Discrete Mathematics lecture notes 12-2

November 25, 2013

32. Dividing with groups.

Let's recall two examples we saw in the previous lecture: First, Γ is the graph

with group of symmetries $C_2 := \{\mathrm{id}, (a\ c)\}$.[1] The name C_2 comes from the fact that this group is called *cyclic*, and is of order 2. The action of C_2 on Γ[2] is not transitive, as there is clearly no symmetry in G that takes the point b to either of the points a or c. The set of C_2-orbits of Γ is $\{[a]_{C_2} = [c]_{C_2} = \{a, c, \}, [b]_{C_2} = \{b\}\}$. In particular, there is more than one orbit.[3]

Second, we have the square with labeled corners Ξ:

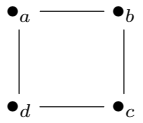

The symmetry group of Ξ is $D_8 = \{\mathrm{id}, R, R^2, R^3, F, FR, FR^2, FR^3\}$, where R is the rotation $(a\ b\ c\ d)$ and F is the flip $(b\ d)$. Ξ is easily seen to be a transitive D_8-set.[4]

Notation 1. If G is a permutation group of the set X, we denote by $G\backslash X$ the set of G-orbits in X.[5]

As the notation suggests, the process of going from a set X to the set of G-orbits of X is a generalization of division: The act of identifying all points of X that lie in the same G-orbit should be thought of as "dividing out the G-action on X."

Let's explore the extent to which our experience with division in $\mathbb{Q}$ gives us good intuition with more general group-division. The first thing we might ask is: Given the sizes of the set X and the group G, we can look at their ration as rational numbers. Does this number agree with the size of the set of G-orbits of X? In other words:

$$\frac{|X|}{|G|} \stackrel{?}{=} \left|G\backslash X\right|$$

In fact, this question is easy to answer: No. In neither of the examples Γ or Ξ does this formula obtain.

Consider the square case first: $|\Xi| = 4$ and $|D_8| = 8$, so the ratio of set size to group size is $\frac{1}{2}$, when we were expecting to get 1 (because the action is transitive). In fact, what would even mean to say that a set has half an element?

Let's defer that question for another class, and instead observe that the size of the stabilizer of any corner happens to be 2. Note that I do not claim that the stabilizers themselves are equal, just that their sizes are. For instance, the stabilizer of a is $\{\mathrm{id}, (b\ d)\}$, while the stabilizer of b is $\{\mathrm{id}, (a\ c)\}$.

---

[1] At least after we identify the symmetries of Γ with the subset of $\Sigma_{\{a,b,c\}}$ that preserves the combinatorial structure of Γ.

[2] This is often expressed with the phrase, "The G-set $\Gamma\ldots$."

[3] We're really just talking about the action of G on the vertices of Γ here, and similarly throughout the rest of this lecture. Note that none of the geometric objects we're discussing could ever be truly transitive under the symmetry group action, as (for example) a vertex can never be sent to a point on the interior of an edge.

[4] Again, the D_8 acts transitively on the *vertices* of Ξ.

[5] The slightly odd notation $G\backslash X$ reflects the fact that we are talking about *left* G-actions, where a symmetry g acts on an element x by $x \mapsto g(x)$.

There is a dual notion of *right* G-actions, where the effect of g on x is written by $x \mapsto (x)g$. The G-orbits of a right action are written X/G.

This may seem like a meaningless distinction, but something odd happens when you compare the effects of two elements of G acting on x, one after the other. For $g, h \in G$ and $x \in X$, associativity of function composition tells us that a left G-action should satisfy $g(h(x)) = (g \circ h)(x)$, while a right G-action should satisfy $((x)g)h = (x)(g \circ h)$. In other words, there's an essential ambiguity about how the expression $g \circ h$ should be interpreted: Does it mean do h first and then g, or g first and then h? If we adopt the former convention, we are naturally led to the notion of a left action; if the latter, to right actions.

Of course, sometimes one wants to consider *both* left and right actions occurring simultaneously....

Let's consider the example of Γ again. Here, not all points have equal stabilizer size, as the stabilizers of a and c are just $\{\mathrm{id}\}$, while that of b is $\{\mathrm{id}, (a\ b)\}$. The observant reader might note that the division happens to coincide with the C_2-orbits. Let's formalize this.

Proposition 2. *Let G be a set of symmetries of the finite set X. If $g \in G$, $x, y \in X$ are such that $g(x) = y$, then $G_y = gG_xg^{-1} := \{g\sigma g^{-1} | \sigma \in G_x\}$.*

Proof. Note in the statement of the Proposition that g, x, and y are all fixed, while σ is free to range over the stabilizer of x.

Consider the function $c_g : G_x \to G_y : \sigma \mapsto g\sigma g^{-1}$. It may not be obvious that the image of c_g actually lies inside G_y, so we need to check that.

If $\sigma \in G_x$, then $\sigma(x) = x$. We want to show that $c_g(\sigma) = g\sigma g^{-1} \in G_y$. Since $g(x) = y$, we have

$$g^{-1}(y) = g^{-1}(g(x)) = (g^{-1} \circ g)(x) = \mathrm{id}(x) = x,$$

or $g^{-1}(y) = x$ for short. Therefore $(g\sigma g^{-1})(y) = g(\sigma(g^{-1}(y))) = g(\sigma(x)) = g(x) = y$, so we are correct in asserting that c_g is a function from G_x to G_y. If we can show that c_g is actually a bijection, we will be done.

To do this we'll show that c_g has an inverse $(c_g)^{-1} : G_y \to G_x$. This inverse is easy to define: Set $c_{g^{-1}} : G_y \to G_x : \tau \mapsto g^{-1}\tau g$. The same argument as in the previous paragraph shows that $c_{g^{-1}}$ is well-defined. We then have to show that $c_{g^{-1}} = (c_g)^{-1}$, i.e., $c_g \circ c_{g^{-1}} = \mathrm{id}_{G_y}$ and $c_{g^{-1}} \circ c_g = \mathrm{id}_{G_x}$. For $\sigma \in G_x$, we have

$$c_{g^{-1}} \circ c_g(\sigma) = c_{g^{-1}}(g\sigma g^{-1}) = g^{-1}(g\sigma g^{-1})g = (g^{-1}g)\sigma(gg^{-1}) = \sigma,$$

so that $c_{g^{-1}}$ is a left inverse to c_g. The other direction is proved in exactly the same way. □

This result is technically important, but for us right now the main point is that it implies the following:

Corollary 3. *If the action of G on X is transitive, then for all $x, y \in X$ we have $|G_x| = |G_y|$.*

Proof. For any $x, y \in X$, because X consists of a single G-orbit there is some $g \in G$ such that $g(x) = y$. Therefore $c_g : G_x \to G_y$ is a bijection by the previous Proposition. □

We should interpret this as saying that sets with transitive permutation group actions are extremely symmetric: Not only is it possible to go from any point to any other via an element of G, but the stabilizers of all points are "the same" in a certain sense.

In fact, even more is true, though we'll need to introduce a new notion to express this fact.

Definition 4. If X is a set with permutation group G, then for any $x, y \in X$ the *G-transporter from x to y* is the set $T_G(x, y) := \{g \in G | g(x) = y\}$.

For example, in the case of the square Ξ with permutation group D_8, it is easy to check that

$$G_a = T_G(a, a) = \{\mathrm{id}, F\}, \qquad T_G(a, b) = \{R, RF\}, \qquad T_G(a, c) = \{R^2, FR^2\}, \qquad T_G(a, d) = \{R^3, R^3F\}.$$

We can represent this as:

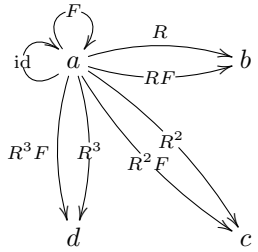

Here, each arrow is a group element, and the action of that group element on the point a is indicated by the terminus of that arrow. There are a couple of things to note: First, for any point $x \in \Xi$, the number of elements in the transporter $T_G(a, x)$ is always 2. Second, given one element in $g \in T_G(a, x)$, we can actually represent the entire transporter as $g \cdot G_a$; i.e., translating the stabilizer of a by g on the left. These observations are generally true.

Proposition 5. *If X is a set with permutation group G and $x, y \in X$, $g \in G$ are such that $g(x) = y$, then $T_G(x, y) = g \cdot G_x := \{g \cdot \sigma | \sigma \in G_x\}$.*

Proof. We define a map $\ell_g : G_x \to T_G(x, y) : \sigma \mapsto g \cdot \sigma$. Let's check that this map is well-defined: If $\sigma \in G_x$, then $\sigma(x) = x$. Therefore $(g \circ \sigma)(x) = g(\sigma(x)) = g(x) = y$, and $g \circ \sigma \in T_G(x, y)$. I claim ℓ_g is a bijection.

We also have a map $\ell_{g^{-1}} : T_G(x, y) \mapsto G_x : \tau \mapsto g^{-1}\tau$. Check yourself that $\ell_{g^{-1}}$ actually takes an element of $T_G(x, y)$ to an element of G_x, and that $\ell_{g^{-1}} = (\ell_g)^{-1}$. □

Note that we need the element g such that $g(x) = y$ in order to make this argument work. Considering the example Γ above, convince yourself that it is possible that in a *nontransitive* situation we can have points $x, y \in X$ such that $T_G(x, y) = \emptyset$.

Again, we're really interested in the numerical implications of this proposition in the transitive case.

Corollary 6. *For X a set with a transitive permutation group G and any $x, y \in X$, we have*

$$|G_x| = |T_G(x, y)| = |G_y|.$$

This is just the formalization of the statement that for a transitive G-set X, not only does every point look the same through the eyes of G, but the action of G on every point also looks the same.

Ok, we're now ready to find a formula for counting the orbits of a transitive G-set.[6]

Theorem 7. *If X is a finite set with transitive permutation group G, then for any $x \in X$,*

$$|G| = |X| \cdot |G_x|.$$

Proof. Define $E := \{(g, y) | g(x) = y\} \subseteq G \times X$. Note that x is fixed in the definition of E, but g and y are both free to vary. We will count the size of E in two different ways, which will lead to the result.

The first way of counting $|E|$ is to fix an element $g \in G$ and ask how many pairs (g, y) lie in E for that particular g, then sum up those answers as g is free to vary. Thus, we define $R_g := |\{y \in X | g(x) = y\}|$. Since g is fixed here, there is exactly one y such that $g(x) = y$ (as g is, at its heart, a bijection $g : X \xrightarrow{\sim} X$), so that $R_g = 1$ for all $g \in G$. As $|E| = \sum_{g \in G} R_g = \sum_{g \in G} 1 = |G|$, we've calculate $|E|$ in a manner that gives the left hand side of the Theorem's conclusion.

The second way of counting $|E|$ is to fix $y \in X$, as how many g correspond to that fixed y, and sum the result as y ranges of X. Set $C_y := |\{g \in G | g(x) = y\}|$. By definition of the transporter, we actually have $C_y = |T_G(x, y)|$, and by the previous Corollary, this is always $|G_x|$, independent of y. Thus $|E| = \sum_{y \in X} C_y = \sum_{y \in X} |G_x| = |X| \cdot |G_x|$. The result follows. □

Going back to our original question of calculating the number of G-orbits, we might restate this result as: For X a transitive G-set and $x \in X$, we have

$$\frac{|X|}{|G|} \cdot |G_x| = 1 = |G \backslash X|$$

In other words, for the transitive case our initial guess for the number of orbits as the ratio of $|X|$ to $|G|$ wasn't that far off; we just had to account for the possibility that a nonidentity group element could fix a point. Thus, to understand $|G\backslash X|$, we need to know $|X|$, $|G|$, and $|G_x|$, at least in the transitive case.

Of course, the transitive case isn't really what we're interested in, as we know $|G\backslash X| = 1$ by definition here. However, we can use this case as a stepping stone to find the answer in general. The problem is that, if X is not transitive, and x, y live in different G-orbits, there is no reason to assume $|G_x| = |G_y|$; cf. the example Γ, where $|G_a| = 1$ and $|G_b| = 2$.

This suggests we need a different way of storing data that doesn't require *a priori* knowledge of the G-orbits of X. The idea is actually hidden within the proof of the previous Theorem:

[6] Exercise: Why does this claim, on its face, seem like a ridiculously circuitous goal? After you've read the rest of the notes, can you explain the point in your own words?

Definition 8. Let X be a set with permutation group G. For $g \in G$, the *fixed point set of* g is

$$X^g := \{x \in X | g(x) = x\}.$$

Note that the condition "$g(x) = x$" is the same that occurs in the definition of the stabilizer of x, G_x. The only difference is whether we're imagining g to be fixed and letting x vary, or *vice versa*. In other words, we should be able to go from statements concerning sizes of stabilizers to statements concerning sizes of fixed point sets by using the trick central to the proof of the last Theorem. Indeed:

Proposition 9. *If X is a set with transitive permutation group G, then*

$$|G| = \sum_{g \in G} |X^g|.$$

Proof. Set $E := \{(g, x) | g(x) = x\} \subseteq G \times X$. Let's count $|E|$ twice and see what we get.

For each $g \in G$, set $R_g := |\{x \in X | (g, x) \in E\}| = |\{x \in X | g(x) = x\}| = |X^g|$. Thus

$$|E| = \sum_{g \in G} R_g = \sum_{g \in G} |X^g|,$$

giving the right hand side of the desired equation.

On the other hand, for each $x \in X$, set $C_x := |\{g \in G | (g, x) \in E\}| = |\{g \in G | g(x) = x\}| = |G_x|$. As X is transitive, $|G_x|$ is constant as a function of x, so

$$|E| = \sum_{x \in X} C_x = \sum_{x \in X} |G_x| = |X| \cdot |G_{x_0}|$$

for $x_0 \in X$ any fixed element. However, but the previous Theorem, $|X| \cdot |G_{x_0}| = |G|$, and comparing the two methods of calculating $|E|$ yields the result. □

Again, we can rephrase this result to say that, for X transitive,

$$|G \backslash X| = 1 = \frac{1}{|G|} \sum_{g \in G} |X^g|.$$

The utility of this change in perspective comes from the passage from the transitive to the nontransitive case:

Lemma 10. *If X is a finite set with permutation group G, and we can decompose X into G-orbits $X_1, X_2, \ldots, X_n$, then for all $g \in G$,*

$$|X^g| = \sum_{i=1}^{n} |X_i^g|.$$

Proof. Exercise. □

In particular the numerics $|X^g|$ give us the sums of the g-fixed points on all orbits, without explicitly knowing how to decompose X into orbits. Putting all of this together, we finally arrive at the punchline:

Theorem 11 (The Lemma that is not Burnside's)**.** *If X is any finite set with permutation group G, then*

$$|G\backslash X| = \frac{1}{|G|}\sum_{g\in G}|X^g|.$$

Proof. We only have one technique—let's hope it doesn't fail us now. Set $E := \{(g,x)|g(x)=x\} \subseteq G\times X$.

If we compute $|E|$ by fixing g first, we define $R_g := |\{x\in X|g(x)=x\}| = |X^g|$, and we compute

$$|E| = \sum_{g\in G} R_g = \sum_{g\in G}|X^g|.$$

On the other hand, if we compute $|E|$ by fixing x first, we set $C_x := |\{g\in G|g(x)=x\}| = |G_x|$. If X has n G-orbits, called $X_1, X_2, \ldots, X_n$, with chosen elements $x_i \in X_i$, we have

$$|E| = \sum_{x\in X} C_x = \sum_{i=1}^{n}\sum_{x\in X_i} C_x = \sum_{i=1}^{n}\sum_{x\in X_i}|G_x| = \sum_{i=1}^{n}\sum_{x\in X_i}|G_{x_i}| = \sum_{i=1}^{n}|G| = n\cdot|G|.$$

Here, the second equality is just grouping the terms C_x according to G-orbit, the fourth is using that fact that within the transitive G-set X_i, the order of the stabilizer is constant, and the fifth is Theorem 7. As $n = |G\backslash X|$, we can put together the two ways of counting $|E|$ to conclude

$$\sum_{g\in G}|X^g| = |G\backslash X|\cdot|G|,$$

from which the Theorem follows. □