) \in N(f(v_1))$ in G. Since by the definition of f we have $v \in N(v_1)$ if and only if $f(v) \in N(f(v_1))$, we have implied that $v \in N(u_1)$ in $G \restriction e_1$ if and only if $g(v) \in N(g(u_1))$ in $G \restriction e_2$. Thus we have finished showing that g is a w-isomorphism of $G \restriction e_1$ and $G \restriction e_2$, and this concludes the proof.

$\square$

4.2.2 Conjectures and Data

Using deletion-contraction, Spirkl and I created a program to compute X_G and XB_G for small simple graphs (ones with ≤ 8 vertices) using data provided by Brendan McKay [30]. For 1000 pairs of these graphs with equal chromatic symmetric function we also output whether the two graphs in the pair are distinguished by the bad-coloring function XB, and whether the graphs contain as an induced subgraph

- A triangle (K_3).

- A hole (a cycle of more than 3 vertices).

- A claw.

- A net.

- An anti-net (the complement of the net).

Checking for triangles and holes is standard in graph theory, as many well-known graph classes are characterized in part by being triangle-free or free of holes of certain lengths. We check for induced claws since this property is closely related to open conjectures involving the chromatic symmetric function. We check for nets and antinets because the net is one of the smallest graphs that is not e-positive that is not the claw (see Section 3.3), and as was noted previously there is some evidence suggesting that e-positivity may be attained in graphs by forbidding claws and nets [15]. The generated data is far too big to attach directly to this thesis, but it may be found in its entirety online at [8].

Based on this data and also on the construction theorems from the previous section, Spirkl and I conjecture

Conjecture 24. *If G and H are simple, triangle-free graphs, and $X_G = X_H$, then G and H are isomorphic.*

This is a substantial strengthening of the conjecture that the chromatic symmetric function distinguishes trees. We now provide some supporting evidence.

First, some data. In the 1000 pairs of graphs with equal chromatic symmetric function noted in [8], every single graph contains a triangle. Second, every one of the three theorems previously mentioned for constructing graphs with equal chromatic symmetric function always produces a pair of graphs containing triangles. In the case of Lemma 21 and Theorem 22 this is explicit. In the case of Lemma 23, suppose that we have a graph G satisfying the assumptions of the lemma. If $N(v_3) = \emptyset$, then $G \cup v_1 v_3$ is isomorphic to $G \cup v_2 v_3$ since by assumption there is an automorphism of $G \backslash v_3$ swapping v_1 and v_2. If there is a vertex $x \in N(v_3)$, then by assumption also $x \in N(v_1)$ and $x \in N(v_2)$, so in $G \cup v_1 v_3$ there is a triangle with vertices v_1, v_3, x and in $G \cup v_2 v_3$ there is a triangle with vertices v_2, v_3, x. Thus, every G satisfying the conditions of Lemma 23 either produces two isomorphic graphs, or two graphs with equal chromatic symmetric function that both contain triangles.

Thus, studying Conjecture 24 will advance knowledge of the chromatic symmetric function. Obviously if it is proved, it would answer a generalization of a heavily-studied question. If it is disproved, it would likely provide new insight into how graphs may have equal chromatic symmetric function and how to construct such graphs.

Furthermore, seven of the pairs of graphs in [8] also have equal bad-coloring chromatic symmetric function. To the best of my knowledge, these are the first known examples of such graphs. They are examples $259, 539, 546, 635, 656, 848,$ and 909 of [8].

4.3 Generalized Degree Sequences

Let $G = (V, E)$ be a simple graph. For i with $0 \le i \le |V(G)|$, let the set $V_i(G)$ consist of all i-element subsets of $V(G)$. For any fixed $X \in V_i(G)$, let $pair(X) = (a, b)$, where a is the number of edges of G with both endpoints in X, and b is the number of edges of G with exactly one endpoint in X. Finally, let $M_i(G)$ be the multiset $\{pair(X) : X \in V_i(G)\}$. We call the sequence $M_0(G), M_1(G), \ldots, M_{|V(G)|}(G)$ the *generalized degree sequence of G*, and in this section we consider what properties of G may be recovered from its generalized degree sequence.

Considering these multisets arises naturally from research regarding the vertex-weighted chromatic symmetric function. The chromatic symmetric function is not a complete graph invariant, nor is the weighted chromatic symmetric function a complete vertex-weighted graph invariant (see Figure 3.2). However, we may ask whether we can determine an unweighted graph up to isomorphism from the weighted chromatic symmetric functions corresponding to different weight functions $w : V(G) \to \mathbb{Z}^+$.

To formulate precisely a question, suppose we are given that $|V(G)| = d$. Given an integer partition $\lambda = (\lambda_1, \ldots, \lambda_n)$ of n parts, we call a weight function $w : V(G) \to \mathbb{Z}^+$ λ-*compatible* if the multiset $\{w(v) : v \in V(G)\}$ is a permutation of the parts of λ. For S a set of integer partitions of n parts, we define G to be S-*distinguishable* if G is uniquely determined up to isomorphism by the multiset $S(G) = \{X_{(G,w)} : w$ is λ-compatible for some λ in $S\}$. We wish to consider for

which S it is the case that all G with $|V(G)| = n$ are S-distinguishable, in which case $S(G)$ is a complete graph invariant on n-vertex simple graphs.

A first question is whether such an S exists at all, and this can be answered affirmatively. Let $S = \{(2^{n-1}, 2^{n-2}, \ldots, 2, 1)\}$. Then we may always reconstruct G from $S(G)$ as follows: choose any $f = X_{(G,w)}$ from S, and start with n vertices labelled $0, 1, \ldots, n - 1$ and no edges. Let vertex i have weight 2^i. Write f in its p-basis expansion using (3.6). There will be one term $p_{(2^{n-1}, \ldots, 1)}$ corresponding to the empty set $\emptyset \subset E(G)$. For each edge of G, there is one term $-p_\lambda$ where λ has $n - 1$ parts, and if the edge connects vertices i and j, those parts will be $2^i + 2^j$ and $2^k, k \neq i, j$. Since the numbers $2^i + 2^j$, $0 \leq i < j \leq n - 1$ are all unique and not powers of 2, these terms uniquely determine which edges exist between our labelled vertices, so we can reconstruct G.

However, the highest weight used is 2^{n-1}, which is quite large. The next question to ask is whether we can minimize the highest weight used; that is, can we find $m = (\min_S \max_{\lambda \in S} \lambda_1)$ among those S that distinguish n-vertex simple graphs? We first try to determine if it is possible that $m = 2$. Let S' be the set of all λ with n parts where every part is equal to 1 or 2. We now consider whether this S' distinguishes n-vertex simple graphs.

It is in this context that the generalized degree sequence of G appears as a subset of the information we get from $S'(G)$. In particular, if $\lambda^i \in S'$ is the partition with i parts equal to 2 and $n - 2i$ parts equal to 1, we may derive $M_i(G)$ from the multiset

of functions $X_{(G,w)}$ as w ranges over all weight functions that are λ^i-compatible. To see this, label the vertices of the unknown graph G with $v_0, v_1, \ldots, v_{n-1}$ and let w_A be the λ^i-compatible weight function on $V(G)$ that assigns weight 2 to the set of i vertices $A = v_{j_1}, \ldots, v_{j_i}$, and weight 1 to all other vertices. Then expanding $X_{(G,w_A)}$ in the p-basis, we will have one term $p_{2^i 1^{n-2i}}$, and we will also get $-p_\lambda$ for every edge of G where λ has $n-1$ parts. These λ will each be one of three possibilities:

- $4^1 2^{i-2} 1^{n-2i}$, in which case the edge connects two vertices in A.

- $3^1 2^{i-1} 1^{n-2i-1}$, in which case the edge connects a vertex in A to a vertex in $V(G) \backslash A$.

- $2^{i+1} 1^{n-2i-2}$, in which case the edge connects two vertices in $V(G) \backslash A$.

Thus we may derive $pair(A) = (a, b)$ where $-a$ is the coefficient of $p_{4^1 2^{i-2} 1^{n-2i}}$ in $X_{(G,w_A)}$, and $-b$ is the coefficient of $p_{3^1 2^{i-1} 1^{n-2i-1}}$.

Therefore given $S'(G)$, we may convert all functions into their p-basis expansion, and then sort out the multiset of those functions that have a $p_{2^i 1^{n-2i}}$ term. Each of those functions corresponds to w_A with $A \in V_i(G)$, and from the above we may then derive the multiset $M_i(G) = \{pair(A) : A \in V_i(G)\}$. We may derive $M_0(G), \ldots, M_{\lfloor n/2 \rfloor}(G)$ in this way, and since clearly we may derive $M_k(G)$ from $M_{n-k}(G)$, this establishes that the entire generalized degree sequence may be derived from $S'(G)$.

4.3.1 Basic Information

We start by discussing a few immediate properties we can glean from the $M_i(G)$.

Firstly, knowing $M_1(G)$ is equivalent to knowing the degree sequence of G, and so

also the number of edges of G. Secondly, we can determine whether G is connected;

G is disconnected if and only if $(a, 0) \in M_i(G)$ for any a and any $i \neq 0$ or $|V(G)|$.

Furthermore, from the $M_i(G)$ we can determine $girth(G)$, defined to be the length

of the smallest cycle in G, or ∞ if G is acyclic.

Lemma 25.

$$girth(G) = \min\{k : \exists (a, b) \in M_k(G) \ with \ a \geq k\}$$

where the min is evaluated as ∞ if no such k exists.

Proof. We determine $girth(G)$ using the following algorithm. First, we determine

if there is a set of three vertices with three edges between them, that is, if there

exists a pair of the form $(3, b)$ in $M_3(G)$. If so, G has a 3-cycle, otherwise it does

not.

Now, iteratively, suppose we have reached $k < |V(G)|$ such that we have de-

termined that G does not have an l-cycle for all $3 \leq l \leq k$. Then we look to see if

any set A of $k+1$ vertices contains at least $k+1$ edges, that is, if there exists a pair

of the form (a, b) in $M_{k+1}(G)$ with $a \geq k + 1$. If so, this implies that the induced

subgraph $G|_A$ contains at least one cycle; since this cycle cannot be of size $\leq k$, it

must be of size $k + 1$, so this is the girth of G. If no such A exists, we continue

trying with sets of the next higher size. If we reach $|V(G)|$ and have still not found any cycles, then we may conclude that G is acyclic, so the girth of G is ∞. □

It is natural to ask whether the generalized degree sequence is a complete graph invariant, or perhaps at least a complete tree invariant. It turns out it is neither; Spirkl and I computed explicitly that there are counterexamples for graphs with as few as 8 vertices (as in Figure 4.3), and there are counterexamples for trees with as few as 11 vertices (as in Figure 4.4).


Figure 4.3: 8-vertex graphs with the same generalized degree sequence

Figure 4.4: 11-vertex trees with the same generalized degree sequence

4.3.2 Trees

Although studying the generalized degree sequence is more difficult for general simple graphs, it is easier to derive information in the specific case of trees. Note that for a general graph G of n vertices we can determine whether it is a tree by checking if it is connected, and if it has $n-1$ edges. Before continuing we briefly note that the following paragraphs use many graph-theoretic terms; if any are unfamiliar, they are defined in Section 2.2.

For a tree T we define its *trunk* T° to be the smallest connected induced subgraph that contains all vertices of degree ≥ 3. For every leaf (vertex of degree 1) l, we define its *twig* $\hat{l}$ to be the longest path P in T containing l such that every interior

vertex (non-endpoint) of P has degree 2 in T. We call a path of T a twig if it is a twig for one of its leaves. Thus, we may view any tree T as the union of its trunk with its twigs, as shown in Figure 4.5.


Figure 4.5: A tree decomposed into its trunk (with thick edges) and twigs (represented by dashed edges).

Note that for any $A \subseteq V(T)$, if there are $|A| - 1$ edges between the vertices of A, then the induced subgraph $T|_A$ is connected. This is because if you start with the vertices of $T|_A$ and insert the edges one at a time, each new edge decreases the number of connected components of $T|_A$ since no cycle is ever created, so inserting $|A| - 1$ edges reduces to a single connected component.

Lemma 26. *Given the generalized degree sequence of T we can determine*

- *The size of T°.*

- *The multiset of lengths of the twigs of T.*

In particular, spiders (trees with exactly one vertex of degree ≥ 3) are uniquely determined by their generalized degree sequence.

Proof. Since we can determine the degree sequence of T, we know how many leaves T has, and thus how many twigs; let a be this number. First, we show that for any connected induced subgraph S of T, it is impossible for S to have more than a incident edges in T (edges with one endpoint in S and one endpoint in $T \backslash S$). Suppose otherwise, and let $v_1, \ldots, v_{a+1}$ be $a+1$ vertices that are endpoints of such edges in $T \backslash S$. By construction each of these vertices is in a different connected component of $T \backslash S$, and each of these (nonempty) components must contain a leaf of T, a contradiction. Furthermore, the only way a connected induced subgraph S of T can have exactly a incident edges is if it contains every vertex of degree ≥ 3. Suppose otherwise, that we have such a subgraph S with a incident edges that doesn't contain some particular vertex x of degree ≥ 3. There is exactly one path connecting x to S; in addition to the edge starting this path at x, there are ≥ 2 further edges coming out of x, each of which leads to at least one leaf of T. Hence the connected component of $T \backslash S$ containing x has at least two leaves, and each of the other $a - 1$ connected components has at least one leaf, again a contradiction.

Thus, the smallest connected induced subgraph of T with a incident edges contains all of T°, and since clearly T° itself must have exactly a incident edges (one to start each twig), we conclude that T° is the smallest connected induced subgraph of T having a incident edges.

Thus, we can determine the number of vertices of T° to be the smallest number k such that there is a set $X \in V_k(G)$ with $pair(X) = (k - 1, a)$.

Let $|V(T°)| = C + 1$, and thus $|E(T°)| = C$. We define c_i, for $1 \leq i \leq n$, to be the number of twigs of T of length i. We will determine the c_i in ascending order of index by an iterative process. For the first step, we evaluate the number of $X \in V_{C+2}$ with $pair(X) = (C + 1, a)$. For each such X, the corresponding induced subgraph $T|_X$ is connected, and by the above arguments $T|_X$ must contain $T°$, as well as one additional edge that starts a twig. Each of the a possible twig-starting edges will work except for those edges that form an entire twig of length 1. Thus $|\{X \in V_{C+2} : pair(X) = (C + 1, a)\}| = a - c_1$, and we may determine c_1.

For the k^{th} step, let the twigs be $t_1, \ldots, t_a$, with lengths $l_1, \ldots, l_a$. Reorder so that for some m we have determined the lengths $l_1 \leq \cdots \leq l_m \leq k - 1$ (equivalently we have determined $c_1, \ldots, c_{k-1}$), and we do not yet know the values $l_{m+1}, \ldots, l_a$, but we do know that all of these are $\geq k$. We determine the number of $X \in V_{C+k+1}$ with $pair(X) = (C + k, a)$. For each such X, the corresponding induced subgraph $T|_X$ will be connected with a incident edges. The set of these $T|_X$ is exactly the set of those connected induced subgraphs of T that contain $T°$ as well as k additional edges, without containing all the edges of any twig (as otherwise it would have fewer than a incident edges). We enumerate these $T|_X$ by splitting into cases based on whether all k edges come from the same twig. First, we can determine the number N of ways to form such a $T|_X$ without choosing k edges from the same twig, because we can use up to $k - 1$ edges from each of $t_{m+1}, \ldots, t_a$, and up to $l_i - 1$ edges from each $t_i, 1 \leq i \leq m$, so N only depends on k and $c_1, \ldots, c_{k-1}$,

all of which are known at this step. Second, the number of ways to use exactly k edges from one twig will be $a - m - c_k$, where c_k is the number of twigs among $t_{m+1}, \ldots, t_a$ with length exactly k, and $m = c_1 + \cdots + c_{k-1}$. Thus, we can evaluate $|\{X \in V_{C+k+1} : pair(X) = (C+k, a)\}| = N + a - c_1 - \cdots - c_k$, and since we know $a, c_1, \ldots, c_{k-1}$, and can determine N only from the information we learned from prior steps, we can determine c_k. We iterate this process until all twig lengths are known.

$\square$

We conclude this section with some additional points that may be useful for future research. First, note that in any connected induced subgraph $T|_X$ of a tree T with k vertices, its number of incident edges is equal to $\left(\sum_{x \in X} d(x)\right) - 2|X| + 2$, or rewritten, $2 - |\{x \in X : x \text{ is a leaf}\}| + \sum_{d(x) \geq 3}(d(x) - 2)$, which can help enumerate the possibilities for connected induced subgraphs with a prescribed number of incident edges.

Another possible use of the $M_i(T)$ for trees is to sort by number of incident edges. Note that for every subset $Y \subseteq E(T)$, there is exactly one way to partition $V(T)$ into two nonempty sets A and B such that the set of edges contributing to the second coordinate of $pair(A)$ and $pair(B)$ (their number of incident edges) is exactly Y, and furthermore, any $X \subseteq V(T)$ having Y as its exact set of incident edges must be one of A or B.

This can give us a different perspective on the information from the M_i. For instance, we can rewrite the information to just be of the form (v, a, b) over all $X \subseteq V(T)$, with $v = |X|$ as a parameter of our tuple. As an example of information we could derive using this, for fixed b, consider an arbitrary X that produces a tuple of the form $(k, k-1, b)$. Then $T|_X$ is an induced connected subgraph, and if the set of b incident edges is denoted B, then $T \backslash B$ contains $b + 1$ connected components. If we take each of these components as vertices of a graph with edge set B, this new graph must be a star (a spider in which every twig has length 1). Thus the number of $X \subseteq V(T)$ that induce tuples $(k, k-1, b)$ is equivalent to how many sets B of b edges exist such that every path between two edges of B does not use a third edge of B. This sort of structural information may be helpful, especially if we are looking at specific types of tree (e.g. caterpillars, trees T in which T° is a path).

4.4   A New Basis for Λ

For any $\lambda = (\lambda_1, \ldots, \lambda_k)$ an integer partition, define G_λ to be the simple graph of $|\lambda|$ vertices labelled $v_{11}, \ldots, v_{1\lambda_1}, v_{21}, \ldots, v_{2\lambda_2}, v_{31}, \ldots, v_{k\lambda_k}$, and such that there is an edge between v_{ij} and v_{ab} if and only if $i \neq a$. Thus, the graph G_λ consists of k stable sets of sizes $\lambda_1, \ldots, \lambda_k$, with all possible edges between them. We define the symmetric function

$$b_\lambda = X_{G_\lambda}$$

Note that $X_{\overline{G_\lambda}} = (\prod_{i=1}^k \lambda_i!)e_\lambda$, so the b_λ are in some sense the graph complement of the e_λ.

Recall that $r_i(\lambda)$ denotes the number of parts of λ equal to i. We define an augmentation of the m-basis by the relation

$$\widetilde{m_\lambda} = \left(\prod_i r_i(\lambda)!\right) m_\lambda.$$

Letting $Stab_\lambda(G)$ be the set of (unordered) partitions of $V(G)$ into k stable sets of sizes $\lambda_1, \ldots, \lambda_k$, recall from (3.2) that for any graph G

$$X_G = \sum_{\lambda \vdash |V(G)|} |Stab_\lambda(G)| \widetilde{m_\lambda}.$$

In considering the $\widetilde{m}$-basis, we will sometimes work in a slightly modified version of the ring of symmetric functions denoted by $\widetilde{\Lambda}$ in which we retain the same addition operation but use a different multiplication operation $\otimes$ defined by $\widetilde{m_\lambda} \otimes \widetilde{m_\mu} = \widetilde{m_{\lambda \sqcup \mu}}$, where $\lambda \sqcup \mu$ is the partition whose parts are the disjoint union of the parts of λ and μ, e.g. $(3,1,1) \sqcup (2,1) = (3,2,1,1,1)$. This ring and the associated multiplication were introduced and studied by Tsujie in [43].

The action of the operation $\otimes$ on the $\widetilde{m}$-basis has an interpretation in terms of the *join* of two graphs G and H, defined as the graph $G \otimes H$ with vertex set $V(G) \sqcup V(H)$ and edge set $E(G) \sqcup E(H) \sqcup \{(v,w) : v \in G, w \in H\}$ (where here all $\sqcup$s mean disjoint union). In what follows, given a symmetric function g and a symmetric function basis $\{f_\lambda\}$ indexed by partitions λ, we use the notation $[f_\mu]g$ to mean the coefficient of f_μ when expressing g in the f-basis.

89

Lemma 27. *Let G, H be graphs. Then*

$$X_G \otimes X_H = X_{G \otimes H}$$

Proof. It suffices to show that both sides have the same coefficient of $\widetilde{m_\lambda}$ for all

$\lambda \vdash |V(G)| + |V(H)|$. For any fixed λ, any stable λ-partition of $G \otimes H$ must be of

the form $A \sqcup B$, where A is some partition of $V(G)$ consisting whose parts are a

submultiset of the multiset of the parts of λ, and B is a partition of $V(H)$ whose

parts are the remainder of λ. Thus the number of stable λ-partitions of $G \otimes H$ and

thus the desired coefficient is

$$[\widetilde{m_\lambda}]X_{G \otimes H} = \sum_{\mu \subseteq \lambda} [\widetilde{m_\mu}]X_G [\widetilde{m_{\lambda \backslash \mu}}]X_H$$

where $\mu \subseteq \lambda$ means that the parts of μ are a submultiset of the parts of λ, and thus

$\lambda \backslash \mu$ denotes the partition comprising the parts of λ with the parts of μ removed.

Clearly this is also equal to

$$[\widetilde{m_\lambda}](X_G \otimes X_H)$$

so we are done.

$\square$

Note that by the definition of the b_λ, we have

$$b_\lambda \otimes b_\mu = b_{\lambda \sqcup \mu}.$$

Since $b_n = p_{1^n}$ is the chromatic symmetric function of the graph of n vertices with no edges, we have the formula

$$b_n = \sum_{\mu \vdash n} \frac{n!}{\prod_i \mu_i! \prod_i r_i(\mu)!} \widetilde{m_\mu}$$

and in the modified ring $\widetilde{\Lambda}$, we thus have

$$b_\lambda = b_{\lambda_1} \otimes b_{\lambda_2} \otimes \cdots \otimes b_{\lambda_k} =$$

$$\left(\sum_{\mu^1 \vdash \lambda_1} \frac{\lambda_1!}{\prod_i \mu_i^1! \prod_i r_i(\mu^1)!} \widetilde{m_{\mu^1}} \right) \otimes \cdots \otimes \left(\sum_{\mu^k \vdash \lambda_k} \frac{\lambda_k!}{\prod_i \mu_i^k! \prod_i r_i(\mu^k)!} \widetilde{m_{\mu^k}} \right) \qquad (4.3)$$

Recall that a *puzzle* of $\mu \to \lambda = (\lambda_1, \ldots, \lambda_k)$ is an ordered tuple of partitions $(\mu^1, \ldots, \mu^k)$ such that

- For all $1 \le i \le k$ we have $\mu^i \vdash \lambda_i$.

- The disjoint union of the parts of the μ^i is μ.

In terms of puzzles, we may extract the coefficient of $\widetilde{m_\mu}$ in the expansion of b_λ by using (4.3) as

$$[\widetilde{m_\mu}]b_\lambda = \frac{\prod_i \lambda_i!}{\prod_i \mu_i!} \sum_{\text{puzzles } \mu \to \lambda} \frac{1}{\prod_{i,j} r_i(\mu^j)!} \qquad (4.4)$$

Note that this value is 0 if the sum is empty, so $[\widetilde{m_\mu}]b_\lambda = 0$ when μ is not a refinement of λ. Therefore the linear transformation mapping the $\widetilde{m}$-basis to the b-basis is (with respect to the reverse lexicographic order on partitions) an upper triangular matrix with 1s on the main diagonal, so it is invertible. Thus, for each

d, the set $\{b_\lambda : \lambda \vdash d\}$ is a basis for $\widetilde{\Lambda}^d$ and so also for Λ^d. Furthermore, passing back to the usual m-basis, we have

$$[m_\mu]b_\lambda = \frac{\prod_i \lambda_i!}{\prod_i \mu_i!} \sum_{\text{puzzles } \mu \to \lambda} \frac{\prod_i r_i(\mu)!}{\prod_{i,j} r_i(\mu^j)!} = \frac{\prod_i \lambda_i!}{\prod_i \mu_i!} R_{\mu\lambda} \tag{4.5}$$

where the $R_{\mu\lambda}$ are the entries of the change-of-basis matrix from the m-basis to the p-basis as given in (2.1). Thus the b-basis fits naturally into the framework of the standardly used symmetric function bases. Furthermore, we may use (4.5) to prove an inverse relationship which has a natural combinatorial interpretation.

First, we need a simple expression for the change-of-basis coefficients from the p-basis to the h-basis. We may verify (e.g. as in [40]) that

$$[h_\mu]p_n = n \sum_{\mu \vdash n} \frac{(-1)^{l(\mu)-1}(l(\mu)-1)!}{\prod_i r_i(\mu)!}$$

Extending this to general p_λ by multiplication gives

$$[h_\mu]p_\lambda = \left(\prod_j \lambda_j\right) \sum_{\text{puzzles } \mu \to \lambda} \prod_{\mu_i} \frac{(-1)^{l(\mu_i)-1}(l(\mu_i)-1)!}{\prod_j r_j(\mu_i)!} \tag{4.6}$$

Theorem 28.

$$m_n = \sum_{\mu \vdash n} (-1)^{l(\mu)-1} c_\mu b_\mu \tag{4.7}$$

where

$$c_\mu = \frac{n!(l(\mu)-1)!}{\prod_i \mu_i! \prod_i r_i(\mu)!}$$

is the number of ways to make a circular necklace with distinguishable beads $B_1, \ldots, B_{l(\mu)}$ such that

- *Each bead B_i contains a nonempty subset $S(B_i)$ of $\{1, 2, \ldots, n\}$.*

- *The $S(B_i)$ form a partition of $\{1, 2, \ldots, n\}$.*

- *The multiset $\{|S(B_1)|, \ldots, |S(B_{l(\mu)})|\}$ is exactly the multiset of parts of μ.*

Proof. By inverting (4.5) we have that

$$[b_\lambda]m_\mu = \frac{\prod_i \mu_i!}{\prod_i \lambda_i!} R_{\mu\lambda}^{-1}.$$

Letting $\mu = (n)$ and comparing to the claimed equation (4.7), it suffices to show that

$$\frac{n!}{\prod_i \lambda_i!} R_{(n)\lambda}^{-1} = n! \frac{(-1)^{l(\lambda)-1}(l(\lambda)-1)!}{\prod_i \lambda_i! \prod_i r_i(\lambda)!}$$

or after simplifying, that

$$R_{(n)\lambda}^{-1} = \frac{(-1)^{l(\lambda)-1}(l(\lambda)-1)!}{\prod_i r_i(\lambda)!}. \tag{4.8}$$

We will show (4.8) directly using the definition of $R_{\mu\lambda}$ as the change-of-basis coefficients from p to m given in (2.1) and the well-known Cauchy identity [27, 40]

$$\prod_{i,j \geq 1} \frac{1}{1 - x_i y_j} = \sum_{\mu} m_\mu(x) h_\mu(y) = \sum_{\lambda} \frac{p_\lambda(x) p_\lambda(y)}{\prod_i i^{r_i(\lambda)} r_i(\lambda)!}. \tag{4.9}$$

where x and y are to be interpreted as shorthand for the countably many variables $x_1, x_2, \ldots$ and $y_1, y_2, \ldots$ respectively. First, we expand the $p_\lambda(y)$ in the right-hand sum of (4.9) into the h-basis using (4.6). Then we consider the middle and right-most sums of (4.9) as symmetric functions in the y variables with coefficients in $\mathbb{R}[[x_1, x_2, \ldots]]$, expand them in the h-basis, and equate the coefficients of $h_\mu(y)$ on both sides. Then we obtain

$$m_\mu(x) = \sum_{\lambda} \frac{p_\lambda(x)}{\prod_i i^{r_i(\lambda)} r_i(\lambda)!} \left(\prod_j \lambda_j \right) \sum_{\text{puzzles } \mu \to \lambda} \prod_{\mu_i} \frac{(-1)^{l(\mu_i)-1}(l(\mu_i)-1)!}{\prod_j r_j(\mu_i)!}.$$

Since by definition $[p_\lambda]m_\mu$ is equal to the corresponding entry of $R_{\lambda\mu}^{-1}$, we thus have

$$R_{\lambda\mu}^{-1} = \frac{1}{\prod_i i^{r_i(\lambda)} r_i(\lambda)!} \left(\prod_j \lambda_j \right) \sum_{\text{puzzles } \mu \to \lambda} \prod_{\mu_i} \frac{(-1)^{l(\mu_i)-1}(l(\mu_i)-1)!}{\prod_j r_j(\mu_i)!}.$$

Setting $\lambda = (n)$ gives exactly (4.8), so we are done.

$\square$

Since $m_n = \widetilde{m_n}$, this result may be used along with multiplication in $\widetilde{\Lambda}$ to compute the values $[b_\lambda]\widetilde{m_\mu}$.

These relationships may be used to derive identities involving partitions and symmetric functions. As an example of one such derivable identity, we compute $[b_\mu]m_n$ in another way. Using only the formula for $[m_\lambda]b_n$ and rearranging we may write

$$\widetilde{m_n} = b_n - \sum_{\lambda \neq (n)} \frac{n!}{\prod_i \lambda_i! \prod_i r_i(\lambda)!} \widetilde{m_\lambda} =$$

$$b_n - \sum_{\lambda \neq (n)} \frac{n!}{\prod_i \lambda_i! \prod_i r_i(\lambda)!} \left(\sum_{\mu_1 \vdash \lambda_1} (-1)^{l(\mu_1)-1} \frac{\lambda_1!(l(\mu_1)-1)!)}{\prod_i \mu_{1i}! \prod_i r_i(\mu_1)!} b_{\mu_1} \right) \cdots$$

$$\left(\sum_{\mu_k \vdash \lambda_k} (-1)^{l(\mu_k)-1} \frac{\lambda_k!(l(\mu_k)-1)!)}{\prod_i \mu_{ki}! \prod_i r_i(\mu_k)!} b_{\mu_k} \right)$$

This tells us that $[b_n]\widetilde{m_n} = 1$, and that for $\mu \neq (n)$ we have

$$[b_\mu]\widetilde{m_n} = - \sum_{\lambda \neq (n)} \frac{n!}{\prod_i \lambda_i! \prod_i r_i(\lambda)!} \sum_{\text{puzzles } \mu \to \lambda} (-1)^{l(\mu)-l(\lambda)} \frac{\prod_i \lambda_i! \prod_i (l(\mu^i)-1))!}{\prod_{i,j} \mu_j^i! \prod_{i,j} r_i(\mu^j)!}$$

Comparing to (4.7), for $\mu \neq (n)$ we have

$$(-1)^{l(\mu)-1} \frac{n!(l(\mu)-1)!}{\prod_i \mu_i! \prod_i r_i(\mu)!} =$$

94

$$-\sum_{\lambda \neq (n)} \frac{n!}{\prod_i \lambda_i! \prod_i r_i(\lambda)!} \sum_{\text{puzzles } \mu \to \lambda} (-1)^{l(\mu)-l(\lambda)} \frac{\prod_i \lambda_i! \prod_i (l(\mu^i)-1)!}{\prod_{i,j} \mu_{ij}! \prod_{i,j} r_i(\mu^j)!}$$

After cancelling common terms from both sides, as well as some common terms from the numerator and denominator of the right-hand side, we reduce to

$$\frac{(l(\mu)-1)!}{\prod_i r_i(\mu)!} = \sum_{\lambda \neq (n)} \frac{(-1)^{l(\lambda)}}{\prod_i r_i(\lambda)!} \sum_{\text{puzzles } \mu \to \lambda} \frac{\prod_{i=1}^{l(\lambda)}(l(\mu^i)-1)!}{\prod_{i,j} r_i(\mu^j)!} \qquad (4.10)$$

We note that the left-hand side of (4.10) is just -1 times the missing $\lambda = (n)$ case of the sum on the right-hand side, so by subtracting it from both sides we derive the identity

$$\sum_{\lambda \vdash n} \frac{(-1)^{l(\lambda)}}{\prod_i r_i(\lambda)!} \sum_{\text{puzzles } \mu \to \lambda} \frac{\prod_{i=1}^{l(\lambda)}(l(\mu^i)-1)!}{\prod_{i,j} r_i(\mu^j)!} = 0. \qquad (4.11)$$

It is worth noting the specific case of $\mu = 1^n$, which simplifies to

$$\sum_{\lambda \vdash n} \frac{(-1)^{l(\lambda)}}{\prod_i \lambda_i \prod_i r_i(\lambda)!} = 0. \qquad (4.12)$$

This equality may also be proved by a generating function argument:

$$\sum_{\text{all } \lambda} \frac{x^{|\lambda|} y^{l(\lambda)}}{\prod_i \lambda_i! \prod_i r_i(\lambda)!} = \sum_{(a_1, a_2, \dots) \in \mathbb{N}^\infty} \frac{x^{\sum i a_i} y^{\sum a_i}}{\prod_{k \geq 1} a_k! k^{a_k}} =$$

$$\sum_{(a_1, a_2, \dots) \in \mathbb{N}^\infty} \prod_{i \geq 1} \frac{x^{i a_i} y^{a_i}}{a_i! i^{a_i}} = \prod_{i \geq 1} \exp\left(\frac{x^i y}{i}\right) = \exp\left(\sum_{i \geq 1} \frac{x^i y}{i}\right) =$$

$$\exp\left(-y \log(1-x)\right) = (1-x)^{-y}$$

Looking at the initial generating function, clearly the left-hand side of (4.12) is the coefficient of x^n when $y = -1$. But from the simplified form of the equation it is clear that this coefficient is always 0 other than in the exceptional case $n = 1$.

Furthermore, this generating function identity itself gives a family of related results. For example, by setting $y = 1$ and looking at the coefficient of x^n, we derive the identity that for any n,

$$\sum_{\lambda \vdash n} \frac{1}{\prod_i \lambda_i \prod_i r_i(\lambda)!} = 1 \tag{4.13}$$

A more familiar form of (4.13) may be recovered by multiplying both sides by $n!$, which produces an equation that is equivalent to enumerating the permutations of S_n by the their cycle type [40].

In a similar vein, numerical computation of small examples suggests

Conjecture 29.

$$\sum_{\lambda \vdash n} \frac{1}{\prod_i r_i(\lambda)!} \sum_{\text{puzzles } \mu \to \lambda} \frac{\prod_{i=1}^{l(\lambda)}(l(\mu^i) - 1)!}{\prod_{i,j} r_i(\mu^j)!} = \frac{l(\mu)!}{\prod_i r_i(\mu)!} \tag{4.14}$$

This is the unsigned analog of (4.11). Notably, the right-hand side of (4.14) is the number of compositions of n whose parts are those of μ.

Bibliography

[1] Alexandersson, Per, and Greta Panova. "LLT polynomials, chromatic quasisymmetric functions and graphs with cycles." Discrete Mathematics 341.12 (2018): 3453–3482.

[2] Aliste-Prieto, José, Anna de Mier, and José Zamora. "On trees with the same restricted U-polynomial and the Prouhet-Tarry-Escott problem." Discrete Mathematics 340.6 (2017): 1435–1441.

[3] Birkhoff, George D. "A determinant formula for the number of ways of coloring a map." The Annals of Mathematics 14.1/4 (1912): 42–46.

[4] Cho, Soojin, and JiSun Huh. "On e-positivity and e-unimodality of chromatic quasisymmetric functions." arXiv preprint arXiv:1711.07152 (2017).

[5] Cho, Soojin, and Stephanie van Willigenburg. "Chromatic classical symmetric functions." arXiv preprint arXiv:1604.03954 (2016).

[6] Chow, Timothy Y. "The path-cycle symmetric function of a digraph." Advances in Mathematics 118.1 (1996): 71–98.

[7] Crew, Logan, and Sophie Spirkl. "A Deletion-Contraction Relation for the Chromatic Symmetric Function." arXiv preprint arXiv:1910.11859 (2019).

[8] Crew, Logan, and Sophie Spirkl. "Examples of Pairs of Graphs of Small Graphs with Equal Chromatic Symmetric Function". Logan Crew's Homepage. Google Sites, https://sites.google.com/site/ltcrewmath/exampl.pdf. PDF download.

[9] Crew, Logan, and Sophie Spirkl. "Using Deletion-Contraction to Construct Graphs with Equal Chromatic Symmetric Function". In preparation.

[10] Dahlberg, Samantha, Adrian She, and Stephanie van Willigenburg. "Schur and e-positivity of trees and cut vertices." arXiv preprint arXiv:1901.02468 (2019).

[11] Dahlberg, Samantha, and Stephanie van Willigenburg. "Lollipop and lariat symmetric functions." SIAM Journal on Discrete Mathematics 32.2 (2018): 1029–1039.

[12] Ellzey, Brittney. "A directed graph generalization of chromatic quasisymmetric functions." arXiv preprint arXiv:1709.00454 (2017).

[13] Fayers, Matthew. "A note on Kostka numbers." arXiv preprint arXiv:1903.12499 (2019).

[14] Földes, Stèphane, and Peter L. Hammer. "Split graphs." Congressus Numerantium 19 (1977): 311–315.

[15] Foley, Angèle M., Chính T. Hoàng, and Owen D. Merkel. "Classes of graphs with e-positive chromatic symmetric function." Electronic Journal of Combinatorics 26.3 (2019): P3–51.

[16] Gasharov, Vesselin. "Incomparability graphs of (3+1)-free posets are s-positive." Discrete Mathematics 157.1–3 (1996): 193–197.

[17] Gebhard, David D., and Bruce E. Sagan. "A chromatic symmetric function in noncommuting variables." Journal of Algebraic Combinatorics 13.3 (2001): 227–255.

[18] Grinberg, Darij. "Double Posets and the Antipode of QSym." Electronic Journal of Combinatorics 24.2 (2017): P2–22.

[19] Guay-Paquet, Mathieu. "A modular relation for the chromatic symmetric functions of (3+1)-free posets." arXiv preprint arXiv:1306.2400 (2013).

[20] Hamel, Angèle M., Chính T. Hoàng, and Jake E. Tuero. "Chromatic symmetric functions and H-free graphs." Graphs and Combinatorics 35.4 (2019): 815–825.

[21] Hasebe, Takahiro, and Shuhei Tsujie. "Order quasisymmetric functions distinguish rooted trees." Journal of Algebraic Combinatorics 46.3–4 (2017): 499–515.

[22] Heil, Sam, and Caleb Ji. "On an Algorithm for Comparing the Chromatic Symmetric Functions of Trees." arXiv preprint arXiv:1801.07363 (2018).

[23] Kaliszewski, Ryan. "Hook coefficients of chromatic functions." arXiv preprint arXiv:1404.7531 (2014).

[24] Klazar, Martin, Martin Loebl, and Iain Moffatt. "The Potts model and chromatic functions of graphs." arXiv preprint arXiv:1311.4348 (2013).

[25] Lenart, Cristian, and Nigel Ray. "Chromatic polynomials of partition systems." Discrete Mathematics 167 (1997): 419–444.

[26] Loebl, Martin, and Jean-Sébastien Sereni. "Isomorphism of weighted trees and Stanley's conjecture for caterpillars." arXiv preprint arXiv:1405.4132 (2014).

[27] Macdonald, Ian Grant. "Symmetric functions and Hall polynomials." Oxford University Press (1998).

[28] Martin, Jeremy L., Matthew Morin, and Jennifer D. Wagner. "On distinguishing trees by their chromatic symmetric functions." Journal of Combinatorial Theory, Series A 115.2 (2008): 237–253.

[29] Martin, Jeremy L., and Jennifer D. Wagner. "On the chromatic symmetric function of a tree." Extended abstract for 18th Annual Conference on Formal Power Series and Algebraic Combinatorics, San Diego. 2006.

[30] McKay, Brendan. "Combinatorial Data". Brendan McKay's Home Page. Australian National University, users.cecs.anu.edu.au/~bdm/data/.

[31] Noble, Steven D., and Dominic J. A. Welsh. "A weighted graph polynomial from chromatic invariants of knots." Annales de l'institut Fourier. Vol. 49. No. 3. 1999.

[32] Orellana, Rosa, and Geoffrey Scott. "Graphs with equal chromatic symmetric functions." Discrete Mathematics 320 (2014): 1–14.

[33] Paunov, Alexander. "Planar graphs and Stanley's Chromatic Functions." arXiv preprint arXiv:1702.05787 (2017).

[34] Pawlowski, Brendan. "Chromatic symmetric functions via the group algebra of S_n." arXiv preprint arXiv:1802.05470 (2018).

[35] Ray, Nigel, and Colin Wright. "Umbral interpolation and the addition/contraction tree for graphs." Discrete Mathematics 103.1 (1992): 67–74.

[36] Sagan, Bruce E. "The symmetric group: representations, combinatorial algorithms, and symmetric functions." Vol. 203. Springer Science & Business Media, 2013.

[37] Stanley, Richard P. "Acyclic orientations of graphs." Discrete Mathematics 5.2 (1973): 171–178.

[38] Stanley, Richard P. "A symmetric function generalization of the chromatic polynomial of a graph." Advances in Mathematics 111.1 (1995): 166–194.

[39] Stanley, Richard P. "Enumerative Combinatorics Volume 1 second edition." Cambridge Studies in Advanced Mathematics (2011).

[40] Stanley, Richard P. "Enumerative Combinatorics, Vol. 2." Cambridge Studies in Advanced Mathematics (1999).

[41] Stanley, Richard P. "Graph colorings and related symmetric functions: ideas and applications. A description of results, interesting applications, & notable open problems." Discrete Mathematics 193.1–3 (1998): 267–286

[42] Shareshian, John, and Michelle L. Wachs. "Chromatic quasisymmetric functions." Advances in Mathematics 295 (2016): 497–551.

[43] Tsujie, Shuhei. "The chromatic symmetric functions of trivially perfect graphs and cographs." Graphs and Combinatorics 34.5 (2018): 1037–1048.

[44] Wang, David G.L., and Monica M.Y. Wang. "Non-Schur-positivity of chromatic symmetric functions." arXiv preprint arXiv:2001.00181 (2020).

[45] Wolfgang, Harry Lewis. "Two interactions between combinatorics and representation theory: monomial immanants and Hochschild cohomology." Diss. Massachusetts Institute of Technology, 1997.