

Enumerative Combinatorics 7: Group actions

Peter J. Cameron

Autumn 2013

How many ways can you colour the faces of a cube with three colours? Clearly the answer is $3^6 = 729$. But what if we regard two colourings as the same if one can be transformed into the other by a rotation of the cube? This is typical of the problems we consider in this chapter.

7.1 The Orbit-Counting Lemma

This chapter of the lectures, unlike most of the others, requires some technical background. I assume that you know the definition of a group. I will run briefly through the theory of group actions, and finally reach the Orbit-Counting Lemma, which solves our introductory problem.

Throughout this section, permutations act “on the right”, that is, the effect of applying a permutation π to an element x of the domain is written $x\pi$. This is not just a matter of notation; it entails the fact that the product $\pi_1\pi_2$ of two permutations is calculated by the rule “first π_1 , then π_2 ”, rather than the other way round. This ensures that $x(\pi_1\pi_2) = (x\pi_1)\pi_2$ for all elements x .

An *action* of a group G on a set X is a map associating to each group element $g \in G$ a permutation π_g of X in such a way that the following two conditions hold:

- (a) $\pi_{gh} = \pi_g\pi_h$ for all $g, h \in G$ (that is, $x\pi_{gh} = x\pi_g\pi_h$ for all $g, h \in G$ and all $x \in X$);
- (b) if 1 denotes the identity element of G , then π_1 is the identity permutation (that is, $x\pi_1 = x$ for all $x \in X$).

Usually we simplify notation by not distinguishing between g and π_g , writing simply xg instead of $x\pi_g$. From a different point of view, an action is a homomorphism from the group G to the symmetric group of all permutations of X .

Two elements $x, y \in X$ are equivalent under the action if there exists an element $g \in G$ such that $xg = y$. It is routine to show that this is really an equivalence relation; its equivalence classes are called *orbits*, and the action is *transitive* if there is just one orbit. Thus we have a first structure theorem: any action can be split uniquely into transitive actions on the sets of the orbit partition of the domain.

In our motivating problem, the group G of 24 rotations of the cube acts on the set X of 729 coloured cubes, and we want to count the orbits. So our immediate goal is to count the orbits in an arbitrary action.

If H is a subgroup of G , then there is a natural partition of G into *right cosets* Hx of H , for $x \in G$. Lagrange's Theorem assures us that each coset has the same cardinality, so the number of cosets is equal to $|G|/|H|$. We denote the set of right cosets of H in G by $\text{cos}(H, G)$. Now there is an action of G on the set $\text{cos}(H, G)$: the group element g induces the permutation $Hx \mapsto H(xg)$. At risk of some confusion, we write this as $(Hx)g = H(xg)$.

Now, given any transitive action of G on a set X , and $x \in X$, the set

$$\{g \in G : xg = x\}$$

is a subgroup of H , called the *stabiliser* of x , and denoted by $\text{Stab}_G(x)$. Now there is a natural bijection between X and $\text{cos}(H, G)$, where the element $y \in X$ corresponds to the set

$$\{g \in G : xg = y\}$$

(it is easily checked that this is a coset of H). This bijection also respects the action of G : if $z \in G$ satisfies $yg = z$, and Hk and Hl are the cosets corresponding to y and z , then $(Hk)g = (Hl)$.

So the so-called "coset spaces" of subgroups of G give a complete list of transitive actions of G , up to a natural notion of isomorphism of actions.

Note in addition that any two points in the same orbit have stabilisers of the same order. (The stabilisers are in fact conjugate subgroups of G .)

In an arbitrary action of G on X , we let $\text{fix}_X(g)$ denote the number of points of X which are fixed by the permutation g . Now we can state the *Orbit-Counting Lemma*, the foundation of enumeration under group action.

Theorem 7.1 *Let G act on the finite set X . Then the number of G -orbits in X is equal to the average number of fixed points of elements of G , that is,*

$$\frac{1}{|G|} \sum_{g \in G} \text{fix}_X(g).$$

The theorem has a probabilistic interpretation. Choose a random element of G (from the uniform distribution). Then its expected number of fixed points is equal to the number of orbits of G .

Proof Construct a bipartite graph as follows. The vertices are of two types: the elements of X , and the elements of G . There is an edge from x to g if $xg = x$. We count the number of edges in two different ways.

Each vertex g lies in $\text{fix}_X(g)$ edges; so the number of edges is $\sum_{g \in G} \text{fix}_X(g)$.

Now we count the other way. Take a point $x \in X$. The number of edges containing it is $|\text{Stab}_G(x)|$. This value is the same for all the points in the orbit $O_G(x)$ containing x . So the number of edges containing points in the orbit is $|\text{Stab}_G(x)| \cdot |O_G(x)| = |G|$. Since each orbit contributes $|G|$ edges, the number of orbits is obtained by dividing the number of edges by $|G|$, as claimed.

Now consider the coloured cubes. In order to do the calculations, we need to classify the elements of the group G of rotations of the cube (a group of order 24). They are of the following types:

- (a) the identity;
- (b) “face rotations” (about an axis through two opposite face centres) through $\pm\pi/2$ (six of these, two for each pair of opposite faces);
- (c) “face rotations” through π (three of these);
- (d) “edge rotations” (about an axis through two opposite edge midpoints) through π (six of these);
- (e) “vertex rotations” (about an axis through two opposite vertices) through $\pm 2\pi/3$ (eight of these, two for each pair of opposite vertices).

For each type of rotation, we have to count the number of coloured cubes it fixes. A cube will be fixed if faces in the same cycle of the permutation have the same colour. So the answer will be 3^c , where c is the number of cycles

of the permutation on faces. For the five types listed above the numbers of cycles are 6 (each single face is a cycle), 3 (for the vertical axis, the top and bottom faces, and the other four in a single cycle), 4 (as the previous except that the 4-cycle splits into two 2-cycles), 3 (the faces are permuted in cycles of two), and 2 (the faces are permuted in cycles of three). So the calculation of the theorem is:

$$\frac{1}{24}(3^6 + 6 \cdot 3^3 + 3 \cdot 3^4 + 6 \cdot 3^3 + 8 \cdot 3^2) = 57.$$

7.2 Labelled and unlabelled

Many combinatorial objects that we want to count are based on an underlying set, which we usually assume to be the set $\{1, 2, \dots, n\}$. Very often the simplest method of counting gives us the total number of objects that can be built on this set. But we may be completely uninterested in the labels $1, 2, \dots, n$, and want to count two objects as being the same if there are some labellings of the underlying set that make them identical.

We distinguish these two problems as counting *labelled* and *unlabelled* objects.

Counting unlabelled objects is thus an orbit-counting problem: we want to know the number of orbits of the symmetric group S_n , acting on the objects in question by permuting the labels.

To take an extreme case: there are $\binom{n}{k}$ labelled k -element subsets of an n -element set, but there is only one unlabelled subset. Here are a few more examples.

Objects	Labelled	Unlabelled
Subsets	2^n	$n + 1$
Partitions	$B(n)$	$p(n)$
Permutations	$n!$	$p(n)$
Linear orders	$n!$	1

Here $B(n)$ is the Bell number (the number of partitions of an n -set) and $p(n)$ the number of partitions of the number n . Note that the numbers of unlabelled structures can agree and those of labelled structures disagree, or *vice versa*.

The third entry needs a little explanation. Any permutation can be written as a product of disjoint cycles; the cycle lengths form a partition of n

called the *cycle structure* of the permutation. Now given two permutations with the same cycle structure, we can replace the entries in one by those in the other. For example, $(1)(2, 3)$ can be transformed into $(2)(1, 3)$ by swapping the labels 1 and 2. (You might recognise this as the argument that shows that two permutations are conjugate in the symmetric group if and only if they have the same cycle structure.)

In the three cases in the table, we can count the unlabelled objects directly; but in more complicated cases, the Orbit-Counting Lemma is required. One example is the number of graphs on n vertices. The labelled number is $2^{n(n-1)/2}$, since for each of the $n(n-1)/2$ pairs of vertices we can choose whether to join it by an edge or not; but the only way to calculate the number of unlabelled graphs is via the Orbit-Counting Lemma.

7.3 Cycle index

There is a way to “mechanise” the counting in many important cases, which we now discuss. This was introduced by Redfield and, independently, by Pólya, and refined by de Bruijn and others. (Incidentally, these early workers found the Orbit-Counting Lemma in Burnside’s group theory book, and called it “Burnside’s Lemma”, a name which is still sometimes used. However, the result is due to Frobenius, and earlier to Cauchy in a special case.)

The set-up is as follows. We have a set X on which a group G acts. We are going to decorate X by placing one of a set of “figures” at each point. Each figure has a weight, which is a non-negative integer. We don’t require the number of figures to be finite, but we ask that there should be only finitely many figures of any given weight. The figures can thus be counted by the *figure-counting series*

$$A(x) = \sum_{n \geq 0} a_n x^n,$$

where a_n is the number of figures of weight n .

Now one of the configurations we want to count consists of the set X with a figure at each point; this can be described by a function from X to the set of figures. Such a function f will have a weight, given by $w(f) = \sum\{w(x) : x \in X\}$. There are only finitely many functions of any given weight, and the action of the group G preserves weight; so we can let b_n be the number of

functions of weight n , and define the *function-counting series*

$$B(x) = \sum_{n \geq 0} b_n x^n.$$

The final ingredient is the *cycle index polynomial* $Z(G)$, defined as

$$Z(G) = \frac{1}{|G|} \sum_{g \in G} s_1^{c_1(g)} s_2^{c_2(g)} \dots s_n^{c_n(g)}.$$

Here s_1, \dots, s_n are indeterminates, and $c_i(g)$ is the number of cycles of length i in the cycle decomposition of g , for $i = 1, \dots, n$.

Now the *Cycle Index Theorem* states:

Theorem 7.2

$$B(x) = Z(G; s_i \leftarrow A(x^i) \text{ for } i = 1, \dots, n).$$

The notation on the right means that we substitute $A(x^i)$ for s_i , for $i = 1, \dots, n$.

I won't prove the theorem here – it follows from the Orbit-Counting Lemma with a certain amount of ingenuity – but will conclude with a simple application which doesn't even hint at the uses of the theorem.

First, let us calculate the cycle index of the rotation group of the cube. The five types of elements mentioned earlier have the following cycle structures in their action on faces:

- (a) Identity: $(1, 1, 1, 1, 1, 1)$ (usually abbreviated to 1^6).
- (b) Face rotations through $\pm\pi/2$: $1^2 4$.
- (c) Face rotations through π : $1^2 2^2$.
- (d) Edge rotations: 2^3 .
- (e) Vertex rotations: 3^2 .

So the cycle index is

$$Z(G) = \frac{1}{24}(s_1^6 + 6s_1^2 s_4 + 3s_1^2 s_2^2 + 6s_2^3 + 8s_3^2).$$

Now any counting problem for which we can write a figure-counting series can be solved by substitution. For example:

- (a) Take each of the three colours to be a figure of weight 0. The figure-counting series is simply 3. We recover our earlier count.
- (b) Take one of the colours (say red) to have weight 1, and all the others weight 0. The figure-counting series is $x + 2$. So substituting $x^i + 2$ for s_i gives a polynomial in which the coefficient of x^k is the number of types of cube which have exactly k red faces.
- (c) A small extension of the Cycle Index Theorem shows that, if we substitute $p_i(x, y, z) = x^i + y^i + z^i$ for s_i , we obtain a trivariate polynomial in which the coefficient of $x^i y^j z^k$ is the number of cubes with i red, j blue, and k green faces.
- (d) The generalisation to an arbitrary number of colours is now routine.

Exercises

- 1** Perform the calculations in the four counting problems above.
- 2** A necklace has ten beads, each of which is either black or white, arranged on a loop of string. A cyclic permutation of the beads counts as the same necklace. How many necklaces are there?
How many are there if the necklace obtained by turning over the given one is regarded as the same?
- 3** Let G be a permutation group on a set X , where $|X| = n$.
For $0 \leq i \leq n$, let p_i be the proportion of elements of G which have exactly i fixed points on X , and let $p(x) = \sum p_i x^i$ be the generating function for these numbers (the *probability generating function for fixed points*).
For $0 \leq i \leq n$, let F_i be the number of orbits of G in its action on the set of i -tuples of distinct elements of X , and let

$$F(x) = \sum \frac{F_i x^i}{i!}$$

be the exponential generating function for these numbers.

Use the Orbit-counting Lemma to show that

$$F(x) = P(x + 1)$$

and deduce that the proportion of fixed-point-free elements in G is $p_0 = F(-1)$.

Taking G to be the symmetric group S_n , show that the number of fixed-point-free permutations (the *derangement number*) is

$$n! \sum_{k=0}^n \frac{(-1)^k}{k!}.$$

Deduce that this number is the closest integer to $n!/e$.

4 Consider the set of all functions from $\{1, \dots, n\}$ to $\{1, \dots, m\}$. There are m^n functions in the set. Now let the symmetric group S_n act on these functions by permuting their arguments: $(f\pi)(x) = f(x\pi^{-1})$. [Incidentally, the inverse is there to make this an action – can you see why?]

Show that orbits correspond to m -tuples of non-negative integers with sum n , so that the number of orbits is $\binom{m+n-1}{n}$. (See the Appendix in Lecture Notes 7.)

Show that a permutation g with k cycles fixes m^k functions. Hence use the Orbit-Counting Lemma to show that

$$\frac{1}{n!} \sum_{k=1}^n u(n, k) m^k = \binom{m+n-1}{n}.$$

Show that we can replace m by an indeterminate x and multiply by $n!$ to get the identity

$$\sum_{k=1}^n u(n, k) x^k = x(x+1) \cdots (x+n-1),$$

from which some sign changes yield

$$\sum_{k=1}^n s(n, k) x^k = x(x-1) \cdots (x-n+1),$$

a formula we met in Lecture Notes 1. (Here $s(n, k)$ and $u(n, k)$ are the signed and unsigned Stirling numbers of the first kind.)