Symmetric Acyclicity Theorems

O. Coppola, M. Reilly
Knots and Graphs

Introduction

Previously we have discussed Stanley's Acyclicity Theorem, which states that the number of acyclic orientations of a graph is equal to its chromatic polynomial evaluated at -1 (up to sign). In 1995, Stanley strengthened this result in order to count acyclic orientations by their number of sinks. To do this, he introduced the symmetric chromatic function of a graph G, defined as

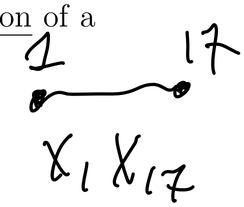

$$X_G = \sum_{\kappa \text{ properly colors } G} x_{\kappa(v_1)} x_{\kappa(v_2)} \cdots x_{\kappa(v_n)}$$

where $v_1, \ldots, v_n$ are the vertices of G. From now on, we will use x^κ to denote $x_{\kappa(v_1)} \cdots x_{\kappa(v_n)}$.

The definition of X_G doesn't provide a useful way of writing down X_G. As it turns out, we are always able to write X_G in terms of elementary symmetric functions, which are of the form $e_n = \sum_{0<i_1<i_2 \cdots <i_n} x_{i_1} x_{i_2} \ldots x_{i_n}$.

Theorem (Stanley). *When X_G is expressed in the elementary basis, the sum of the coefficients of the terms with k factors, is equal to the number of acyclic orientations of G with k sinks.*

Example. Let G be the star graph with 3 edges. It has the following orientations (all of which are acyclic):

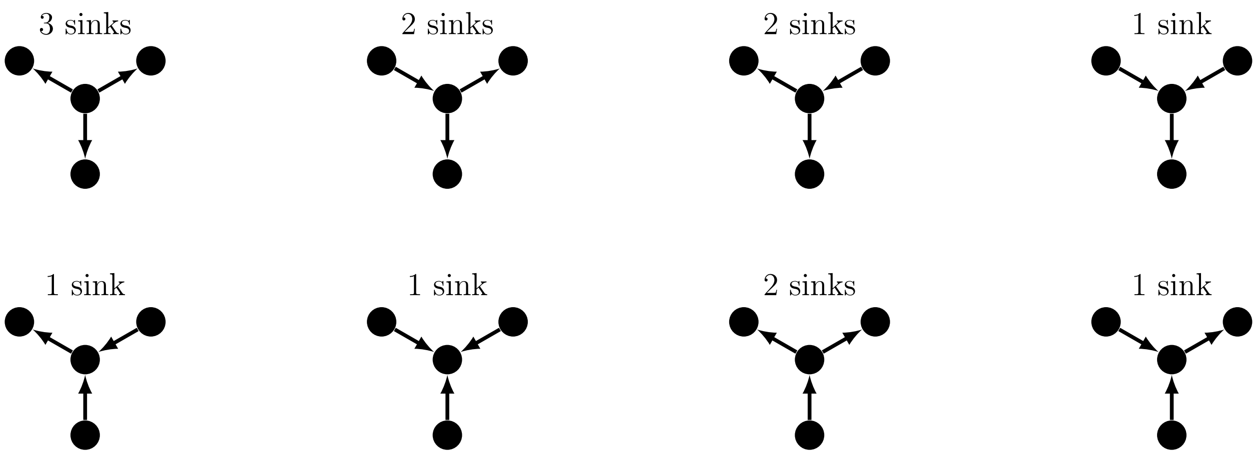

Trivial calculation shows that

$$\begin{aligned} X_G &= 5e_1e_3 + e_1^2e_2 - 2e_2^2 + 4e_4 \\ &= (4e_4) + (5e_1e_3 - 2e_2e_2) + (e_1e_1e_2) \end{aligned}$$

We can apply Stanley's theorem to see this means G has 4 orientations with 1 sink, $5-2=3$ orientations with 2 sinks, and 1 orientation with 3 sinks, exactly as there should be.

The Signed Case

The goal of our project is to generalize Stanley's theorem to signed graphs. To do this we need to start with the B-symmetric chromatic function of a signed graph Σ.

$$X_\Sigma = \sum_{\kappa \text{ properly colors } \Sigma} x^\kappa$$

$X_\bullet = \sum_{i \in \mathbb{N}} X_i$ $\qquad$ $X_\bullet = \sum_{i \in \mathbb{Z}} X_i$

The only difference with the usual symmetric chromatic function is that κ can assign negative colors since Σ is a signed graph. As an aside, this means that the B-symmetric chromatic function of a signed graph with all positive edges is trivially different from the symmetric chromatic function of the graph we get by pretending that it is an unsigned graph.

In this presentation we will present a proof of the following theorem.

Theorem. *There exists a set of functions*

$$\left\{\xi_n \,\middle|\, n = 0, 1, 2, \dots\right\} \cup \left\{q_{a,b} \,\middle|\, a, b \geq 1\right\} \cup \left\{z_n \,\middle|\, n = 0, 1, \dots\right\}$$

in the variables x_i for $i \in \mathbb{Z}$, such that any X_Σ can be uniquely expressed in terms of these functions and when written this way, the sum of the coefficients of the all of the terms of the form $\left(\prod q_{a,b} \cdot \prod z_n\right) \cdot \xi_{n_1} \cdots \xi_{n_k}$, is the number of acyclic orientations of Σ with $n_1 + \cdots + n_k$ sinks.

Example. Let Σ be the triangle with two negative edges.

It has the following orientations (all of which are acyclic):

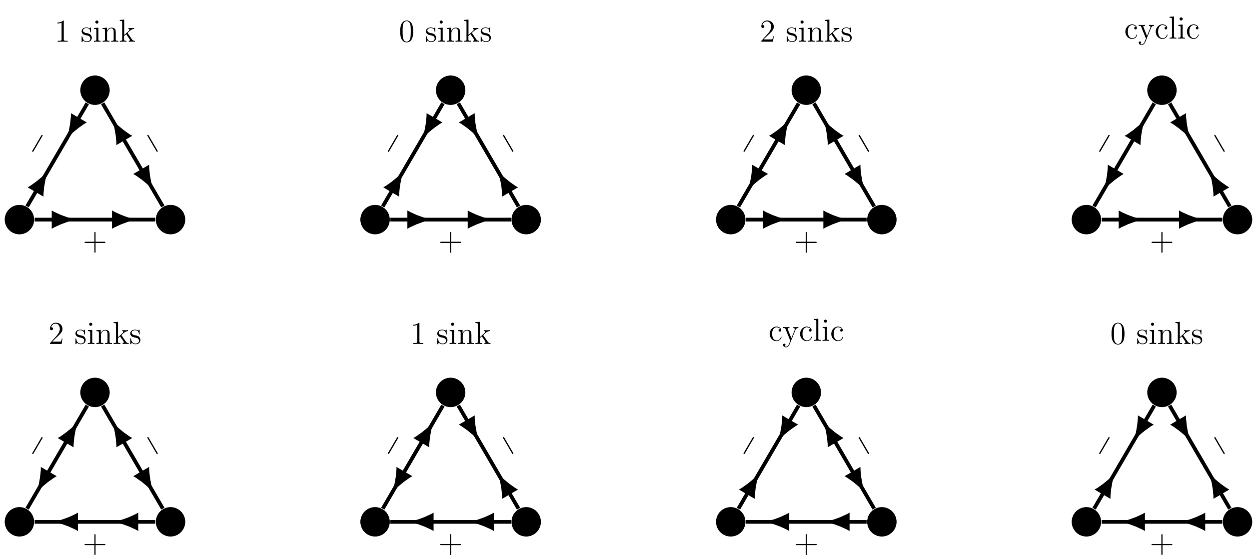

Trivial calculation shows that

$$
\begin{aligned}
X_\Sigma &= \xi_1^3 + 2q_{1,1}\xi_1 - \xi_1\xi_2 + 2\xi_1^2 + 2q_{2,1} \\
&= (\xi_1\xi_1\xi_1 - \xi_1\xi_2) + (2\xi_1\xi_1) + (2q_{1,1}\xi_1) + (2q_{2,1})
\end{aligned}
$$

We can apply the theorem to see this means Σ has $1 - 1 = 0$ orientations with 3 sinks, 2 orientations with 2 sinks, 2 orientations with 1 sink and 2 orientations with 0 sinks, exactly as there should be.

Proof Sketch

The proof is a culmination of work from this year as well as important contributions by Jake Huryn and previous years' research groups.

We begin as before by considering the connection between colorings and orientations. Each proper coloring preserves (induces) a unique acyclic orientation, and by allowing for enough colors, each acyclic orientation can be obtained from a proper coloring.

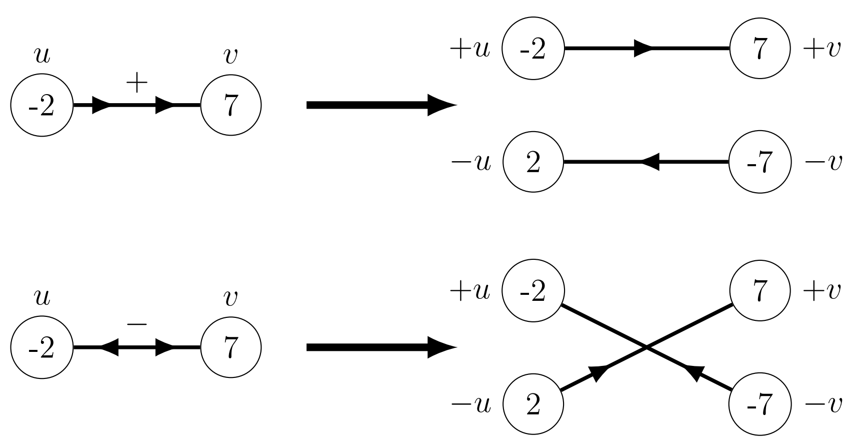

In other words, we can partition the set of proper colorings of Σ into sets of colorings which correspond to acyclic orientations of Σ. So,

$$\{\kappa : \kappa \text{ is a proper coloring of } \Sigma\} = \bigsqcup_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \{\kappa : \kappa \text{ preserves } P\}$$

We can introduce the notation $Y_P = \sum_{\kappa \text{ preserves } P} x^\kappa$ so that we may write

$$X_\Sigma = \sum_{\kappa \text{ properly colors } \Sigma} x^\kappa = \sum_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \left(\sum_{\kappa \text{ preserves } P} x^\kappa \right) = \sum_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} Y_P$$

Why use the letter P for orientations? Well, a directed (acyclic) unsigned graph can be thought of as a partial order on the set of vertices, and we can do a similar thing with directed signed graphs by considering the covering graph.

For example, the following directed graph imposes a partial ordering on its vertex set $\{a, b, c, d\}$.

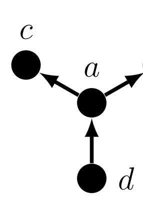

The relations are $c > a$, $b > a$, and $a > d$, as well as the implicit relations that arise from transitivity, $c > d$ and $b > d$. Note however that b and c are incomparable. (It could be that $b > c$ or $c > b$.) The structure of a set together with a partial ordering is called a poset.

For a directed (acyclic) signed graph, we can look at the covering graph and think of it as a poset on its vertices.

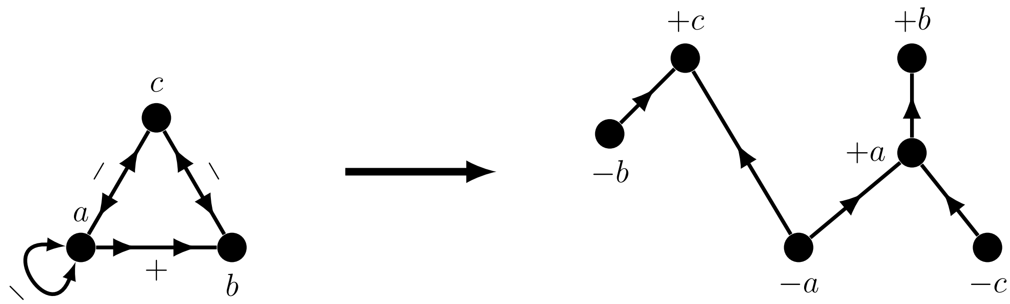

The shift from orientations to posets allows us to formalize an important insight: if two vertices are incomparable, then their colors do not depend on each other at all. We can further get a handle on this incomparability by introducing linear extensions.

Whereas partial orderings can leave some elements incomparable, linear extensions impose a total ordering so that any two elements are comparable (i.e. for any two vertices a, b either $a > b$ or $b > a$). Any poset can be extended to a linear extension that agrees with it by simply ranking each pair of incomparable vertices (while making sure not to introduce any cycles). Because of this choice, posets can typically be extended to more than one linear extension.

So, if you linearly extend a poset P to different linear extensions L_1 and L_2, and L_1 places $a > b$ while L_2 places $b > a$, then a and b were originally incomparable in P.

We can make use of this to consider only linear extensions instead of posets. Specifically, given any orientation P of Σ, consider the set of all of its linear extensions $\mathcal{L}(P)$. We can arbitrarily pick some linear extension ω of P to serve as a base point, and by noticing where every other linear extension of P disagrees with ω, we can recover all information about our initial poset P.

This will allow us to partition the sets $\{\kappa : \kappa \text{ preserves } P\}$ into sets of colorings which correspond to linear extensions of P. Formally, for two linear extensions α and ω, let $\mathcal{K}(\alpha, \omega)$ be the set of all colorings which by and large preserve α, but are allowed to be improper where α and ω disagree.

For example, if we have α and ω such that

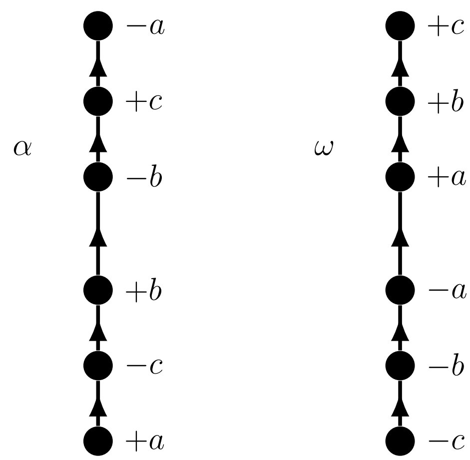

Then $\mathcal{K}(\alpha, \omega)$ is the set of all colorings κ such that $|\kappa(a)| \geq |\kappa(c)| > |\kappa(b)| \geq 0$ with $\kappa(a) \leq 0$, $\kappa(c) > 0$ and $\kappa(b) \leq 0$. With a bit of thought, we can find that $\mathcal{K}(\alpha_1, \omega) \cap \mathcal{K}(\alpha_2, \omega) = \emptyset$ when $\alpha_1 \neq \alpha_2$, which allows us to partition the colorings according to these $\mathcal{K}$ sets.

So, for each orientation P of Σ and for any fixed $\omega \in \mathcal{L}(P)$,

$$\{\kappa : \kappa \text{ preserves } P\} = \bigsqcup_{\alpha \in \mathcal{L}(P)} \mathcal{K}(\alpha, \omega)$$

so that in total we have

$$\begin{aligned} \{\kappa : \kappa \text{ is a proper coloring of } \Sigma\} &= \bigsqcup_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \{\kappa : \kappa \text{ preserves } P\} \\ &= \bigsqcup_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \left(\bigsqcup_{\alpha \in \mathcal{L}(P)} \mathcal{K}(\alpha, \omega) \right) \end{aligned}$$

Let's use this framework to introduce functions in the variables x_i for $i \in \mathbb{Z}$. We'll associate each $\mathcal{K}(\alpha, \omega)$ to a function $F_{\alpha,\omega}$ of degree d, simply by saying $F_{\alpha,\omega} = \sum_{\kappa \in \mathcal{K}(\alpha,\omega)} x^\kappa$, so that $Y_P = \sum_{\kappa \text{ preserves } P} x^\kappa = \sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}$. This lets us re-express the B-symmetric chromatic function:

$$X_\Sigma = \sum_{\kappa \text{ properly colors } \Sigma} x^\kappa = \sum_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} Y_P = \sum_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega} \right).$$

Hence these $F_{\alpha,\omega}$'s confer some immediate advantages. They serve as a basis set for studying B-symmetric chromatic functions while also explicitly connecting colorings to orientations.

And now we introduce a magic function, φ with the property that

$$\varphi(Y_P) = \varphi\left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega} \right) = t^{(\text{Number of sinks of } P)}$$

$\varphi(X_\Sigma) = \sum \left(\#\text{ of orientations with } k \text{ sinks} \right) \cdot t^k$

This function is linear and multiplicative over B-symmetric chromatic functions, and turns them into polynomials of an indeterminate variable, t. The polynomial's coefficients, when read off, are the number of orientations with a certain number of sinks. This is a powerful result that brings us closer to the statement of the main theorem. (Although we haven't met the ξ_n's, $q_{a,b}$'s, or z_n's yet.)

Example. Let's consider a single orientation P, of this signed graph Σ.

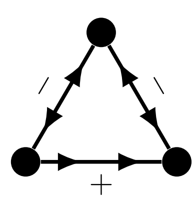

P has an associated B-symmetric function $Y_P = \sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}$ for some fixed $\omega \in \mathcal{L}(P)$. And

$$\varphi(Y_P) = \varphi\left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}\right) = t^{\operatorname{sink} P} = t^2.$$

This was just one orientation of Σ. If we sum over every orientation of Σ, then we would obtain

$$\begin{aligned}
\varphi(X_\Sigma) &= \varphi\left(\sum_{\substack{P \text{ is an acyclic} \\ \text{orientation of } \Sigma}} \left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}\right)\right) \\
&= \operatorname{acyc}_3(\Sigma)t^3 + \operatorname{acyc}_2(\Sigma)t^2 + \operatorname{acyc}_1(\Sigma)t + \operatorname{acyc}_0(\Sigma) \\
&= 0t^3 + 2t^2 + 2t + 2
\end{aligned}$$

indicating that for $k = 0, 1, 2$, there are 2 acyclic orientations of Σ which have k sinks, for a total of 6 acyclic orientations.

Applying φ to X_Σ is very difficult when we do not have a reasonable way of writing X_Σ down. Is there a way to easy calculate X_Σ in terms of simpler functions? To better investigate this, let's return to the unsigned case for a moment.

Recall that for the regular chromatic polynomial, we have the deletion-contraction condition, $\chi_G = \chi_{G-e} - \chi_{G/e}$. We might optimistically hope that this is also true for the symmetric chromatic function, that we may write $X_G = X_{G-e} - X_{G/e}$, but this cannot be correct, since terms in X_{G-e} have larger degree than terms in $X_{G/e}$, and every term in X_G has the same degree.

We can resolve this issue by placing weights on each vertex and performing weighted deletion-contraction. What this means is that every vertex starts with a weight of 1 and whenever an edge is contracted, the weights of the vertices are added to obtain the weight of the new vertex. Then we have that $X_{G^w} = X_{G^w - e} - X_{(G/e)^{w'}}$ where G^w denotes the graph G with weight distribution w and w' is the weight distribution we get by contracting e. When all edges are gone, we replace each vertex v by $p_{(\text{weight of } v)}$ (recall that $p_n = \sum_{k=1}^{\infty} x_k^n$) and multiply together each of the terms we get from the remaining vertices.

For example, let G be the star graph with 3 edges and for convenience we will write just a graph in place of writing the graph in the subscript of X.

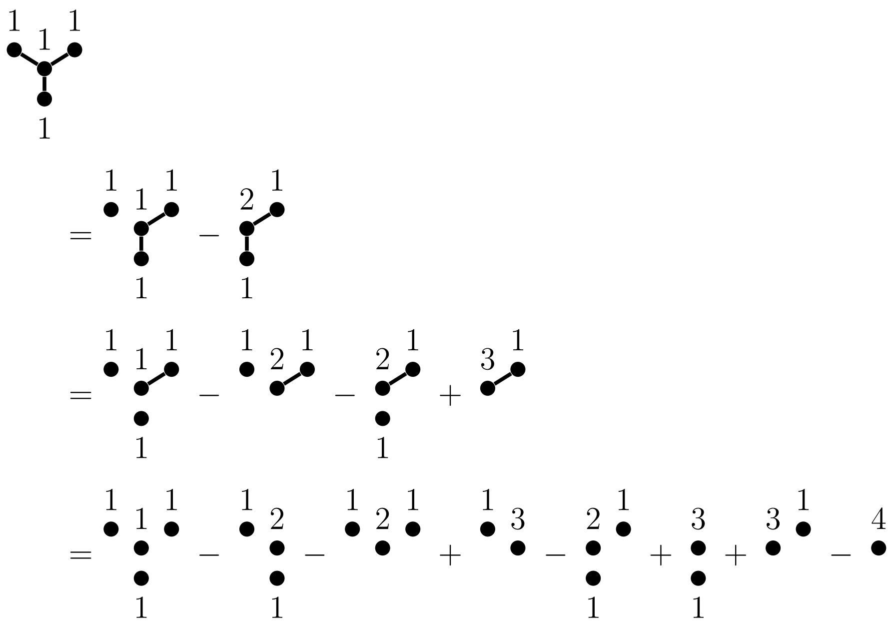

$$=p_1^4 - p_1^2p_2 - p_1^2p_2 + p_1p_3 - p_1^2p_2 + p_1p_3 + p_1p_3 - p_4$$

$$=p_1^4 - 3p_1^2p_2 + 3p_1p_3 - p_4$$

So $X_G = p_1^4 - 3p_1^2p_2 + 3p_1p_3 - p_4$.

In the signed case we would hope to have a similar procedure. A Knots and Graphs group from several years ago (composed of James Enouen, Eric Fawcett, Rushil Raghavan and Ishaan Shaw) found that an almost identical process works for signed graphs, where initially vertices are weighted with $(1, 0)$. Then we have weights add component-wise when contracting over a positive edge and when switching a vertex with weight (a, b), the weight becomes (b, a) (recall that switching a vertex means flipping the sign of all edges it is connected to; this is

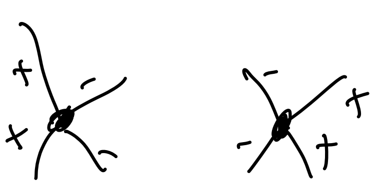

necessary because negative edges cannot be contracted).

When there are no edges left, a vertex with weight (a, b) becomes $p_{a,b}$ where $p_{a,b} = \sum_{i\in\mathbb{Z}} x_i^a x_{-i}^b$ and a vertex with weight (a, b) which has a negative loop becomes $p^*_{a,b} =$ $\sum_{i\in\mathbb{Z}\setminus\{0\}} x_i^a x_{-i}^b = p_{a,b} - x_0^{a+b}$.

For example, take Σ to be the triangle with one negative edge.

Then we have

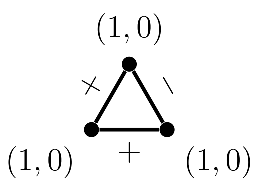

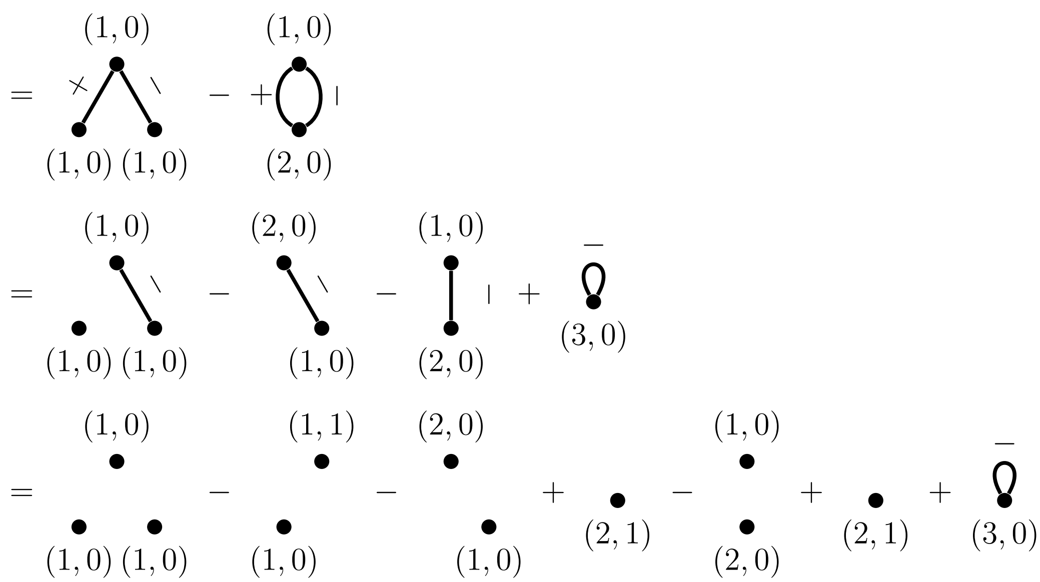

$$=p_{1,0}^3 - p_{1,0}p_{1,1} - p_{1,0}p_{2,0} + p_{2,1} - p_{1,0}p_{2,0} + p_{2,1} + p^*_{3,0}$$

$$=p_{1,0}^3 - p_{1,0}p_{1,1} - 2p_{1,0}p_{2,0} + 2p_{2,1} + p^*_{3,0}$$

So $X_\Sigma = p_{1,0}^3 - p_{1,0}p_{1,1} - 2p_{1,0}p_{2,0} + 2p_{2,1} + p^*_{3,0}$.

Our magic function φ plays nicely enough with the $p_{a,b}$'s. It turns out that we have $\varphi(p_{a,0}) = (t-1)^a - (-1)^a$ and $\varphi(p_{a,b}) = (-1)^{a+b+1}$ for $b \geq 1$. The $q_{a,b}$'s we saw in the theorem are defined (for $a, b \geq 1$) to be $q_{a,b} \stackrel{\text{def}}{=} (-1)^{a+b+1} p_{a,b}$, so that $\varphi(q_{a,b}) = 1$.

It also works well with x_0^n (and so also $p^*_{a,b}$). We have $\varphi(x_0^n) = (-1)^n$, and so $z_n \stackrel{\text{def}}{=} (-1)^n x_0^n$, so that $\varphi(z_n) = 1$.

The only thing we haven't defined is the ξ_n's. $1 + \sum_{a=0}^{n} \binom{n}{a} p_{a,0}$

$$\xi_0 \stackrel{\text{def}}{=} 1$$

$$\xi_n \stackrel{\text{def}}{=} 1 + \sum_{a=0}^{n} \binom{n}{a} p_{a,0} (-1)^{n-a} \quad (p_{0,0} = 1)$$

$$= \sum_{a=0}^{n} \binom{n}{a} (t-1)^a (-1)^a - \sum_{a=0}^{n} \binom{n}{a} (-1)^a$$

Most importantly, $\varphi(\xi_n) = t^n$. Together with the $q_{a,b}$'s, z_n's, and definition of φ, we can pretty easily arrive at the theorem.

Any B-symmetric chromatic function of a signed graph Σ can be uniquely expressed in terms of $p_{a,b}$'s and $p^*_{a,b}$'s, and we have simple substitutions to convert the expression to ξ_n's.

$$p_{1,0} = \xi_1$$

$$p_{2,0} = \xi_2 - 2\xi_1$$

$$p_{3,0} = \xi_3 - 3\xi_2 + 3\xi_1$$

$$\vdots$$

$$p_{a,0} = \sum_{k=0}^{a-1} (-1)^k \binom{a}{k} \xi_{a-k}$$

We know that we can uniquely write X_Σ in terms of ξ_n's, $q_{a,b}$'s and z_n's and we can see that the terms that φ sends to t^m are precisely those of the form $\left(\prod q_{a,b} \cdot \prod z_j\right) \cdot \xi_{n_1} \cdots \xi_{n_k}$ where $n_1 + \cdots + n_k = m$, which gives us the theorem.

$\xi =$ 

$X_\xi =$ 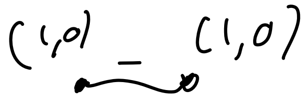 $=$ 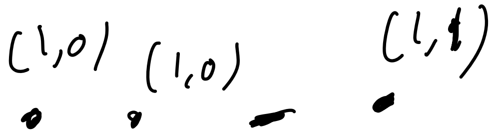

$$= p_{1,0}^2 - p_{1,1}$$

$$= \xi_1^2 - (-\xi_{1,1})$$

$$= \xi_1^2 + \xi_{1,1}$$

$$= (\xi_1 \xi_1) + (\xi_{1,1})$$

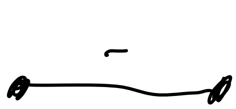

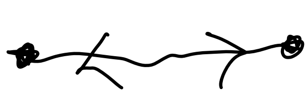

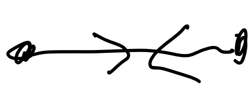

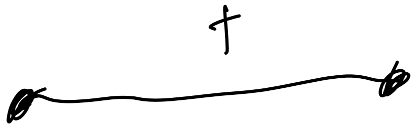

$\chi_\xi =$ 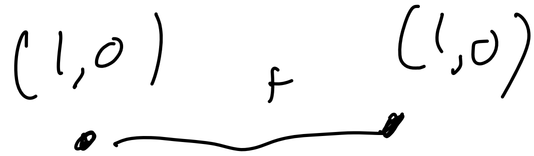

$=$ 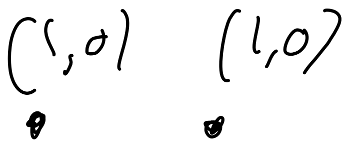 $-$ 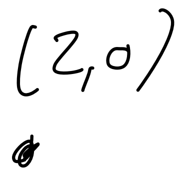

$$= p_{1,0}^2 - p_{2,0}$$

$$= \xi_1^2 - (\xi_2 - 2\xi_1)$$

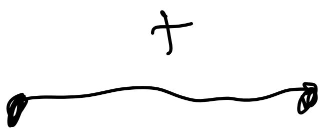 $= (\xi_1^2 - \xi_2) + (2\xi_1)$

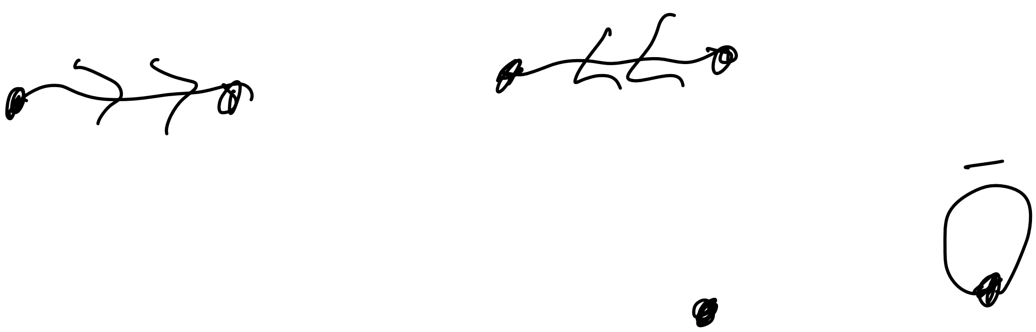

Theorem. *There exists a set of functions*

$$\left\{\xi_n \,\middle|\, n = 0, 1, 2, \ldots\right\} \cup \left\{q_{a,b} \,\middle|\, a, b \geq 1\right\} \cup \left\{z_n \,\middle|\, n = 0, 1, \ldots\right\}$$

in the variables x_i for $i \in \mathbb{Z}$, such that any X_Σ can be uniquely expressed in terms of these functions and when written this way, the sum of the coefficients of the all of the terms of the form $\left(\prod q_{a,b} \cdot \prod z_n\right) \cdot \xi_{n_1} \cdots \xi_{n_k}$, is the number of acyclic orientations of Σ with $n_1 + \cdots + n_k$ sinks.

Key Points from Formal Proof

We will find it very convenient to treat linear extensions of signed posets as functions. If α is a linear extension of some signed poset, we can define $\alpha : \overline{V} \to \{-d, -d+1, \ldots, -1, 1, \ldots, d-1, d\}$ such that $\varepsilon_1 u >_\alpha \varepsilon_2 v \implies \alpha(\varepsilon_1 u) > \alpha(\varepsilon_2 v)$. By the nature of covering graphs and their orientations, we have that $\alpha(\varepsilon v) = \varepsilon\alpha(+v)$.

Additionally, we define $\operatorname{sgn}_\alpha(v) = \dfrac{|\alpha(+v)|}{\alpha(+v)}$.

For any ω, let $\mathrm{BSymm}_d = \mathrm{span}(\{F_{\alpha,\omega} : \alpha \text{ is a linear extension of } d \text{ vertices}\})$

Definition. Fix any total ordering ω. Given a total ordering α, relabel the vertices such that $|\alpha(v_i)| = i$ for all i. Also, let $\varepsilon_i = \operatorname{sgn}_\alpha(v_i)$. Let $\varphi : \mathrm{BSymm}_d \to \mathbb{Q}[t]$ be a linear function which satisfies

$$\varphi(F_{\alpha,\omega}) = \begin{cases} t(t-1)^k & \text{if } \omega(v_{n-k}) > \omega(v_{n-k+1}) > \cdots > \omega(v_n),\ \varepsilon_i = + \text{ for } i \in [n-k, n] \\ & \text{and } 0 < \varepsilon_1\omega(v_1) < \varepsilon_2\omega(v_2) < \cdots < +\omega(v_{n-k});\ \text{for } 0 \le k < n \\ (t-1)^k & \text{if } -\omega(v_{n-k}) > \omega(v_{n-k+1}) > \cdots > \omega(v_n),\ \varepsilon_i = + \text{ for } i \in [n-k+1, n], \\ & \varepsilon_{n-k} = - \text{ and } 0 < \varepsilon_1\omega(v_1) < \varepsilon_2\omega(v_2) < \cdots < -\omega(v_{n-k});\ \text{for } 0 \le k < n \\ (t-1)^n & \text{if } 0 > \omega(v_1) > \omega(v_2) > \cdots > \omega(v_n) \text{ and } \varepsilon_i = + \text{ for all } i \\ 0 & \text{otherwise.} \end{cases}$$

We will take for granted that φ is well defined (this follows from the fact that the set of $F_{\alpha,\omega}$'s is linearly independent and whenever we have $F_{\alpha,\omega} = F_{\beta,\omega}$, we also have $\varphi(F_{\alpha,\omega}) = \varphi(F_{\beta,\omega})$).

Definition. Given a signed poset P, let the number of positive maximal elements be denoted $\operatorname{sink}(P)$. This is the same as the number of vertices which have no arrows coming pointing away from them.

Lemma. *Let P be a signed poset. Then $\varphi(Y_P) = t^{\operatorname{sink}(P)}$.*

Proof. Given some poset P, fix your favorite linear extension ω of P.

We will begin by noting that in the first case of φ the maximal element under ω is positive, in the second case the maximal element under ω is negative with $k < n$ and the third case is what we would get by using $k = n$ in the second case along with the convention $v_0 = 0$.

It is also important to note that the vertex v_j in P is a sink iff $+v_j$ is the maximal element of some linear extension of P.

Consider the case where the largest element under ω is positive. That means that this vertex (call it s) must be a sink in P. Now select any k of the remaining sinks of P other than s (there are $\mathrm{sink}(P) - 1$ to choose from) and call these vertices $u_1, u_2, \ldots, u_k$.

Now for the remaining $n - k - 1$ vertices we will label then $v_1, v_2, \ldots, v_{n-k-1}$ such that $|\omega(v_i)| < |\omega(v_j)|$ iff $i < j$.

Consider the linear extension α, which has $\alpha(u_i) = n - i + 1$, $\alpha(s) = n - k$ and for $i < n - k$, $|\alpha(v_i)| = i$ with $\mathrm{sgn}_\alpha(v_i) = \mathrm{sgn}_\omega(v_1)$.

i.e. $+u_1$ is the largest element under α, $+u_2$ is the second largest and so on. Then $+s$ is directly under $+u_k$ and under $+s$, the vertices are arranged based on their ordering under ω and each with the same sign as under ω.

Note that after initially choosing the k sinks, no more choices are made, meaning that any α constructed this way is unique when given a choice of k sinks.

To see that α is a linear extension of P, we will suppose for the sake of contradiction that there is some relation in P that α violates.

Note that for any two vertices other than $u_1, \ldots, u_k$ or s, the absolute value of α respects the ranking of the absolute value of ω, meaning that $|\alpha(v_i)| < |\alpha(v_j)|$ iff $|\omega(v_i)| < |\omega(v_j)|$. In addition the signs of these vertices under α is the same as under ω. In particular, this means that α respects all relations of P, except possibly those involving $u_1, \ldots, u_k$ or s.

So the relation that α violates must contain at least one of $u_1, \ldots, u_k$ or s. The relation must contain at most one of these vertices, since they are all sinks and so there is no directed positive edge or inward facing directed negative edge between any two of them. If there is an outward facing negative edge between two of them, this translates to the relation

$\alpha(p) > -\alpha(q)$ for $p, q \in \{u_1, \ldots, u_k\} \cup \{s\}$. But this relation holds in α, since $\alpha(p) > 0$ for any $p \in \{u_1, \ldots, u_k\} \cup \{s\}$.

So this relation must contain exactly one of $u_1, \ldots, u_k$ or s and exactly one non-sink (the case where it contains a sink other than $u_1, \ldots, u_k$ or s is impossible for the same reasons as before). Call the sink it contains p and the non-sink r. Then the edges it could come from are a directed positive edge from r to p or an inward facing directed negative edge. These directed edges invoke the relations $\alpha(p) > \alpha(r)$ and $\alpha(p) > -\alpha(r)$ respectively. In either case we can see that α satisfies these relations since $|\alpha(p)| > |\alpha(r)|$ by construction.

This is a contradiction and so α is a linear extension of P.

So we have seen that for any $k < \operatorname{sink}(P)$ and for each choice of k sinks of P (other than s), there is exactly one linear extension α which satisfies the first case of φ. This means that there are $\binom{\operatorname{sink} P - 1}{k}$ linear extensions α for which $\varphi(F_{\alpha,\omega}) = t(t-1)^k$. This holds for all $k \leq \operatorname{sink} P - 1$ and any α not of this form must have $\varphi(F_{\alpha,\omega}) = 0$ by our observation at the beginning of the proof.

So then we have that

$$\begin{aligned}
\varphi(Y_P) &= \varphi\left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}\right) \\
&= \sum_{k=0}^{\operatorname{sink} P - 1} \binom{\operatorname{sink} P - 1}{k} t(t-1)^k \\
&= t \cdot \sum_{k=0}^{\operatorname{sink} P - 1} \binom{\operatorname{sink} P - 1}{k} (t-1)^k \\
&= t \cdot t^{\operatorname{sink} P - 1} = t^{\operatorname{sink} P}
\end{aligned}$$

The second to last equality easily follows from the binomial theorem. This prove the theorem in the case where the maximal element under ω is positive.

In the case where the maximal element under ω is negative the argument is almost identical, except that we have $\operatorname{sink} P$ vertices to choose from when selecting $u_1, \ldots, u_k$. The verification that the previous construction still works is left as an exercise for the reader. Note that the third case of φ occurs when $\operatorname{sink} P = n$.

So for any $k \leq \operatorname{sink} P$, and for each choice of k sinks of P, there is exactly one linear extension α which satisfies the second or third case of φ. This means that there are $\binom{\operatorname{sink} P}{k}$ linear extensions α for which $\varphi(F_{\alpha,\omega}) = (t-1)^k$. This holds for all $k \leq \operatorname{sink} P$ and any α not of this form must have $\varphi(F_{\alpha,\omega}) = 0$ by our observation at the beginning of the proof.

So then we have that

$$\begin{aligned}
\varphi(Y_P) &= \varphi\left(\sum_{\alpha \in \mathcal{L}(P)} F_{\alpha,\omega}\right) \\
&= \sum_{k=0}^{\operatorname{sink} P} \binom{\operatorname{sink} P}{k} (t-1)^k \\
&= \sum_{k=0}^{\operatorname{sink} P} \binom{\operatorname{sink} P}{k} (t-1)^k \\
&= t^{\operatorname{sink} P}
\end{aligned}$$

This completes the proof. $\square$