

Bijjective proofs of proper coloring theorems

Bruce E. Sagan

Department of Mathematics, Michigan State University,
East Lansing, MI 48824, USA, sagan@math.msu.edu

Vincent Vatter*

Department of Mathematics, University of Florida,
Gainesville, FL 32601, USA, vatter@ufl.edu

October 15, 2019

Key Words: acyclic orientation, broken circuit, chromatic polynomial, chromatic symmetric function, proper coloring

AMS subject classification (2010): 05C31, 05E05 (Primary) 05C15 (Secondary)

Abstract

The chromatic polynomial and its generalization, the chromatic symmetric function, are two important graph invariants. Celebrated theorems of Birkhoff, Whitney, and Stanley show how both objects can be expressed in three different ways: as sums over all subgraphs, as sums over subgraphs with no broken circuits, and in terms of acyclic orientations with compatible colorings. We establish all six of these expressions bijectively. In fact, we do this with only two bijections, as the proofs in the symmetric function setting are obtained using the same bijections as in the polynomial case and the bijection for broken circuits is just a restriction of the one for all subgraphs.

1 Introduction

First defined by Birkhoff in 1912 for planar graphs [4], the chromatic polynomial failed at its inventor's original purpose, the proof of what is now the four color theorem. Indeed, in a 1946 survey on the chromatic polynomial, Birkhoff and Lewis divided attacks on this then-conjecture into two types, qualitative (Type 1) and quantitative (Type 2), placing the study of chromatic polynomials at the center of the efforts of Type 2. Even then, the outlook was not rosy, as they admitted [5, p. 357] (emphasis in original):

These researches have not been as successful as the researches of Type 1 in yielding results that are directly connected with the four-color problem. It is certain that the greater generality of the problem here considered has introduced

*Vatter's research was partially supported by the Simons Foundation via award number 636113.




Figure 1: A graph, two colorings, and a (spanning) subgraph.

complications which have *so far* rendered the solution of the classical four-color problem more remote by the methods characteristic of Type 2 than it is by the methods of Type 1.

Despite its failure in leading to a proof of the four color theorem, the chromatic polynomial has nevertheless become a fundamental object in graph theory. We refer to Read's still-excellent 1968 survey [10] or the second part of Biggs' *Algebraic Graph Theory* [3] for more background on the chromatic polynomial and its importance. Our aim here is to explain, via simple bijections, three classic expressions for the chromatic polynomial—as a sum over spanning subgraphs, as a sum over spanning subgraphs without broken circuits, and in terms of acyclic orientations together with compatible colorings. In each case, we show how these bijections can also be used to prove analogous results for the chromatic symmetric function, a generalization of the chromatic polynomial defined by Stanley [13] in 1995. First, though, we need definitions.

Let $G = (V, E)$ be a graph with vertices V and edges E . A *coloring* of G is a function κ from V to a set of *colors*. Here our colors are exclusively the positive integers \mathbb{P} and initial segments thereof, for which we use the notation $[t] = \{1, 2, \dots, t\}$. Thus a \mathbb{P} -*coloring* of G is a function $\kappa : V \rightarrow \mathbb{P}$ and a $[t]$ -*coloring* is a function $\kappa : V \rightarrow [t]$. An example of a graph is shown on the left of Figure 1 and two $[4]$ -colorings of it are shown in the center.

The edge $e = uv \in E$ is *monochromatic* in the coloring κ if $\kappa(u) = \kappa(v)$, and a coloring is *proper* if none of its edges is monochromatic. Note that only the second coloring shown in Figure 1 is proper. The *chromatic polynomial* of the graph G is defined by

$$\chi(G; t) = \text{the number of proper } [t]\text{-colorings of } G$$

for all positive integers t . Although it is not obvious from this definition, the chromatic polynomial is indeed a polynomial in t . In fact, this was the main result of Birkhoff's seminal paper [4], and he established it by proving the result we present as Theorem 2.1.

A *spanning subgraph* of a graph $G = (V, E)$ is a graph on the same set of vertices but with a subset of edges. As all of the subgraphs used in this work are spanning, we simply call them subgraphs. For example, the graph on the right in Figure 1 is a subgraph of the graph on the left. We associate each subset $S \subseteq E$ of edges with the (spanning) subgraph of G with edge set S , so the subgraph on the right of Figure 1 is denoted by $S = \{uv, vx\}$. A *connected component*, or simply *component*, K of a subgraph S is a maximal subset of vertices such that one can reach any vertex in K from any other by traveling vertex-to-vertex via the edges of S . The components of any graph (or subgraph) form a partition of its vertices. Returning to the subgraph on the right of Figure 1, we see that it has two components, one

with vertex set $\{u, v, w\}$ and the other consisting of the single vertex x . We let

$$c(S) = \text{the number of components of } S.$$

In our proofs we frequently consider certain *improper* colorings. Given a subgraph $S \subseteq E$ and a coloring κ , we say that κ is *monochromatic on the components of S* if $\kappa(u) = \kappa(v)$ for all vertices u and v in the same component of S . These are precisely the colorings such that $\kappa(u) = \kappa(v)$ for all $uv \in S$. If κ is a proper coloring, it follows that κ can only be monochromatic on the components of S if $S = \emptyset$. The following lemma, while trivial to prove, is extremely useful in what follows.

Lemma 1.1. *For every graph $G = (V, E)$ and subgraph $S \subseteq E$, the number of $[t]$ -colorings of G that are monochromatic on the components of S is $t^{c(S)}$.*

Proof. Every such coloring is obtained by independently choosing one of the t available colors for each component of S . \square

Since some of the results we prove involve signs, it is natural to use sign-reversing involutions. Let \mathcal{S} be a finite set with an associated sign function $\text{sgn} : \mathcal{S} \rightarrow \{+1, -1\}$. A *sign-reversing involution* on \mathcal{S} is an involution $\iota : \mathcal{S} \rightarrow \mathcal{S}$ that satisfies the following two conditions.

- If s is a fixed point of ι , then $\text{sgn}(s) = +1$.
- If (s, t) is a two-cycle of ι , then $\text{sgn}(s) = -\text{sgn}(t)$.

Immediately from the definition we see that

$$\sum_{s \in \mathcal{S}} \text{sgn}(s) = |\text{Fix } \iota|,$$

where $\text{Fix } \iota = \{s \in \mathcal{S} \mid \iota(s) = s\}$ is the set of fixed points of ι and the vertical bars denote cardinality.

In order to generalize the chromatic polynomial to the chromatic symmetric function, let $\mathbf{x} = \{x_1, x_2, \dots\}$ be an infinite set of commuting variables indexed by the positive integers. To every \mathbb{P} -coloring of the graph $G = (V, E)$ we associate the monomial or *weight*

$$\mathbf{x}^\kappa = \prod_{v \in V} x_{\kappa(v)}.$$

For example, both colorings in Figure 1 have monomial $\mathbf{x}^\kappa = x_1 x_3 x_3 x_4 = x_1 x_3^2 x_4$. The *chromatic symmetric function* of G is then

$$X(G; \mathbf{x}) = \sum_{\kappa} \mathbf{x}^\kappa,$$

where the sum is over all proper \mathbb{P} -colorings of G . That $X(G; \mathbf{x})$ is a symmetric function follows from the observation that permuting the labels of the colors in a proper coloring results in another proper coloring. For our purposes, however, the fact that $X(G; \mathbf{x})$ is a

symmetric function is not very important, and the reader more comfortable with generating functions than symmetric functions will lose nothing by viewing $X(G; \mathbf{x})$ as the weight generating function for proper \mathbb{P} -colorings of G .

By making the substitution

$$x_i = \begin{cases} 1 & \text{if } i \in [t], \\ 0 & \text{if } i \notin [t], \end{cases}$$

all terms of $X(G; \mathbf{x})$ vanish except those corresponding to proper $[t]$ -colorings of G , which become equal to 1. Thus $X(G; \mathbf{x})$ reduces to $\chi(G; t)$ under this substitution, so the chromatic symmetric function is indeed a generalization of the chromatic polynomial.

To describe the analogue of Lemma 1.1 in this context, we must introduce a particular basis for the ring of symmetric functions, the *power sum symmetric functions*. All bases of the ring of symmetric functions can be indexed by integer partitions, and these are no exception. Suppose that $\lambda = (\lambda_1, \dots, \lambda_\ell)$ is a partition of the positive integer n , meaning that the λ_i are positive integers called *parts* satisfying $\lambda_1 + \dots + \lambda_\ell = n$ and $\lambda_1 \geq \dots \geq \lambda_\ell$. We define $p_\lambda = p_\lambda(\mathbf{x})$ as the product

$$p_\lambda = p_{\lambda_1} p_{\lambda_2} \cdots p_{\lambda_\ell},$$

where, for a positive integer k ,

$$p_k = x_1^k + x_2^k + x_3^k + \cdots.$$

For example,

$$p_{(3,1)} = p_3 p_1 = (x_1^3 + x_2^3 + x_3^3 + \cdots)(x_1 + x_2 + x_3 + \cdots).$$

Finally, for a subgraph $S \subseteq E$ of the graph $G = (V, E)$, we let $\lambda(S)$ denote the integer partition defined by listing the number of vertices in each component of S in weakly decreasing order, so $\lambda(S)$ is a partition of $|V|$ into $c(S)$ parts. For example, the subgraph shown on the right of Figure 1 has $\lambda(S) = (3, 1)$. We then have the following analogue of Lemma 1.1.

Lemma 1.2. *For every graph $G = (V, E)$ and subgraph $S \subseteq E$, we have*

$$\sum_{\kappa} \mathbf{x}^\kappa = p_{\lambda(S)},$$

where the sum is over all \mathbb{P} -colorings of G that are monochromatic on the components of S .

Proof. The symmetric function for monochromatic colorings of a component with k vertices is p_k . Since each component of S may be colored independently, the product of the p_k for all parts k of λ accounts for every coloring that is monochromatic on the components of S . \square

We could have established Lemma 1.2 first and then derived Lemma 1.1 from it by substitution as previously described. Indeed, throughout this paper we present results in pairs like this where the first is a specialization of the second. However, instead of establishing the first from the second by substitution, we establish the second by adapting the bijection used to prove the first. These adaptations are quite minor. The only differences in the

symmetric function case are that the set of colorings is infinite, and that we have a weight function that associates a monomial to each coloring. The cancellation we desire, which results in the equality

$$\sum_{s \in \mathcal{S}} \text{sgn}(s) \text{wt}(s) = \sum_{s \in \text{Fix } \iota} \text{wt}(s),$$

still holds so long as the number of elements in \mathcal{S} of a given weight is finite and ι is *weight-preserving*, meaning that $\text{wt}(s) = \text{wt}(\iota(s))$ for all $s \in \mathcal{S}$.

With this preparation complete, we are ready to construct our bijective proofs. In Section 2, we show how $\chi(G; t)$ and $X(G; \mathbf{x})$ can both be expressed as sums over all subgraphs of G . Then in Section 3, we show (using a restriction of the bijection from Section 2) how they can both be expressed as sums over only those subgraphs without broken circuits. Finally, in Sections 4 and 5 we give bijective proofs of Stanley's results linking $\chi(G; t)$ and $X(G; \mathbf{x})$ to acyclic orientations of graphs. It should be noted that these two sections are inspired by the work of Blass and Sagan [6]. However, our approach differs from theirs in several respects. First, we treat both the chromatic polynomial and the chromatic symmetric function while they only considered the former (the latter not having been defined when [6] was written). Second, we give an explicit definition of the inverse of the bijection used in Section 4, while they only described one direction and then showed it was one-to-one and onto. Finally, we establish Stanley's theorem on acyclic orientations and compatible colorings for all negative values of t in Section 5, whereas Blass and Sagan only considered the case of $t = -1$.

2 Arbitrary subgraphs

Both $\chi(G; t)$ and $X(G; \mathbf{x})$ can be expressed as sums over the (spanning) subgraphs of G . In the case of $\chi(G; t)$, this expression is given by the following result of Birkhoff, from which it follows immediately that the chromatic polynomial is indeed a polynomial.

Theorem 2.1 (Birkhoff [4]). *For every graph $G = (V, E)$ and every positive integer t , we have*

$$\chi(G; t) = \sum_{S \subseteq E} (-1)^{|S|} t^{c(S)}.$$

To give a bijective proof of this result, we start by giving a combinatorial interpretation to the right-hand side. If $S \subseteq E$ is a subgraph of G , then Lemma 1.1 shows the number of $[t]$ -colorings of V that are monochromatic on the components of S is $t^{c(S)}$. Thus the right-hand side of Theorem 2.1 can be written as

$$\sum_{(S, \kappa) \in \mathcal{S}} (-1)^{|S|},$$

where \mathcal{S} is the set of pairs defined by

$$\mathcal{S} = \{(S, \kappa) : S \subseteq E \text{ and } \kappa : V \rightarrow [t] \text{ is monochromatic on the components of } S\}.$$

At this point the reader is welcome to derive Theorem 2.1 by inclusion–exclusion (it won't be difficult), but to practice for our later proofs, we employ a sign-reversing involution ι .

For the rest of this proof—in fact, for the rest of the paper—it is necessary to fix a total ordering on the edges of G . With this order fixed, we describe the edges of G as being *first*, *last*, *earlier*, or *later*, referring to their positions in this order. We also define a sign function on the pairs in \mathcal{S} by

$$\text{sgn}(S, \kappa) = (-1)^{|S|},$$

so

$$\sum_{S \subseteq E} (-1)^{|S|} t^{c(S)} = \sum_{(S, \kappa) \in \mathcal{S}} \text{sgn}(S, \kappa).$$

We may now define ι , which is the star of the rest of this section and the next. If κ is a proper coloring, we define $\iota(S, \kappa) = (S, \kappa)$. As remarked earlier, this case can only occur if $S = \emptyset$, because otherwise S must have at least one nontrivial component and thus the proper coloring κ cannot be monochromatic on its components.

Now suppose that κ has at least one monochromatic edge. We take e to be the last such edge (in the ordering on the edges of G fixed above) and define

$$\iota(S, \kappa) = (S \triangle e, \kappa),$$

where \triangle is the symmetric difference operator, removing e from S if it is present and adding it to S otherwise.

Our proof of Theorem 2.1 is completed with the following proposition, which implies that

$$\sum_{S \subseteq E} (-1)^{|S|} t^{c(S)} = \sum_{(S, \kappa) \in \mathcal{S}} \text{sgn}(S, \kappa) = |\text{Fix } \iota| = \chi(G; t).$$

Proposition 2.2. *The mapping ι is a sign-reversing involution on \mathcal{S} with*

$$\text{Fix } \iota = \{(\emptyset, \kappa) \in \mathcal{S} \mid \kappa \text{ is a proper coloring of } G\}.$$

Proof. First we verify that ι maps \mathcal{S} to \mathcal{S} . This is clear if κ is a proper coloring, so we must show that if $(S, \kappa) \in \mathcal{S}$ where κ is not proper, then $(S \triangle e, \kappa) \in \mathcal{S}$, i.e., κ is still monochromatic on the components of $S \triangle e$. If $S \triangle e = S - e$, then this holds because each component of $S - e$ is contained in a component of S . If $S \triangle e = S \cup e$, then it is possible that two components of S were merged by the addition of the edge e . If that occurs, then since κ is monochromatic on e , κ must have been the same color on both components that were merged. In either case, κ is monochromatic on the components of $S \triangle e$, so $\iota(S, \kappa) \in \mathcal{S}$ for all $(S, \kappa) \in \mathcal{S}$.

The fixed points of ι are precisely the pairs (\emptyset, κ) where κ is a proper $[t]$ -coloring of S . Moreover, these fixed points have positive sign, as desired. On the other elements of \mathcal{S} , ι is sign-reversing because $|S \triangle e| = |S| \pm 1$. Finally, ι is an involution because κ , and thus the definition of the edge e , does not change when passing from (S, κ) to $\iota(S, \kappa)$. \square

The proof of the symmetric function analogue requires only minor modifications.

Theorem 2.3 (Stanley [13, Theorem 2.5]). *For every graph $G = (V, E)$, we have*

$$X(G; \mathbf{x}) = \sum_{S \subseteq E} (-1)^{|S|} p_{\lambda(S)}.$$

Proof. Let \mathcal{S} denote the set of pairs (S, κ) where $S \subseteq E$ and κ is a \mathbb{P} -coloring of G that is monochromatic on the components of S . Since ι does not depend on the range of κ , its definition extends to pairs where κ has range \mathbb{P} without modification, as do both the statement and the proof of Proposition 2.2.

By defining the weight of a pair $(S, \kappa) \in \mathcal{S}$ as

$$\text{wt}(S, \kappa) = \mathbf{x}^\kappa$$

and appealing to Lemma 1.2, we see that

$$\sum_{S \subseteq E} (-1)^{|S|} p_{\lambda(S)} = \sum_{(S, \kappa) \in \mathcal{S}} \text{sgn}(S, \kappa) \text{wt}(S, \kappa).$$

Since $\text{wt}(S, \kappa)$ depends only on κ , which does not change in passing from (S, κ) to $\iota(S, \kappa)$, it follows that ι is weight-preserving. Because ι is (still) sign-reversing, every term in the sum on the right-hand side of the equation above cancels except those terms corresponding to pairs $(\emptyset, \kappa) \in \mathcal{S}$ where κ is a proper \mathbb{P} -coloring of G , and thus

$$\sum_{S \subseteq E} (-1)^{|S|} p_{\lambda(S)} = \sum_{(S, \kappa) \in \mathcal{S}} \text{sgn}(S, \kappa) \text{wt}(S, \kappa) = \sum_{(S, \kappa) \in \text{Fix } \iota} \text{wt}(S, \kappa) = X(G; \mathbf{x}),$$

proving the theorem. □

3 NBC subgraphs

In this section we show how the bijective proof of Theorem 2.1 given in the previous section leads quickly to Whitney's no broken circuits theorem and its chromatic symmetric function analogue.

First we must define broken circuits. A *walk* from the vertex v_0 to the vertex v_k in the graph $G = (V, E)$ is an alternating sequence $v_0, e_1, v_1, e_2, v_2, \dots, e_k, v_k$ of vertices and edges such that for all $i \in [k]$ we have $e_i = v_{i-1}v_i \in E$. This walk is a *path* if it does not repeat any vertices or edges, except possibly the first and the last vertices. A walk or a path is *closed* if $v_0 = v_k$, that is, if its first and last vertices are the same. A *cycle*, also known as a circuit, is a closed path. Just as we conflate subgraphs with their edge sets, we also conflate paths and cycles with their edge sets. For example, the graph G in Figure 2 has a unique cycle which we denote as $C = \{e_1, e_2, e_3\}$. Graphs which do not contain cycles can be referred to as *acyclic*, but are more commonly called *forests*.

Given a fixed total ordering on the edges of G , a *broken circuit* is a set $B \subseteq E$ of the form

$$B = C - \max C$$

where C is a cycle and $\max C$ is the largest edge of C . For example, if the edges of the graph in Figure 2 are ordered $e_1 < e_2 < e_3 < e_4$, then the cycle $C = \{e_1, e_2, e_3\}$ corresponds to the broken circuit $B = \{e_1, e_2\}$. A subgraph $S \subseteq E$ is *NBC* (short for *no broken circuits*) if it does not contain any broken circuits. It is frequently helpful to note that NBC subgraphs




Figure 2: A graph with ordered edges $e_1 < e_2 < e_3 < e_4$, a cycle in this graph, the corresponding broken circuit, and an NBC subgraph (with colors).

are necessarily forests; since they don't contain broken circuits, they certainly don't contain cycles. The rightmost picture in Figure 2 represents an NBC subgraph (together with colors that the reader may ignore for now).

Consider a particular pair (S, κ) from the set \mathcal{S} in the proof of Theorem 2.1, so S is an arbitrary subgraph of G and κ is monochromatic on the components of S . Suppose that S contains a broken circuit, say $B = C - \max C$. Since the vertices of C lie in the same component of S (because $B \subseteq S$), it follows that κ is monochromatic on the vertices of C , and thus in particular on the edge $\max C$. Therefore, if S contains the broken circuit B , then we have $\iota(S, \kappa) = (S \triangle e, \kappa)$ for an edge e that is either equal to $\max C$ or occurs after it in the total ordering on edges of G . In particular, if the broken circuit B is contained in S , then it is also contained in $S \triangle e$. We record this fact below.

Proposition 3.1. *The mapping ι restricts to a fixed-point-free sign-reversing involution on the set of pairs (S, κ) where S is a subgraph containing a broken circuit and κ is monochromatic on the components of S .*

This result shows that the contributions of all such pairs cancel in the summation in Theorem 2.1, leaving us with only the NBC subgraphs, and establishing Whitney's no broken circuits theorem.

Theorem 3.2 (Whitney [16, Section 7]). *For every graph $G = (V, E)$, every total ordering of its edges, and every positive integer t , we have*

$$\chi(G; t) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} (-1)^{|S|} t^{c(S)}. \quad \square$$

The symmetric function analogue of the no broken circuit theorem follows immediately from the observation, made in the proof of Lemma 2.3, that the quantity $\text{wt}(S, \kappa) = \mathbf{x}^\kappa$ is preserved by ι because κ does not change.

Theorem 3.3 (Stanley [13, Theorem 2.9]). *For every graph $G = (V, E)$ and every total ordering of its edges, we have*

$$X(G; \mathbf{x}) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} (-1)^{|S|} p_{\lambda(S)}. \quad \square$$




Figure 3: A graph, an acyclic orientation, and a compatible coloring (which can be ignored until Section 5).

4 Acyclic orientations

Thus far we have only defined $\chi(G; t)$ for positive integers t , where Theorem 2.1 shows that it is equal to a polynomial. For the rest of the paper, we take this polynomial as the *definition* of $\chi(G; t)$. With this change, we may evaluate $\chi(G; t)$ at negative values of t . Our first such evaluation, below, is a consequence of Whitney's no broken circuits theorem.

Corollary 4.1. *For every graph $G = (V, E)$, every total ordering of its edges, and every positive integer t , we have*

$$\chi(G; -t) = (-1)^{|V|} \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} t^{c(S)}.$$

Proof. Theorem 3.2 shows us that

$$\chi(G; -t) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} (-1)^{|S|+c(S)} t^{c(S)}.$$

Since every NBC subgraph is a forest, an NBC subgraph S of G has $|V| - |S|$ components. The result follows immediately. \square

Stanley has established a beautiful interpretation of the value of $\chi(G; t)$ for negative integers t . Indeed, this relationship is one of the most striking examples of the phenomenon he later termed *combinatorial reciprocity* [12]. In this section we give a bijective proof of this result in the special case of $t = -1$, following in the footsteps of Blass and Sagan [6]. In the next section we demonstrate the counter-intuitive fact that this special case implies the theorem for all negative values of t , and also establish the analogous theorem for the chromatic symmetric function.

Stanley's interpretation of $\chi(G; t)$ for negative t involves orientations of G , and so we must introduce oriented (or directed) edges, which we call *arcs*. Given a pair of distinct vertices u and v , there are two possible arcs between them, which we denote by $\overrightarrow{uv} = \overleftarrow{vu}$ and $\overleftarrow{uv} = \overrightarrow{vu}$. An *orientation* of the (undirected) graph $G = (V, E)$ is obtained by replacing each edge $uv \in E$ by one of the arcs \overrightarrow{uv} or \overrightarrow{vu} . Each such orientation is then an *oriented graph*, which we denote by $O = (V, A)$, where A is the set of arcs. Figure 3 show a graph on the left and one of its orientations in the center.

The bijection employed in this section proceeds edge-by-edge, so in the intermediate stages our graphs have both edges and arcs. It is therefore convenient to set our proof in

the context of mixed graphs, whose study dates to a 1966 paper of Harary and Palmer [7]. A *mixed graph* is a triple $M = (V, E, A)$ where V is a set of vertices, E is a set of edges, and A is a set of arcs. To eliminate the possibility of confusion, we briefly adapt our definitions of walks, paths, and cycles to this context.

A *walk* from v_0 to v_k in the mixed graph $M = (V, E, A)$ is an alternating sequence $v_0, c_1, v_1, c_2, v_2, \dots, c_k, v_k$ of vertices and edges/arcs such that for all $i \in [k]$, either $c_i = v_{i-1}v_i \in E$ or $c_i = \overrightarrow{v_{i-1}v_i} \in A$. We say that this walk *traverses* each of these edges and arcs. We further call this walk a *path* if it does not repeat any vertices, arcs, or edges, except possibly the first and last vertices. A path (or more generally, a walk) is *closed* if its first and last vertices are the same, and a closed path is a *cycle*. Finally, a mixed graph is called *acyclic* if it does not contain a cycle. For example, the mixed graph (which also happens to be an oriented graph) in the center of Figure 3 is acyclic.

One of the basic lemmas of graph theory is that if a graph contains a walk between two vertices, then it also contains a path between them (to prove this, simply shorten the walk each time it repeats a vertex). For mixed graphs, we have the following result whose proof is similar in spirit.

Proposition 4.2. *If a mixed graph contains a closed walk that traverses at least one arc, then it also contains a cycle.*

Proof. We prove the proposition by induction on the number of vertices in such a walk. The base case with a single vertex is trivial. Suppose that for some $k \geq 1$, a mixed graph contains the walk $v_0, c_1, v_1, c_2, v_2, \dots, c_k, v_k$, which we denote by W . Further suppose that W is closed (so $v_0 = v_k$) and that W traverses at least one arc. If W is not itself a cycle, then we must have $v_i = v_j$ for indices $i < j$ satisfying either $i > 0$ or $j < k$. Thus W contains two shorter closed walks:

$$v_i, c_{i+1}, v_{i+1}, \dots, c_j, v_j$$

and

$$v_j, c_{j+1}, v_{j+1}, \dots, c_k, v_k, c_1, v_1, \dots, c_i, v_i.$$

These walks together traverse all edges and arcs traversed by W , so at least one of them must traverse an arc. The proof is completed by appealing to induction with that shorter walk. \square

Stanley presented the following result as a corollary of Theorem 5.1. However, we appeal to Proposition 4.3 in our proof of Theorem 5.1.

Proposition 4.3 (Stanley [11, Corollary 1.3]). *For every graph $G = (V, E)$,*

$$(-1)^{|V|} \chi(G; -1) = \text{the number of acyclic orientations of } G.$$

Proof. Fix a total ordering $e_1 < e_2 < \dots < e_m$ of the edges of G , so Corollary 4.1 shows that

$$\chi(G; -1) = (-1)^{|V|} \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} 1,$$

or in words, that $(-1)^{|V|} \chi(G; -1)$ is equal to the number of NBC subgraphs of G .

It therefore suffices to exhibit a bijection between the NBC subgraphs of G and its acyclic orientations. Define \mathcal{M}_i to be the set of acyclic mixed graphs comprised of




Figure 4: Applying the mappings ϕ_1, ϕ_2, ϕ_3 , and ϕ_4 to an acyclic orientation of the graph shown in the top left, where the normal orientations of the edges are as shown in the top right.

- an NBC subset of the edges $\{e_1, \dots, e_i\}$ and
- an orientation of the edges $\{e_{i+1}, \dots, e_m\}$.

Recall that an NBC subset of edges is automatically acyclic. Thus \mathcal{M}_0 and \mathcal{M}_m consist of the acyclic orientations and NBC subgraphs of G , respectively. We establish the proposition by constructing bijections $\phi_i : \mathcal{M}_{i-1} \rightarrow \mathcal{M}_i$ for all indices $1 \leq i \leq m$.

To define these bijections, we need to distinguish between the two orientations of each edge of G . How we make this distinction is immaterial, so we simply fix one orientation of each edge e_i as *normal*, denoted by \vec{e}_i , and call the opposite orientation \overleftarrow{e}_i *abnormal*. We may now define $\phi_i : \mathcal{M}_{i-1} \rightarrow \mathcal{M}_i$. Take $M \in \mathcal{M}_{i-1}$ and suppose that the edge e_i of G appears as the arc a_i in M . The mixed graph $\phi_i(M)$ is obtained by either *unorienting* a_i (replacing the arc a_i with the edge e_i) or *removing* it (deleting the arc a_i). The rules by which we make this choice are as follows. We unorient a_i if both

- (A) a_i is the normal orientation of e_i and
- (B) unorienting a_i does not create a cycle.

Otherwise, we remove a_i . An example of applying the mappings ϕ_i is shown in Figure 4, where the label above each arrow indicates which of (A) or (B) was violated if the arc was removed (there is no such label if the arc satisfied both (A) and (B) and thus was unoriented instead).

We must check that ϕ_i produces members of \mathcal{M}_i . It follows immediately that $M' = \phi_i(M)$ is acyclic because rule (B) prevents the creation of a cycle. However, how does ϕ_i avoid creating broken circuits? The answer is that were M' to contain a broken circuit, then—since $M \in \mathcal{M}_{i-1}$ and thus does not itself contain a broken circuit—the last two edges in the corresponding cycle of G must be e_i and e_j for some $j > i$. However, M contains some orientation of e_j , so unorienting a_i would result in a cycle. Therefore rule (B) ensures that a_i is removed, and thus M' cannot contain a broken circuit. This verifies that $M' \in \mathcal{M}_i$.

Our proof is completed by constructing the inverse of ϕ_i , which we denote by $\psi_i : \mathcal{M}_i \rightarrow \mathcal{M}_{i-1}$. Given a mixed graph $M' \in \mathcal{M}_i$, the mixed graph $M = \psi_i(M')$ is obtained by adding one of the orientations of edge e_i to M' (and removing e_i if it is present in M'). We give e_i the *abnormal* orientation if both

(A') e_i is not an edge of M' and

(B') adding $\overleftarrow{e_i}$ to M' does not create a cycle.

Otherwise, we give e_i the normal orientation.

Again we must check that $M \in \mathcal{M}_{i-1}$. Orienting an edge can not cause M to contain an NBC set, so we need only check that M is acyclic. This follows from (B') if e_i was oriented abnormally. If e_i was oriented normally, then either (A') or (B') was violated. If it was (A') that was violated, then $e_i \in M'$, and M is acyclic because orienting an edge that is already present cannot create a cycle.

This leaves us to consider the case where e_i was oriented normally because condition (B') was violated. Suppose $\overleftarrow{e_i} = \overleftarrow{uv}$, so the fact that condition (B') was violated means that M' contains a path from u to v , which we denote by P . If P were to consist entirely of edges, then $P \subseteq \{e_1, e_2, \dots, e_{i-1}\}$ because $M' \in \mathcal{M}_i$. However, in this case e_i is the greatest edge of the cycle $P \cup e_i$ of G , so P is a broken circuit. As this contradicts our assumption that $M' \in \mathcal{M}_i$, it follows that P must traverse at least one arc. Therefore, if M' were to also contain a path from v to u , then this path together with P would be a closed walk that traversed at least one arc. Proposition 4.2 would then imply that M' contains a cycle, a contradiction. We may therefore conclude that M' does not contain a path from v to u . Thus orienting e_i as $\overrightarrow{e_i} = \overrightarrow{uv}$, as it is oriented in $M = \psi_i(M')$, does not create a cycle.

It remains only to check that $\psi_i \circ \phi_i$ and $\phi_i \circ \psi_i$ are the identity mappings on \mathcal{M}_{i-1} and \mathcal{M}_i , respectively. We show only one of these, as the other is similar. Consider $\psi_i(\phi_i(M))$ where $M \in \mathcal{M}_{i-1}$, and let $M' = \phi(M)$. There are two cases, depending on whether M' is obtained from M by unorienting e_i or by removing it. In the first case e_i was oriented normally in M and e_i is an edge of M' . Thus (A') is violated and $\psi(M')$ is obtained by orienting e_i normally, as desired.

For the second case, suppose that e_i was removed, so either (A) or (B) was violated. If (A) was violated, then e_i was abnormally oriented in M and we must show that $\psi_i(M')$ adds the arc back in that orientation, or in other words, that both (A') and (B') are satisfied. We know that (A') is satisfied because e_i was removed from M due to its abnormal orientation. We also see that (B') holds because orienting e_i as $\overleftarrow{e_i}$ results in the mixed graph M which was assumed to be acyclic. For the final subcase, assume that (A) holds but (B) does not. This means that e_i was oriented normally in M but was removed because (B) was violated. In this case (B') is violated, and e_i is added back to M' in its normal orientation, as desired. \square

5 Multi-colored acyclic orientations

Proposition 4.3 is a consequence of a more general result in Stanley's 1973 paper [11]. Here, by adapting an idea essentially due to Vo [15, Section 3], we show how this more general result follows bijectively from Proposition 4.3, and then establish the chromatic symmetric function analogue.

Suppose that $O = (V, A)$ is an orientation of the graph $G = (V, E)$ and that κ is a $[t]$ - or \mathbb{P} -coloring of G . We say that O and κ are *compatible* if

$$\overrightarrow{uv} \in A \text{ implies } \kappa(u) \leq \kappa(v),$$

that is, if all the arcs of O point in the direction of weakly increasing values of κ . It is convenient to display such pairs by labeling the vertices of O with their values under κ . Using this convention, we see a compatible pair on the right in Figure 3. We can now state the generalization of Proposition 4.3 to all negative values of t .

Theorem 5.1 (Stanley [11, Theorem 1.2]). *For every graph $G = (V, E)$ and every positive integer t we have*

$$(-1)^{|V|}\chi(G; -t) = |\{(O, \kappa) \mid O \text{ is an acyclic orientation of } G \text{ and } \kappa \text{ a compatible } [t]\text{-coloring}\}|.$$

Proof. Fix an total ordering of the edges of G , so by Corollary 4.1 we have

$$(-1)^{|V|}\chi(G; -t) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} t^{c(S)}.$$

Lemma 1.1 shows that this sum counts pairs (S, κ) where S is an NBC subgraph of G and κ is a $[t]$ -coloring that is monochromatic on the components of S . We would like to construct a bijection between such pairs and pairs of the form (O, κ) , where O is an acyclic orientation of G and κ is a compatible $[t]$ -coloring. To prove the theorem, we construct, for every $[t]$ -coloring κ of G , a bijection Φ between the sets

$$\begin{aligned} \mathcal{O}_\kappa &= \{O : O \text{ is an orientation of } G \text{ compatible with } \kappa\} \text{ and} \\ \mathcal{S}_\kappa &= \{S : S \subseteq E \text{ is NBC and } \kappa \text{ is monochromatic on the components of } S\}. \end{aligned}$$

Note that when $t = 1$, there is only one coloring of G (the constant coloring), and thus in this case Φ must restrict to a bijection from acyclic orientations of G to its NBC subgraphs. Indeed, we constructed just such a bijection in the previous section, although any such bijection will do for our purposes here.

For each $i \in [t]$, define

$$V_i = \{v \in V \mid \kappa(v) = i\},$$

and let E_i denote the set of (monochromatic) edges of G between vertices in V_i . Note that $E_1 \cup E_2 \cup \dots \cup E_t$ needn't contain all of the edges of G , but for every $S \in \mathcal{S}_\kappa$ we have

$$S \subseteq E_1 \cup E_2 \cup \dots \cup E_t$$

because κ is monochromatic on the components of S . Similarly, for each orientation $O \in \mathcal{O}_\kappa$, we know how the edges outside $E_1 \cup E_2 \cup \dots \cup E_t$ are oriented because O and κ are compatible. Indeed, any edge uv where $u \in V_i$, $v \in V_j$, and $i < j$ must be oriented as \overrightarrow{uv} . Let C denote the set of these arcs.

We may now define Φ . Take some orientation $O \in \mathcal{O}_\kappa$, and for each $i \in [t]$, let O_i denote the sub-orientation consisting of all arcs of O between vertices of color i , so O_i is an orientation of the subgraph E_i . Because we are assuming Proposition 4.3, we know that for each E_i , there is a bijection from the acyclic orientations of E_i to its NBC subgraphs. Let

this bijection be denoted by $\phi^{(i)}$. Thus $\phi^{(i)}(O_i)$ is an NBC subgraph of E_i for every $i \in [t]$. We may now define $\Phi(O)$ simply by

$$\Phi(O) = \phi^{(1)}(O_1) \cup \phi^{(2)}(O_2) \cup \dots \cup \phi^{(t)}(O_t).$$

Note that the components of $\phi^{(i)}(O_i)$ are subsets of E_i . This implies that κ is monochromatic on the components of $\Phi(O)$. This fact also allows us to see that $\Phi(O)$ is an NBC subgraph—if $\Phi(O)$ were to contain a broken circuit, then that broken circuit would lie in a single component, and thus it would be contained in $\phi^{(i)}(O_i)$ for some i , but this is a contradiction. It follows that $\Phi(O) \in \mathcal{S}_\kappa$.

We construct a mapping Ψ which will be the inverse of Φ . To define it, consider an NBC subgraph $S \in \mathcal{S}_\kappa$. Because κ is monochromatic on the components of S , we know that $S \subseteq E_1 \cup E_2 \cup \dots \cup E_t$. For each $i \in [t]$, let $S_i = S \cap E_i$, so S_i is an NBC subgraph of E_i . For each $i \in [t]$, let $\psi^{(i)}$ denote the inverse of the bijection $\phi^{(i)}$. We then define $\Psi(S)$ by

$$\Psi(S) = C \cup \psi^{(1)}(S_1) \cup \psi^{(2)}(S_2) \cup \dots \cup \psi^{(t)}(S_t)$$

where C is the set of arcs defined earlier (whose orientations are determined by compatibility).

We need to make sure that $\Psi(S)$ is acyclic. Suppose, toward a contradiction, that this orientation contains a cycle. The vertices visited by this cycle cannot all have the same color since each $\psi^{(i)}(S_i)$ is acyclic. Thus we may label the vertices visited by the cycle as $v_0, v_1, v_2, \dots, v_k = v_0$ where $\kappa(v_0) < \kappa(v_1)$. However, the fact that κ is compatible with this orientation implies that

$$\kappa(v_0) < \kappa(v_1) \leq \kappa(v_2) \leq \dots \leq \kappa(v_k) = \kappa(v_0).$$

This contradiction—that $\kappa(v_0) < \kappa(v_0)$ —establishes that the orientation is acyclic, showing that $\Psi(S) \in \mathcal{O}_\kappa$. It follows that Φ is a bijection with inverse Ψ since the analogous statement is true of the $\phi^{(i)}$ and $\psi^{(i)}$. This completes the proof. \square

To state the symmetric function analogue of Theorem 5.1, we first need to introduce a well-known involution ω on the ring of symmetric functions. When $\lambda = (\lambda_1, \dots, \lambda_\ell)$ is a partition of n into ℓ parts, we define

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

We then extend ω to the entire ring linearly.

Theorem 5.2 (Stanley [13, Theorem 4.2]). *For every graph G , we have*

$$\omega(X(G; \mathbf{x})) = \sum_{(O, \kappa)} \mathbf{x}^\kappa,$$

where the sum is over all pairs (O, κ) where O is an acyclic orientation of G that is compatible with the \mathbb{P} -coloring κ .

Proof. From Theorem 3.3 we have that

$$X(G; \mathbf{x}) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} (-1)^{|S|} p_{\lambda(S)}.$$

All subgraphs S appearing in this sum are forests (because they are NBC), so $\lambda(S)$ is a partition of $|V|$ into $|V| - |S|$ parts. Thus applying ω to the equation above yields

$$\omega(X(G; \mathbf{x})) = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} (-1)^{|S|} (-1)^{|V| - (|V| - |S|)} p_{\lambda(S)} = \sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} p_{\lambda(S)}.$$

The right-hand side of this equation is, by Lemma 1.2, the weight generating function for pairs (S, κ) where S is an NBC set and κ is monochromatic on the components of S . Since the bijection Ψ of the previous proof preserves colorings, it also preserves the quantity \mathbf{x}^κ , so this bijection shows that

$$\sum_{\substack{S \subseteq E, \\ S \text{ is NBC}}} p_{\lambda(S)} = \sum_{(O, \kappa)} \mathbf{x}^\kappa,$$

finishing the proof of the theorem. □

6 Concluding remarks

Even after more than a century of study, several questions about the chromatic polynomial remain open. To name a particularly natural and tantalizing one, which was raised by Read in his 1968 survey [10, Section 9] and then again (independently, it seems) by Wilf [17] eight years later: which polynomials are chromatic polynomials? Of course, the analogous question for chromatic symmetric functions is just as natural.

As we have seen, mixed graphs arise naturally in the study of the chromatic polynomial. Beck, Bogart, and Pham [2] have extended the definition of chromatic polynomials to this context. Specifically, they define the *strong chromatic polynomial* of the mixed graph $G = (V, E, A)$ to be the number of $[t]$ -colorings κ satisfying

- (a) $\kappa(u) \neq \kappa(v)$ if $uv \in E$ and
- (b) $\kappa(u) < \kappa(v)$ if $\overrightarrow{uv} \in A$.

They show that this is a polynomial in t . They also prove an analogue of Proposition 4.3 for these polynomials, phrased in the language of hyperplane arrangements. (Given any graph G with vertex set $[n]$, one associates a set of hyperplanes where $ij \in E$ corresponds to the hyperplane $x_i = x_j$. It is then easy to see that the acyclic orientations of G correspond to the regions obtained by removing the hyperplanes from \mathbb{R}^n .) Beck, Blado, Crawford, Jean-Louis, and Young [1] later introduced the *weak chromatic polynomial* of a mixed graph, which counts colorings satisfying the weak version of the inequality in (b). They establish a mixed graph analogue of Theorem 5.1 for these polynomials.




Figure 5: From left to right, the butterfly (or bowtie), the kite, the net, and X_{169} .

Finally, we would be remiss if we did not mention a question of Stanley’s that remains open despite significant interest. First, in terms of the definitions we have established, a *tree* is a forest with a single connected component. For any tree T on n vertices, it is easy to see that $\chi(T; t) = t(t - 1)^{n-1}$. Thus the chromatic polynomial does not distinguish trees. However, Stanley [13, p. 170] asked whether the opposite is true for the chromatic symmetric function. Specifically, if T_1 and T_2 are non-isomorphic trees, is it true that $X(T_1; \mathbf{x}) \neq X(T_2; \mathbf{x})$? Stanley repeated this question as Exercise 7.47.b in Volume 2 of *Enumerative Combinatorics* [14], and it has since become known as the *tree conjecture*. If true, the tree conjecture would show that the chromatic symmetric function is a *complete tree invariant*.

Note that the chromatic symmetric function does *not* distinguish all graphs. The example Stanley [13] gives is that the chromatic symmetric functions of the butterfly and the kite are equal (see Figure 5). Orellana and Scott [9] have observed that the chromatic symmetric function also does not distinguish unicyclic graphs (those with a single cycle). The example they present is the net and the graph that we refer to—following the *Information System on Graph Classes and their Inclusions* [8] for lack of a better name—as X_{169} (also shown in Figure 5).

References

- [1] M. Beck, D. Blado, J. Crawford, T. Jean-Louis, and M. Young. On weak chromatic polynomials of mixed graphs. *Graphs Combin.*, 31(1):91–98, 2015.
- [2] M. Beck, T. Bogart, and T. Pham. Enumeration of Golomb rulers and acyclic orientations of mixed graphs. *Electron. J. Combin.*, 19(3):Paper 42, 13 pp., 2012.
- [3] N. L. Biggs. *Algebraic Graph Theory*. Cambridge University Press, Cambridge, England, second edition, 1993.
- [4] G. D. Birkhoff. A determinant formula for the number of ways of coloring a map. *Ann. of Math. (2)*, 14(1-4):42–46, 1912/13.
- [5] G. D. Birkhoff and D. C. Lewis. Chromatic polynomials. *Trans. Amer. Math. Soc.*, 60:355–451, 1946.

- [6] A. R. Blass and B. E. Sagan. Bijective proofs of two broken circuit theorems. *J. Graph Theory*, 10(1):15–21, 1986.
- [7] F. Harary and E. M. Palmer. Enumeration of mixed graphs. *Proc. Amer. Math. Soc.*, 17:682–687, 1966.
- [8] Information System on Graph Classes and their Inclusions (ISGCI). Published electronically at <http://www.graphclasses.org/>.
- [9] R. Orellana and G. Scott. Graphs with equal chromatic symmetric functions. *Discrete Math.*, 320:1–14, 2014.
- [10] R. C. Read. An introduction to chromatic polynomials. *J. Combin. Theory*, 4:52–71, 1968.
- [11] R. P. Stanley. Acyclic orientations of graphs. *Discrete Math.*, 5:171–178, 1973.
- [12] R. P. Stanley. Combinatorial reciprocity theorems. *Adv. Math.*, 14:194–253, 1974.
- [13] R. P. Stanley. A symmetric function generalization of the chromatic polynomial of a graph. *Adv. Math.*, 111(1):166–194, 1995.
- [14] R. P. Stanley. *Enumerative Combinatorics. Vol. 2*, volume 62 of *Cambridge Stud. in Advanced Math.* Cambridge University Press, Cambridge, England, 1999.
- [15] K. P. Vo. Graph colorings and acyclic orientations. *Linear and Multilinear Algebra*, 22(2):161–170, 1987.
- [16] H. Whitney. A logical expansion in mathematics. *Bull. Amer. Math. Soc.*, (38):572–579, 1932.
- [17] H. S. Wilf. Which polynomials are chromatic? In *Colloquio Internazionale sulle Teorie Combinatorie (Roma, 3-15 Settembre 1973)*, volume 1, pages 247–256. Accademia Nazionale dei Lincei, Rome, Italy, 1976.