Combinatorial Reciprocity: Counting Through the Looking-Glass

Will Dana

August 2–3, 2022

1 The chromatic polynomial

For the purposes of this class, a **graph** is a finite collection of vertices V and a collection of edges E, where each edge corresponds to a pair of vertices. We do not allow loops or multiple edges between vertices.

Definition. *A **proper coloring with** n **colors** of a graph is an assignment of a label from $\{1, \ldots, n\}$ to every vertex such that no two adjacent vertices have the same color.*

Example. Below are two colorings. The one on the left is proper, but the one on the right is not, because two adjacent vertices are both a rich hue of 3.

Some of the interest in proper colorings comes from the infamous 4-Color Theorem, which states that any graph that can be drawn without edge crossings can be properly colored with at most 4 colors. But there's also interest in counting how many different ways there are to do this. Let $\chi_G(n)$ be the number of ways of properly coloring a graph G with n colors. It's called the **chromatic polynomial**, because it turns out to be a polynomial.

Proposition 1. *$\chi_G(n)$ is a polynomial in n.*

Proof. We can prove this inductively, based on the number of edges of G. If there are no edges at all, then $\chi_G(n) = n^{\#V}$, where $\#V$ is the number of vertices.

Now for an arbitrary graph G, pick an edge e. Define two new graphs: the **deletion** $G - e$, obtained by removing e, and the **contraction** G/e, obtained by collapsing the two endpoints of e into a single vertex.

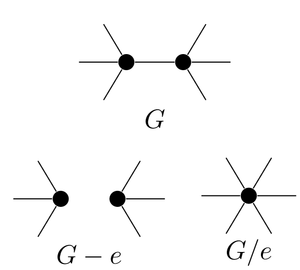

Figure 1: The deletion and contraction operations.

Then you can show the following relationship (exercise!):

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

Since the terms on the right are polynomials by the induction hypothesis, so is the term on the left. □

Example. Sometimes in simple cases we can calculate this polynomial directly. For example, consider the graph

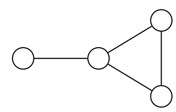

We can start coloring it by picking one of n colors for the left vertex. The middle vertex can't have the same color, but we can pick any other one, for a total of $n-1$ options. Similarly, we can pick any color for the upper right vertex which doesn't match the middle one, for a total of $n-1$ options. Finally, we need to avoid the colors of both the middle and upper right vertices when coloring the lower right one, so there are $n-2$ options. Putting it all together, there are $n(n-1)^2(n-2)$ ways we could color this graph.

2 Through the looking-glass

There are a lot of things we can do with polynomials. For instance, as defined the function $\chi_G(n)$ only makes sense for positive integer n. But knowing that it's a polynomial, we can plug in other numbers. For example, what does the number $\chi_G(-1)$ tell us?

I love getting sensible information from a nonsensical input, which is why the answer to this question is one of my all-time favorite theorems.

Theorem 1. *Say that an **acyclic orientation** of a graph assigns to each edge a direction such that there are no oriented cycles (i.e., any path following the arrows must eventually stop.) Then $(-1)^{\#V}\chi_G(-1)$ is the number of acyclic orientations of G.*

Example. We can once again look at the graph above. The total number of orientations (acyclic or otherwise) is $2^4 = 16$: we make a binary choice of orientation for each of the 4 edges. But there are $2 \cdot 2 = 4$ orientations which contain cycles: we can choose an arbitrary orientation for the left edge, and then either a clockwise or counterclockwise cycle on the right. Thus there are $16 - 4 = 12$ acyclic orientations, and indeed,

$$(-1)^4\chi_G(-1) = (-1)(-2)^2(-3) = 12.$$

You can prove this theorem inductively using the same kind of argument as we used to show that $\chi_G(n)$ is a polynomial. But why stop here — why not try plugging in other negative integer values[1]?

We can begin to see what such a result might look like by trying to connect proper colorings to acyclic orientations. In fact, there is a natural map from proper colorings to acyclic orientations: orient each edge so it points from the higher color to the lower one:

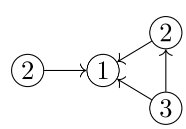

This perspective makes it a little clearer how the following, more general theorem connects back to the chromatic polynomial.

Theorem 2. *Let $\overline{\chi}_G(n)$ count the number of ways of choosing the following:*

[1]You could also try plugging in nonintegral, even complex values, which is way outside the scope of this class. However, there has been some interesting work on the roots of chromatic polynomials.

- *An acyclic orientation of G.*
- *A labeling σ of the vertices of G by $\{1, \dots, n\}$ such that if $v_1 \to v_2$ in our orientation, then $\sigma(v_1) \geq \sigma(v_2)$.*

Then

$$\overline{\chi}_G(n) = (-1)^{\# V} \chi_G(-n)$$

Notice how this generalizes the theorem above: if $n = 1$, then each acyclic orientation appears in this count exactly once, as we're forced to apply the same color to everything.

Note that the labelings considered in this theorem are, in a sense, a slightly looser version of proper coloring, where we don't require that labels strictly decrease in the direction of edges, just that they don't increase. This is just one instance of a pattern that shows up repeatedly across many combinatorial problems.

3 Subsets of a set

The simplest example of this pattern comes from binomial coefficients: the number of ways of choosing k distinct objects from a set of n is given by $\binom{n}{k}$. If we fix k, this is a polynomial expression in n. So we can try plugging in a negative number, as before.

$$\begin{aligned}\binom{-n}{k} &= \frac{(-n)(-n-1)\cdots(-n-k+1)}{k!} = (-1)^k \frac{(n)(n+1)\cdots(n+k-1)}{k!} \\ &= (-1)^k \binom{n+k-1}{k}\end{aligned}$$

It turns out that the binomial coefficient on the right side is also important.

Theorem 3. $\binom{n+k-1}{k}$ *counts the number of ways of selecting k* ***not necessarily distinct*** *objects from among n.*

Proof. Suppose that we've selected k numbers, potentially with repetition, from among $\{1, \dots, n\}$. Suppose that among these numbers there are k_1 1's, k_2 2's,..., and k_n n's. We can encapsulate this information in a diagram consisting of k_1 dots, followed by a divider, followed by k_2 dots, followed by a divider, and so on. Consider, for instance, the case $n = 4$:

$$1, 1, 2, 4, 4, 4 \mapsto \circ \circ \mid \circ \mid \mid \circ \circ \circ$$

On the other hand, how many diagrams of the form on the right can occur? There will be k dots and $n - 1$ dividers sorting them into n buckets, so $n + k - 1$ symbols in total. We specify the diagram by choosing which k of these symbols are dots, and there are $\binom{n+k-1}{k}$ ways of doing this. □

So again, plugging in a negative number also counts something, but with a tweak: instead of requiring all the elements we select to be distinct, we're relaxing that condition.

4 Polytopes and Ehrhart polynomials

Here's another, more complicated situation in which plugging in negative numbers gives us good information. Consider a polytope[2] whose vertices all have integer coordinates. The motivating situation here is that there are two different ways of measuring how big this polytope is:

- we could measure its volume, or
- we could count how many points with integer coordinates it contains.

[2] Essentially, a higher-dimensional version of a polygon or polyhedron. If you're not familiar with this idea, you can just think about polygons and polyhedra.

How are this continuous measurement and this discrete one related to each other?

The secret turns out to be that we should perform the point count for different **dilations** of a polytope. Given a polytope $\mathcal{P}$, denote by $n\mathcal{P}$ the new polytope where we've scaled up all the coordinates of the vertices by a factor of n.

We'd like to count how the number of points changes as we scale up the polytope, though there's one decision we need to make before we start counting: are we including the points on the boundary of the polytope, or just the ones that are entirely contained within it? For reasons that will become clear shortly, I'm going to define a function for each one.

So let

$$j(\mathcal{P}, n) = \text{\# of points with integer coordinates in } n\mathcal{P}\text{, including its boundary}$$
$$i(\mathcal{P}, n) = \text{\# of points with integer coordinates in } n\mathcal{P}\text{, excluding its boundary}$$

Example. Let's consider the unit square, with vertices $(0,0)$, $(1,0)$, $(0,1)$, and $(1,1)$. If we scale it up by a factor of n, we get an $n \times n$ square looking like this (here, $n = 4$):

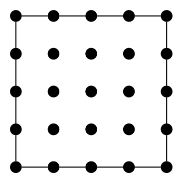

If we include the points around the edge, we get $j(\mathcal{P}, n) = (n+1)^2$. If we don't include those points, we get $i(\mathcal{P}, n) = (n-1)^2$.

In general, it is a fact that both $j(\mathcal{P}, n)$ and $i(\mathcal{P}, n)$ are polynomials. This gives a neat, if roundabout, way of computing volume. Asymptotically, as n gets big both $i(\mathcal{P}, n)$ and $j(\mathcal{P}, n)$ will behave like $\mathrm{Vol}(\mathcal{P}) n^{\dim \mathcal{P}}$, so we know this is the leading coefficient. On the other hand, since they are polynomials of degree $\dim \mathcal{P}$, we can pin them down exactly by computing them at $\dim \mathcal{P} + 1$ values, which can be done just by counting points.

However, as should be foreshadowed by now, we're mainly interested in what happens if we plug $-n$ into these polynomials. For example, in this case

$$j(\mathcal{P}, -n) = (-n+1)^2 = (n-1)^2 = i(\mathcal{P}, n)$$

And in fact this is a general theorem:

Theorem 4 (Ehrhart-Macdonald Reciprocity)**.**

$$j(\mathcal{P}, -n) = (-1)^{\dim \mathcal{P}} i(\mathcal{P}, n)$$

5 Commonalities

What do the three results above have in common?

- In the case of the Ehrhart polynomial, plugging in $-n$ toggles whether we count the boundary points.
- In the case of the binomial coefficient, plugging in $-n$ toggles whether we allow repeated items in our selection.
- In the case of the chromatic polynomial, plugging in $-n$ toggles whether we require colors to be strictly decreasing or simply nonincreasing along arrows in our acyclic orientation.

In all three cases, we have counting functions related by whether the structure we're counting is defined using strict or nonstrict inequalities. This sets the stage for a big theorem which ends up generalizing all of these.

6 Stanley's reciprocity theorem

Theorem 5. *Suppose we have a homogeneous[3] system of linear equations with integer coefficients*

$$\begin{aligned} E_1(x_1,\ldots,x_s) &= 0 \\ E_2(x_1,\ldots,x_s) &= 0 \\ &\vdots \\ E_r(x_1,\ldots,x_s) &= 0 \end{aligned}$$

Let N be the set of all solutions to this system using ***nonnegative*** *integers, and let P be the set of all solutions using* ***positive*** *integers. Define two power series:*

$$\begin{aligned} F(X_1,\ldots,X_s) &= \sum_{(\alpha_1,\ldots,\alpha_s)\in N} x_1^{\alpha_1}\cdots x_s^{\alpha_s} \\ \overline{F}(X_1,\ldots,X_s) &= \sum_{(\beta_1,\ldots,\beta_s)\in P} x_1^{\beta_1}\cdots x_s^{\beta_s} \end{aligned}$$

Then:

(1) F and $\overline{F}$ define rational functions.

(2) If P is nonempty (i.e., $\overline{F}\neq 0$), then

$$F(X_1,\ldots,X_s) = (-1)^\kappa \overline{F}\left(\frac{1}{X_1},\ldots,\frac{1}{X_s}\right)$$

where κ is the ***nullity*** *of the system (the dimension of its solution space over $\mathbb{R}$).*

Example. We can see what's going on here when there's just one equation: suppose $E_1(x_1,x_2) = x_1 - x_2$. Then N is just the set of pairs $(0,0)$, $(1,1)$, $(2,2)$,... while P is almost the same, but omits $(0,0)$.

If we work out the two power series above, we get geometric series.

$$\begin{aligned} F(X_1,X_2) &= 1 + X_1X_2 + X_1^2X_2^2 + \ldots = \frac{1}{1-X_1X_2} \\ \overline{F}(X_1,X_2) &= X_1X_2 + X_1^2X_2^2 + \ldots = \frac{X_1X_2}{1-X_1X_2} \end{aligned}$$

And indeed, when we plug in the reciprocals as suggested by the theorem, things rearrange nicely:

$$F\left(\frac{1}{X_1},\frac{1}{X_2}\right) = \frac{1}{1-\frac{1}{X_1X_2}} = \frac{X_1X_2}{X_1X_2-1} = -\frac{X_1X_2}{1-X_1X_2} = -\overline{F}(X_1,X_2)$$

This is the sign we expect, because we used 1 equation in 2 variables: the dimension of the solution space is $2-1=1$.

I claim that this theorem is somehow a generalization of *all* of the results described above. It looks similar if you squint, but there are some big gaps: we need to somehow get from subsituting $1/X$ for X to substituting $-n$ for n, and we need to somehow take the tremendous amount of information contained in these power series and boil them down to counting things.

The first of these is accomplished by a lemma.

Lemma 1. *Let $H(i)$ be a polynomial. Define*

$$F(X) = \sum_{r=0}^{\infty} H(r)X^r \qquad \overline{F}(X) = \sum_{r=1}^{\infty} H(-r)X^r$$

Then $F(X)$ and $\overline{F}(X)$ are rational functions, and $F(X) = -\overline{F}(1/X)$.

[3] That is, there are no constant terms.

Proof. One way to prove this is to start with the case that $H(r) = \binom{r}{k}$ for some k, in which case F and $\overline{F}$ are particularly well-behaved.

Specifically, if we start from the geometric series

$$\frac{1}{1-X} = \sum_{r=0}^{\infty} X^r$$

and take the derivative of both sides, we get

$$\frac{1}{(1-X)^2} = \sum_{r=1}^{\infty} rX^{r-1} = \sum_{r=0}^{\infty} (r+1)X^r$$

If we go on to do this k times, we get

$$\frac{k!}{(1-X)^{k+1}} = \sum_{r=0}^{\infty} (r+1)(r+2)\cdots(r+k)X^r.$$

(exercise: check this!)

Then dividing both sides by $k!$, we get

$$\frac{1}{(1-X)^{k+1}} = \sum_{r=0}^{\infty} \binom{r+k}{k} X^r$$

From here, we can take two paths. If we multiply both sides by X^k and reindex to push the first term back to 0, we get

$$\frac{X^k}{(1-X)^{k+1}} = \sum_{r=0}^{\infty} \binom{r+k}{k} X^{r+k} = \sum_{r'=0}^{\infty} \binom{r'}{k} X^{r'}$$

which is thus $F(X)$. On the other hand, if we multiply both sides by X and reindex, we get

$$\frac{X}{(1-X)^{k+1}} = \sum_{r=0}^{\infty} \binom{r+k}{k} X^{r+1} = \sum_{r'=1}^{\infty} \binom{r'+k-1}{k} X^{r'}$$

Importantly, we showed above that $\binom{r'+k-1}{k} = (-1)^k \binom{-r'}{k}$, so

$$\overline{F}(X) = (-1)^k \frac{X}{(1-X)^{k+1}}$$

Then, making the substitution in the theorem,

$$-\overline{F}(1/X) = (-1)^{k+1} \frac{1/X}{(1-1/X)^{k+1}} = (-1)^{k+1} \frac{X^k}{(X-1)^{k+1}} = \frac{X^k}{(1-X)^{k+1}} = F(X)$$

as required.

We can then build other polynomials as linear combinations of binomial coefficients and thus prove that the identity works for them as well; the details of this are left as an exercise. □

With this lemma in hand, we can move to proving more specific results, including the ones above. In each case, by choosing a particular system of linear equations tailored to the situation, we can come up with a generating function which condenses down to count the quantity we want.

7 Combinatorial reciprocity and the Ehrhart polynomial

First, assume without loss of generality that the vertices of our polytope $\mathcal{P}$ have *positive* coordinates; we can just shift it over if necessary, which won't change the Ehrhart polynomial.

There are a couple different ways of defining polytopes, but one way which works is to use inequalities: for each face of the polytope, we can write down an equation for the plane containing that face, then use an inequality to specify that the polytope has to be on one side of that plane. Suppose that the inequalities defining our polytope look like this:

$$\begin{aligned} P_1(x_1, \ldots, x_s) &\le c_1 \\ P_2(x_1, \ldots, x_s) &\le c_2 \\ &\vdots \\ P_r(x_1, \ldots, x_s) &\le c_r \end{aligned}$$

How can we set up a system of equations where the integer solutions capture the points in scaled versions of this polytope?

The system we want. Let's introduce one more variable y, which will represent our scale factor, and then r additional "dummy variables" $z_1, \ldots, z_r$. Then consider the equations

$$\begin{aligned} P_1(x_1, \ldots, x_s) - c_1 y + z_1 &= 0 \\ P_2(x_1, \ldots, x_s) - c_2 y + z_2 &= 0 \\ &\vdots \\ P_r(x_1, \ldots, x_s) - c_r y + z_r &= 0 \end{aligned}$$

What does a solution in nonnegative integers mean? We choose a scale factor y and some nonnegative coordinates $x_1, \ldots, x_s$ such that for each index $1 \le i \le r$,

$$P_i(x_1, \ldots, x_s) - c_i y = -z_i \le 0$$

But since there are no other constraints on the z_i, any point which satisfies the inequalities $P_i(x_1, \ldots, x_s) \le c_i y$ will contribute a solution — and these are exactly the points in $y\mathcal{P}$, including the boundary. (Importantly, because we specified our polytope must already consist of points with positive coordinates, the requirement that the x_i are nonnegative is irrelevant.)

Notice in particular how, using the dummy variables z_i, we turned inequalities into equalities — we're not concerned with the actual values of z_i, only that they are nonnegative. This will be used in all the examples that follow.

What does a solution in positive integers mean? We choose a scale factor y and some positive coordinates $x_1, \ldots, x_s$ such that for each index $1 \le i \le r$,

$$P_i(x_1, \ldots, x_s) - c_i y = -z_i < 0$$

That is, we must satisfy the inequalities $P_i(x_1, \ldots, x_s) < c_i y$ — this is like the above, but now we're not allowed to include the faces of the polytopes. So this gives us points of $y\mathcal{P}$ which are not in the boundary.

As it stands, if we worked with the power series F and $\overline{F}$ from the reciprocity theorem, they would have too much information — they would list out every single point, while we just want to count the number of points for each scale factor. But there's a way around this, by setting most of the variables in our power series equal to 1.

Extracting the information we want. Let Y be the variable in the power series F and $\overline{F}$ whose exponent tracks the scale factor y. If we set all the other variables equal to 1, we can collect together all the terms corresponding to a given scale factor — and the number of such terms is precisely the number of points in the appropriately scaled polytope. Specifically,

$$f(Y) := F(1,\ldots,1,Y,1,\ldots,1) = \sum_{y=0}^{\infty} j(\mathcal{P},y)Y^y$$
$$\overline{f}(Y) := \overline{F}(1,\ldots,1,Y,1,\ldots,1) = \sum_{y=1}^{\infty} i(\mathcal{P},y)Y^y$$

Applying the reciprocity theorem. The theorem tells us that

$$f(1/Y) = (-1)^{\kappa}\overline{f}(Y)$$

where in this case κ is equal to 1 plus the dimension of the polytope, since we took the inequalities defining the polytope and added one more free variable in the form of y. But then Lemma 1 implies that

$$j(\mathcal{P},y) = (-1)^{\dim \mathcal{P}} i(\mathcal{P},-y)$$

which is precisely Ehrhart-Macdonald reciprocity!

8 Combinatorial reciprocity and binomial coefficients

It's kind of overpowered compared to our ad hoc proof above, but the binomial coefficients also crumble before the mighty power of the reciprocity theorem!

The system we want. Let's introduce k variables $x_1,\ldots,x_k$, which will correspond to the elements we choose from our set, a variable y which will specify the size of our set, and dummy variables $z_1,\ldots,z_{k-1},w_1,\ldots,w_k$. (As above, the purpose of these variables is to turn inequalities into equalities.) Then consider the equations

$$\begin{aligned}
x_1 - x_2 + z_1 &= 0\\
x_2 - x_3 + z_2 &= 0\\
&\vdots\\
x_{k-1} - x_k + z_{k-1} &= 0\\
x_1 - y + w_1 &= 0\\
x_2 - y + w_2 &= 0\\
&\vdots\\
x_k - y + w_k &= 0
\end{aligned}$$

What does a solution in nonnegative integers mean? We choose an integer y and some integers $x_1,\ldots,x_k$, which are:

- in nondecreasing order (because $x_i - x_{i+1} = -z_i \le 0$ for each i)
- all between 0 and y inclusive (because they are nonnegative themselves, and $x_i = y - w_i \le y$ for each i)

But this is the same thing as a selection of k unordered elements, potentially with repetition, from among the $y+1$ options $0,\ldots,y$.

What does a solution in positive integers mean? Now the $x_1, \ldots, x_k$ are

- in increasing order, and
- all between 0 and y exclusive.

But this is the same thing as a selection of k unordered, distinct elements from among the $y-1$ options $1, \ldots, y-1$.

Extracting the information we want. To clear up the notation a little bit, let $\widetilde{\binom{n}{k}}$ denote the number of ways of selecting k not necessarily distinct elements from n options. Then as in the Ehrhart polynomial case, we can focus on just the variable Y in our power series corresponding to y, set all the others to 1, and collect like terms to count the number of solutions for a fixed y. Specifically,

$$f(Y) := F(1, \ldots, 1, Y, 1, \ldots, 1) = \sum_{y=0}^{\infty} \widetilde{\binom{y+1}{k}} Y^y$$
$$\overline{f}(Y) := \overline{F}(1, \ldots, 1, Y, 1, \ldots, 1) = \sum_{y=1}^{\infty} \binom{y-1}{k} Y^y$$

Applying the reciprocity theorem. The theorem tells us that

$$f(1/Y) = (-1)^{\kappa} \overline{f}(Y)$$

where in this case $\kappa = (3k) - (2k-1) = k+1$ — we have $2k-1$ independent equations in $3k$ variables. But then Lemma 1 implies that

$$\widetilde{\binom{y+1}{k}} = (-1)^k \widetilde{\binom{-y-1}{k}}$$

which is exactly the result we saw above through simpler means.

9 Combinatorial reciprocity and chromatic polynomials

Finally, we return to the theorem that got me delving into this subject in the first place. It's a little more complicated, but follows the same basic principle. The trick is to fix an orientation, count only the colorings corresponding to that orientation, prove the identity for each of those counts separately, and then add them all up.

To that end, let ω be a particular acyclic orientation of the edges of the graph G. Define

$$\chi_{G,\omega}(n) = \text{number of labelings } \sigma \text{ by } 1, \ldots, n \text{ such that if } v_1 \to v_2 \text{ in } \omega,\ \sigma(v_1) > \sigma(v_2)$$
$$\overline{\chi}_{G,\omega}(n) = \text{number of labelings } \sigma \text{ by } 1, \ldots, n \text{ such that if } v_1 \to v_2 \text{ in } \omega,\ \sigma(v_1) \geq \sigma(v_2)$$

If we take the sum of $\overline{\chi}_{G,\omega}(n)$ over all acyclic orientations ω, we get the function $\overline{\chi}_G(n)$ from Theorem 2 by definition. If we take the sum of $\chi_{G,\omega}(n)$ over all acyclic orientations ω, we get $\chi_G(n)$ — all the labelings concerned are proper colorings, and each proper coloring is compatible with exactly one acyclic orientation.

So it will suffice to show that $\overline{\chi}_{G,\omega}(n) = (-1)^{\#V} \chi_{G,\omega}(-n)$.

The system we want. We'll use variables x_v and w_v indexed by the vertices; x_v will correspond to the colors of the vertices, while w_v will be dummy variables turning inequalities into equalities. We'll use a dummy variable z_e for every edge. Finally, we'll use an additional variable y, which will keep track of the number of colors we're using. Then consider the equations

$$x_{v_2} - x_{v_1} + z_e = 0 \text{ for each edge } v_1 \xrightarrow{e} v_2$$
$$x_v - y + w_v = 0 \text{ for each vertex } v$$

What does a solution in nonnegative integers mean? We're selecting a label x_v for each vertex, and because of each of the constraints

$$x_{v_2} - x_{v_1} = -z_e \leq 0$$

we know that the labels don't increase in the direction of edges. Additionally, because

$$x_v = y - w_v \leq y$$

and the x_v are themselves nonnegative, we know that each label lies between 0 and y inclusive. For a fixed y, we get $\overline{\chi}_{G,\omega}(y+1)$ solutions.

What does a solution in positive integers mean? Now we require that the labels are decreasing in the directions of edges, and that they are between 0 and y exclusive. For a fixed y, we get $\chi_{G,\omega}(y-1)$ solutions.

Extracting the information we want. As above, we set every variable in the power series F and $\overline{F}$, except for the variable Y whose exponent tracks the parameter y, equal to 1. Collecting together all of the terms for each fixed y, we get

$$f(Y) := F(1, \ldots, 1, Y, 1, \ldots, 1) = \sum_{y=0}^{\infty} \overline{\chi}_{G,\omega}(y+1) Y^y$$
$$\overline{f}(Y) := \overline{F}(1, \ldots, 1, Y, 1, \ldots, 1) = \sum_{y=1}^{\infty} \chi_{G,\omega}(y-1) Y^y$$

Applying the reciprocity theorem. The theorem tells us that

$$f(1/Y) = (-1)^{\kappa} \overline{f}(Y)$$

where in this case $\kappa = (\#E + 2\#V + 1) - (\#E + \#V) = \#V + 1$ — again subtracting off the number of equations from the number of variables. But then Lemma 1 implies that

$$\overline{\chi}_{G,\omega}(y+1) = (-1)^{\#V} \chi_{G,\omega}(-y-1)$$

which is exactly what we wanted!

10 Conclusions

We were faced with 3 different combinatorial polynomials which all exhibited weirdly similar behavior when we plugged in negative numbers — somehow expressing the distinction between "strict" and "non-strict" solutions.

The key feature binding all these examples together is that they can all be described as counting integer solutions to a system of equations — and Richard Stanley managed to show that *all* such problems have a similar kind of reciprocity property. And in the case of chromatic polynomials, expressing them from this perspective reveals that the connection to acyclic orientations isn't that strange after all.

References

[1] Richard Stanley. "Linear homogeneous diophantine equations and magic labelings of graphs." Duke Mathematical Journal, vol. 40 (1973), issue 3, p. 607–632.

[2] Richard Stanley. "Combinatorial Reciprocity Theorems." Advances in Mathematics, vol. 14 (1974) p. 194–253.