Discrete Mathematics lecture notes 13-1

December 6, 2013

33. Colorings and the cycle index

In the Lemma that is not Burnside's, we now have a powerful tool for counting orbits of a permutation group acting on a set. In these last two lectures, we'll look at a specific type of set and specific type of action: Instead of simply considering a permutation group G on a set X, we will examine how G acts on the set of *colorings* of X, and as a result obtain a formula for the number of essentially distinct colorings possible. The advantage of this point of view is that we will unify what had formerly been several seemingly unrelated computations, hinting at a deeper structure encoded within the group theory we've developed.

Let's start out by fixing a set X of size n and a permutation group $G \leq \Sigma_X$. We want to color the points of X, say red, blue, green, ..., and ask how many distinct colorings are possible, up to the symmetries of G. For example, if X is the corners of the square and we have only the colors black and white to choose from, there are a total of $2^4 = 16$ different possible colorings:

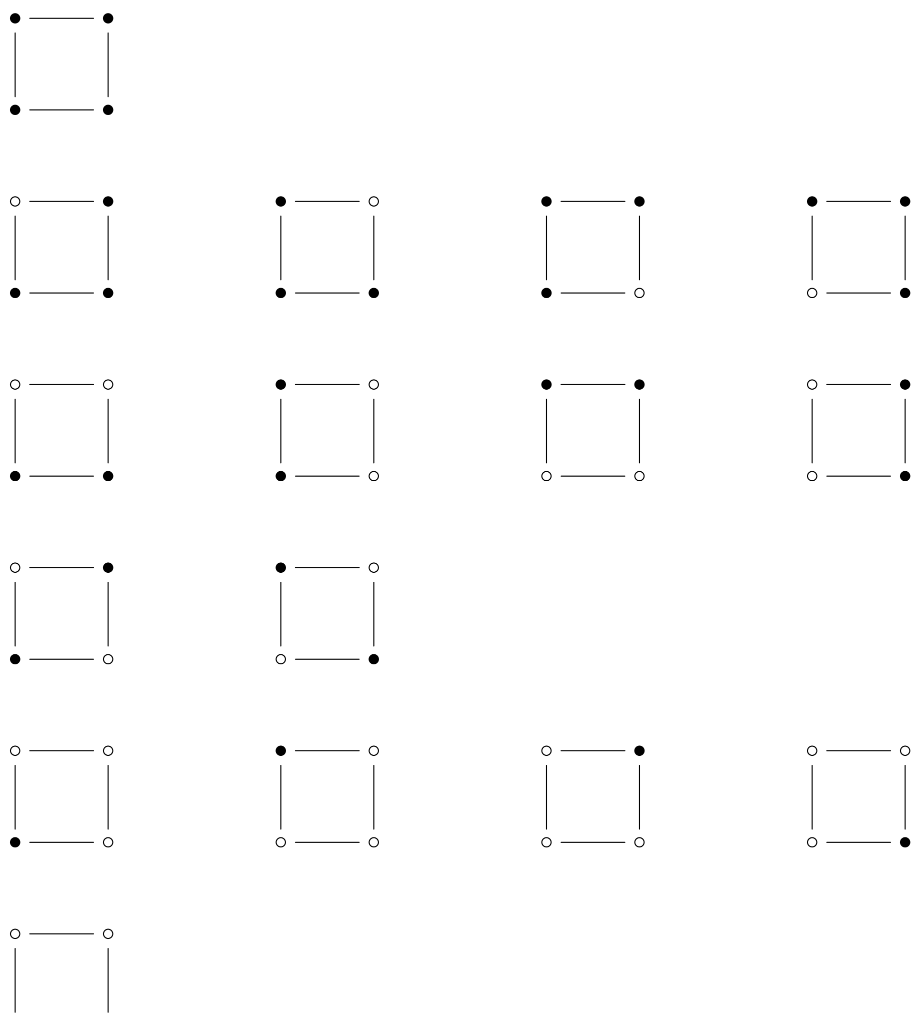

Here we've been careful in our organization of the colorings: Each row consists of the same number of black and white dots, but more importantly, they are arranged in such a way that one one can pass from one coloring to another via some element of D_8, the symmetry group of the square.[1] Let's generalize this.

[1] Indeed, within each row, passing to the right corresponds to the 90° clockwise rotational symmetry.

In addition to fixing X of size n and a permutation group G, let's also fix a set of r "colors," called $f_1, f_2, \ldots, f_r$.[2] Let F denote this set of r colors.

Definition 1. An *r-coloring* of X is a function $\omega : X \to F$. We denote by Ω the set of all r-colorings of X.

Exercise. Show that $|\Omega| = r^n$.

Our goal is to now understand how G can transform one coloring into another, and then count the number of G-orbits of this action. We will do this by first explicitly turning G into a group of symmetries of the set Ω.

Definition 2. For $g \in G$, let $\widehat{g} \in \Sigma_\Omega$ be the permutation of Ω defined by

$$(\widehat{g}(\omega))(x) := \omega(g^{-1}(x)).$$

Let $\widehat{G} \subseteq \Sigma_\Omega$ denote the set of all permutations of Ω that arise in this manner: $\widehat{G} = \{\widehat{g} \mid g \in G\}$.

This definition is a bit convoluted; the key to understanding it is remembering where each of the pieces live. $x \in X$ is a point of the underlying set. ω is a coloring of X, so $\omega(x) \in F$ is a color. $g \in \Sigma_x$ is a permutation of x, so g^{-1} is as well, hence $g^{-1}(x) \in X$ is a point and $\omega(g^{-1}(x)) \in F$ is a color. Finally, $\widehat{g}$, what we're defining, is a permutation of Ω: When fed a coloring ω, it returns another coloring $\widehat{g}(\omega)$ (which is in turn only understood by describing where points of X are sent).

One might reasonably ask why $\widehat{g}$ is defined using g^{-1} instead of g. An example-based explanation would be to note that we've already seen this in the colorings of the square used above: If the vertices of the square are labeled

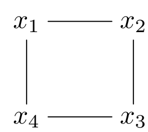

then the clockwise 90° rotation R is given by the cycle $(x_1\ x_2\ x_3\ x_4)$. If we examine the second row above,

the idea is that ω_2 is obtained from ω_1 by applying R to the underlying square, which we express by saying $\omega_2 = \widehat{R}(\omega_1)$. Similarly $\omega_3 = \widehat{R}(\omega_2)$, $\omega_4 = \widehat{R}(\omega_3)$, and $\omega_1 = \widehat{R}(\omega_4)$. What does this actually mean in terms of colorings, though? If we use the symbol b for the color black and w for the color white, we can express all the colorings in the table below:

	x_1	x_2	x_3	x_4
ω_1	w	b	b	b
ω_2	b	w	b	b
ω_3	b	b	w	b
ω_4	b	b	b	w

Now, where does $\omega_2 = \widehat{R}(\omega_1)$ send x_i, in terms of the rotation $R = (x_1\ x_2\ x_3\ x_4)$ and ω_1? Since $\omega_2(x_2) = w$ and the only point colored white by ω_1 is x_1, we see that ω_2 colors by first *undoing* the action of R and the coloring via ω_1. As $R(x_1) = x_2$, we have $x_1 = R^{-1}(x_2)$, so

$$w = \omega_1(x_1) = \omega_1(R^{-1}(x_2)) = \omega_2(x_2),$$

which we in turn generalize to say $\omega_2(x_i) = \omega_1(R^{-1}(x_i))$. From the original observation that $\omega_2 = \widehat{R}(\omega_1)$, we see that we've actually recreated the strange looking definition of $\widehat{g}$ given above.

There is a deeper reason for defining $\widehat{g}(\omega)$ in terms of g^{-1}, which is the content of the following:

[2] By definition, a *color* is just one of these f_i. We use the letter f instead of the more suggestive (unless you're German, say) c for color because we will be talking about cycles shortly.

Proposition 3. *For all $g, h \in G$:*

(a) $\widehat{g} \circ \widehat{h} = \widehat{g \circ h}$.

(b) $(\widehat{g})^{-1} = \widehat{g^{-1}}$.

Proof. It's important to remember *where* the compositions and inversions of this Proposition are taking place: For part (a), in the expression $\widehat{g} \circ \widehat{h}$ the composition is in Σ_Ω, while in $\widehat{g \circ h}$ it is in Σ_X.[3] The content of this result is summarized with the phrase $\widehat{\ }: G \to \widehat{G} : g \mapsto \widehat{g}$ is a *group homomorphism*, i.e., a set-map that respects the multiplicative structure of groups.

(a) We must show $(\widehat{g} \circ \widehat{h})(\omega) = \widehat{g \circ h}(\omega)$ for all $\omega \in \Omega$, which in turn means we need

$$((\widehat{g} \circ \widehat{h})(\omega))(x) \stackrel{?}{=} (\widehat{g \circ h}(\omega))(x)$$

for all $\omega \in \Omega$ and $x \in X$. Using the definition of $\widehat{\ }$, the left hand side becomes

$$((\widehat{g} \circ \widehat{h})(\omega))(x) = (\widehat{g}(\widehat{h}(\omega)))(x) = (\widehat{h}(\omega))(g^{-1}(x)) = \omega(h^{-1}(g^{-1}(x))) = \omega((h^{-1} \circ g^{-1})(x)),$$

while the right is

$$(\widehat{g \circ h}(\omega))(x) = \omega((g \circ h)^{-1}(x)).$$

Since we know for all $g, h \in G$ that $(g \circ h)^{-1} = h^{-1} \circ g^{-1}$, the result follows.

(b) By definition of the inverse, this amounts to showing $\widehat{g} \circ \widehat{g^{-1}} = \mathrm{id}_\Omega$.[4] By part (a), we know

$$\widehat{g} \circ \widehat{g^{-1}} = \widehat{g \circ g^{-1}} = \widehat{\mathrm{id}_X}.$$

Show from the definition of $\widehat{\ }$ that $\widehat{\mathrm{id}_X} = \mathrm{id}_\Omega$.

□

Exercise. If we had defined $\widehat{g}(\omega)$ using g instead of g^{-1}, show that part (a) of the previous Proposition would not hold.[5]

By definition, the map $G \to \widehat{G}$ is a surjection. In fact, it is an injection as well:

Proposition 4. *If $r \geq 2$ and $g, h \in G$ are such that $\widehat{g} = \widehat{h}$, then $g = h$.*[6]

Proof. $\widehat{g} = \widehat{h}$ iff for all $\omega \in \Omega$ we have $\widehat{g}(\omega) = \widehat{h}(\omega)$ iff for all $\omega \in \Omega$ and $x \in X$ we have

$$(\widehat{g}(\omega))(x) = (\widehat{h}(\omega))(x) \Leftrightarrow \omega(g^{-1}(x)) = \omega(h^{-1}(x)).$$

I claim that this means $g^{-1}(x) = h^{-1}(x)$ for all $x \in X$: If not, then for some x_0 such that $g^{-1}(x_0) \neq h^{-1}(x_0)$, there is a coloring ω_0 with $\omega_0(g^{-1}(x_0)) = f_1$ and $\omega_0(h^{-1}(x_0)) = f_2$, contrary to our assumption that $\widehat{g} = \widehat{h}$. But if $g^{-1}(x) = h^{-1}(x)$ for all $x \in X$, we have $g^{-1} = h^{-1} \in \Sigma_X$, and hence $g = h$, as desired. □

In other words, not only does the permutation group G of X induce a permutation group $\widehat{G}$ of Ω, but G and $\widehat{G}$ are essentially the same group.[7]

---

[3]Since $g \in \Sigma_X$ and $\widehat{g} \in \Sigma_\Omega$, this is in fact the only interpretation that makes sense.

[4]Actually, we also need that $\widehat{g^{-1}} \circ \widehat{g} = \mathrm{id}_\Omega$ as well; can you show in general that this second equality follows from the first?

[5]This is actually a reflection of the "left action vs. right action" dichotomy hinted at in the previous notes.

[6]What happens in the degenerate case $r = 1$? Recall that r is the number of colors in our colorings.

[7]We say that the bijective homomorphism $G \to \widehat{G}$ is an *isomorphism* of groups. The existence of an isomorphism means that G and $\widehat{G}$ are two different *representations* of the same *abstract group*, which should be thought of as a platonic ideal of a set of symmetries, divorced from the mundane reality of the object being symmatrized.

Ok, let's step back. The goal originally was to count the number of colorings of X, up to the symmetries induce by the permutation group G. We've now described explicitly how G acts on the set of colorings Ω, and in so doing rephrased the question as calculating the number of orbits of $\widehat{G}$ on Ω: We seek a description of $|\widehat{G}\backslash\Omega|$. Of course, we already have such a numerical description in the Lemma that is not Burnside's:

$$|\widehat{G}\backslash\Omega| = \frac{1}{|\widehat{G}|}\sum_{\widehat{g}\in\widehat{G}}|\Omega^{\widehat{g}}| = \frac{1}{|G|}\sum_{g\in G}|\Omega^{\widehat{g}}|$$

Here, the second expression is just a reformulation of the result that the map $G \to \widehat{G}$ is a bijection, so it does not matter whether we choose to index over G or $\widehat{G}$. The question is now whether we can improve on this expression by, for instance, finding a way to deal with r- and r'-coloring simultaneously.

Let's consider, for fixed $g \in G$, the set $\Omega^{\widehat{g}}$. If ω is a fixed point of $\widehat{g}$, then by definition $\widehat{g}(\omega) = \omega$, or for all $x \in X$

$$\omega(g^{-1}(x)) = \omega(x).$$

If we express g in terms of its disjoint cycles,

$$g = (x_1^1\ x_2^1\ \dots\ x_{\ell_1}^1)(x_1^2\ x_2^2\ \dots\ x_{\ell_2}^2)\dots(x_1^k\ x_2^k\ \dots\ x_{\ell_k}^k),$$

then an easy induction shows that

$$\omega(x_1^i) = \omega(x_2^i) = \dots = \omega(x_{\ell_i}^i)$$

for all $1 \le i \le k$. In other words, $\omega \in \Omega^{\widehat{g}}$ if and only if ω is *constant* on each of g's cycles. As we can choose any of the r colors of F for each cycle of g, we have proved

Lemma 5. *If $g \in G$ has k cycles, then $|\Omega^{\widehat{g}}| = r^k$.*

We will see in the next lecture that it is worthwhile to distinguish between the lengths of the various cycles, so that we are lead to define:

Definition 6. For $g \in G$, the *cycle type* of g is the expression $[1^{c_1}2^{c_2}\dots n^{c_n}]$, where c_i is the number of cycles of length i of g.

Note that for all $g \in G$, we have $\sum_{i=1}^{n} i \cdot c_i = n$. If g has k cycles, then $\sum_{i=1}^{n} c_i = k$.

We wish to manipulate cycle types algebraically, so we introduce the *formal variables* $x_1, x_2, \dots, x_n$, and define $\zeta_g(x_1,\dots,x_n)$ to be the polynomial $x_1^{c_1}x_2^{c_2}\dots x_n^{c_n}$.

As the x_i are formal variables (being manipulated like variables in a polynomial ring, because that's what they are) we can *evaluate* $\zeta_g(x_1,\dots,x_n)$ at an n-tuple of real numbers $(a_1,\dots,a_n) \in \mathbb{R}^n$ by plugging in a_i for x_i. If we evaluate at the constant n-tuple $(r, r, \dots, r)$ (r is the number of colors in F), we have

$$\zeta_g(r, r, \dots, r) = r^{c_1}r^{c_2}\dots r^{c_n} = r^{c_1+c_2+\dots+c_n} = r^k = |\Omega^{\widehat{g}}|.$$

Definition 7. The *cycle index* of G is the polynomial

$$\zeta_G(x_1,\dots,x_n) := \frac{1}{|G|}\sum_{g\in G}\zeta_g(x_1,\dots,x_n).$$

If we evaluate ζ_G at $(r, r, \dots, r)$, we get

$$\zeta_G(r, r, \dots, r) = \frac{1}{|G|}\sum_{g\in G}\zeta_g(r, r, \dots, r) = \frac{1}{|G|}\sum_{g\in G}|\Omega^{\widehat{g}}| = |\widehat{G}\backslash\Omega|,$$

and we have proved

Theorem 8. *The number of G-distinct r-colorings of X is $\zeta_G(r, r, \dots, r)$.*

Example 9. Let's look at the square

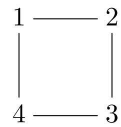

with group of symmetries D_8 generated by $R = (1234)$ and $F = (24)$. In terms of this labeling, the elements of D_8, with their corresponding cycle indices are given by:[8]

g	cycles	ζ_g
id	(1)(2)(3)(4)	x_1^4
R	(1234)	x_4^1
R^2	(13)(24)	x_2^2
R^3	(1432)	x_4^1
F	(1)(24)(3)	$x_1^2 x_2^1$
FR	(14)(23)	x_2^2
FR^2	(13)(2)(4)	$x_1^2 x_2^1$
FR^3	(12)(34)	x_2^2

and we compute

$$\zeta_{D_8} = \frac{1}{8}(x_1^4 + 3x_2^2 + 2x_4 + 2x_1^2 x_2).$$

Evaluating each x_i at 1, we conclude that there is

$$\zeta_{D_8}(1,1,1,1) = \frac{1}{8}(1+3+2+2) = 8$$

distinct ways to color the edges of a square with a single color.[9] If we we want to ask how many distinct 2-colorings are possible, we find:

$$\zeta_{D_8}(2,2,2,2) = \frac{1}{8}(16 + 3\cdot 4 + 2\cdot 2 + 2\cdot 8) = 6,$$

which agrees with our opening example (count the number of rows of colored squares on the first page). What's nice is that we get for free the number of 3-colorings, and 13-colorings, and etc. in the same manner:

$$\begin{aligned} \zeta_{D_8}(3,3,3,3) &= \frac{1}{8}(81 + 3\cdot 9 + 2\cdot 3 + 2\cdot 27) = 21 \\ \zeta_{D_8}(13,13,13,13) &= \frac{1}{8}(13^4 + 3\cdot 13^2 + 2\cdot 13 + 2\cdot 13^3) = 3679 \\ &\vdots \end{aligned}$$

[8]While we normally suppress 1-cycles from our notation, in this context it is important not to forget that they have never left us.

[9]This should not be surprising. Actually, we haven't explicitly proved that this formula works for the case $r = 1$, as Proposition 4 assumed that $r \geq 2$, and this was actually a central technical piece of the argument. Can you give a separate argument to show that $\zeta_G(1, 1, \ldots, 1) = 1$ in general?