Permutation groups

> Whatever you have to do with a structure-endowed entity Σ try to determine its group of automorphisms ... You can expect to gain a deep insight into the constitution of Σ in this way.
>
> Hermann Weyl, *Symmetry*

1 Automorphism groups, permutation groups, abstract groups

Let Σ be a mathematical object or structure of any kind whatever, based on a set Ω of "points". The object Σ may be combinatorial (a graph, a design, etc.,), algebraic (a group a field, a vector space, etc.), topological, or indeed anything at all.

An *automorphism* of Σ is an isomorphism from Σ to itself; that is, a bijective map g from Ω to itself such that g and g^{-1} preserve the structure of Σ.

A bijection from Ω to itself is a *permutation* of Ω. We write permutations on the right, so that αg is the image of α under the permutation g.

The set of all automorphisms of Σ is called the *automorphism group* of Σ, denoted by $\mathrm{Aut}(\Sigma)$. It has the following easily checked properties (where G is written for $\mathrm{Aut}(\Sigma)$):

(PG1) for $g_1, g_2 \in G$, the composition $g_1 \circ g_2$ (obtained by applying first g_1, then g_2) is in G;

(PG2) the identity map (which fixes every point) belongs to G;

(PG3) if $g \in G$, then the inverse function g^{-1} belongs to G.

An arbitrary set G of permutations of Ω satisfying (PG1)–(PG3) is called a *permutation group* on Ω. Thus, the automorphism group of any object is a permutation group. Conversely, given any permutation group, it is possible to concoct an object of which it is the automorphism group; but there may be no "natural" object (e.g. graph) with this property.

According to (PG1), if G is a permutation group, then composition is a binary operation on G. This operation has the properties

(AG1) The *associative law* holds:

$$(g_1 \circ g_2) \circ g_3 = g_1 \circ (g_2 \circ g_3)$$

for all $g_1, g_2, g_3 \in G$.

(AG2) There is an element $1 \in G$ such that $g \circ 1 = 1 \circ g = g$ for all $g \in G$.

(AG3) For any $g \in G$, there is an element $g^{-1} \in G$ such that $g \circ g^{-1} = g^{-1} \circ g = 1$, where 1 is the element in (AG2).

The elements in (AG2) and (AG3) are those guaranteed by (PG2) and (PG3), while composition of maps is always associative.

An arbitrary set G with a binary operation $\circ$ satisfying (AG1)–(AG3) is called an *(abstract) group*. Thus, every permutation group is an abstract group.

Two groups G and G' are *isomorphic* if there is a bijection $\phi : G_1 \to G_2$ satisfying $(g_1 \circ g_2)\phi = (g_1\phi) \circ (g_2\phi)$. The map ϕ is an *isomorphism* from G to H. Isomorphic groups are indistinguishable with regard to their algebraic properties. *Cayley's Theorem* asserts that the converse of our earlier remark holds: every abstract group is isomorphic to a permutation group.

Thus our three types of group are in a sense all the same.

We normally suppress the group operation from the notation and write g_1g_2 instead of $g_1 \circ g_2$. The associative law allows us to write $g_1g_2g_3$ unambiguously.

2 Group actions

A *subgroup* of a group G is a subset H of G such that H equipped with the restriction of the operation of G is itself a group.

The *symmetric group* $\mathrm{Sym}(\Omega)$ on a set Ω is the set of all permutations of Ω. Obviously it is a group (with the operation of composition), and a permutation group on Ω) is precisely a subgroup of $\mathrm{Sym}(\Omega)$.

A *homomorphism* from a group G to a group H is a function $\phi : G \to H$ satisfying $(g_1g_2)\phi = (g_1\phi)(g_2\phi)$ for all $g_1, g_2 \in G$. Thus an isomorphism is a bijective homomorphism. The image of a homomorphism ϕ is a subgroup of H, while its *kernel*, the set $\{g \in G : g\phi = 1_H\}$ is a normal subgroup of G.

An *action* of a group G on the set Ω is a homomorphism from G to $\mathrm{Sym}(\Omega)$. We usually suppress the name of the homomorphism and write αg for the image of α under the permutation $g\phi$, where ϕ is the action.

The image of an action is a permutation group. The concept of an action is more flexible than that of permutation groups for two reasons: a given group may have many actions; and the actions may not be faithful (may have non-trivial kernel). For example, the automorphism group of a design acts on the points, point pairs, blocks, flags, etc., of the design.

However, most of the concepts we discuss below apply as well to group actions as to permutation groups.

3 Orbits, transitivity, stabiliser

Suppose that G acts on Ω. Define a relation $\sim$ on Ω by the rule that $\alpha \sim \beta$ if $\alpha g = \beta$ for some $g \in G$. Then $\sim$ is an equivalence relation: the reflexive, symmetric and transitive laws follow immediately from (PG2), (PG3) and (PG1) respectively. Thus Ω decomposes as a disjoint union of equivalence classes. These classes are called *orbits*, and we say that G acts *transitively* on Ω if there is only one orbit.

The *stabiliser* G_α of the point α is defined to be the subset

$$G_\alpha = \{g \in G : \alpha g = \alpha\}$$

of G. It is straightforward to show that it is a subgroup of G, and that for any β, the set

$$\{g \in G : \alpha g = \beta\}$$

is either empty (if α and β lie in different orbits) or a right coset of G_α. Thus, the orbit containing α is in one-to-one correspondence with the set of right cosets of G_α. Moreover, this correspondence respects the action of G (by right multiplication on the set of right cosets). So the action can be described “internally”, within G.

A consequence of this is a version of *Lagrange’s Theorem*: the number $|G|/|G_\alpha|$ of cosets of G_α is equal to the size of the orbit containing α.

4 Multiple transitivity

Let G act on Ω, and let t be a positive integer not greater than Ω. We say that G is *t-transitive* if, given any two t-tuples $(\alpha_1, \ldots, \alpha_t)$ and $(\beta_1, \ldots, \beta_t)$ of distinct elements of Ω, there is an element g of G carrying the first to the second (that is, $\alpha_i g = \beta_i$ for $i = 1, \ldots, t$).

A t-transitive group automatically gives us various t-designs: take Ω to be the point set, and as blocks the images under G of some base block, an arbitrary k-set B.

It is easy to see that a t-transitive group is $(t-1)$-transitive; so any transitive group has a degree of transitivity, the maximum t for which it is t-transitive. The symmetric group of degree n is n-transitive, and the alternating group is $(n-2)$-transitive but not $(n-1)$-transitive. It is known that the largest degree of transitivity of a finite permutation group other than a symmetric or alternating group is 5; but the proof of this requires the Classification of Finite Simple Groups (CFSG), and so is far from elementary. Moreover, using CFSG, all the 2-transitive groups have been determined. In particular, the only 5-transitive groups apart from symmetric and alternating groups are the *Mathieu groups* M_{12} and M_{24}. These are the automorphism groups of 5-$(12,6,1)$ and 5-$(24,8,1)$ designs respectively.

5 Primitivity

Let G act transitively on Ω. A *G-congruence* on Ω is a partition of Ω which is preserved by all elements of G. There are always two trivial congruences: the partition into singletons, and the partition with a single part. The action is called *primitive* if these are the only G-congruences.

A transitive action of G is primitive if and only if the stabiliser G_α is a maximal proper subgroup of G. It is easy to see that any 2-transitive action is primitive. A non-trivial normal subgroup of a primitive group is transitive.

Indeed, more can be said about normal subgroups of primitive groups. Such a group has at most two minimal normal subgroups; if there are two, they are isomorphic. Furthermore, the *O'Nan–Scott Theorem* classifies primitive groups into several types according to the action of the *socle* (the product of the minimal normal subgroups). This result allows CFSG to be applied to many problems about primitive groups, although we are far from a complete classification of them similar to that of the 2-transitive groups. This is discussed further in another topic essay in this series.

6 Permutation character and orbit-counting

The *permutation character* π associated with a permutation group G or an action of G is the function defined by

$$\pi(g) = |\{\alpha \in \Omega : \alpha g = \alpha\}|$$

giving the number of fixed points of the elements of G.

It is a character of G (that is, the trace of a matrix representations), and so can be decomposed as a linear combination of irreducible characters with non-negative integer coefficients. One irreducible which always occurs is the *principal character*, the function taking the value 1 on all elements.

The *Orbit-counting Lemma* (mis-titled Burnside's Lemma) asserts that the number of orbits of G is equal to the average number of fixed points of its elements:

$$\#\, G\text{-orbits} = \frac{1}{|G|} \sum_{g \in G} \pi(g).$$

This is the basis of the theory of enumeration under group action, as we shall see. It has other implications too. For example, it implies that the number of orbits of G is equal to the multiplicity of the principal character in π. Also, G is 2-transitive if and only if π is the sum of just two irreducible characters (one of which is the principal character).

7 Cycle index

More advanced applications of the Orbit-Counting Lemma depend on the *cycle index*, the normalised generating function for the cycle structure of elements of G. More precisely, if $|\Omega| = n$, the cycle index $Z(G)$ is the polynomial in indeterminates $s_1, \ldots, s_n$ given by

$$Z(G) = \frac{1}{|G|} \sum_{g \in G} s_1^{c_1(g)} \cdots s_n^{c_n(g)},$$

where $c_k(g)$ is the number of cycles of length k in the decomposition of g into disjoint cycles.

This polynomial solves many problems involving counting configurations on Ω up to the action of G. We give a general result known as the *Cycle Index Theorem*.

Suppose that F is a set of “figures”, each with a non-negative integer weight. We allow infinitely many figures but require that the number a_n of figures of weight n is finite for all n. The *figure-counting series* is the generating function

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

A configuration, obtained by placing a figure at each point of Ω, can be described as a function from Ω to F. Any such function ϕ has a weight, the sum of the weights of $\phi(\alpha)$ for $\alpha \in \Omega$. The weight is invariant under the action of G on the set of functions given by

$$(\phi g)(\alpha) = \phi(\alpha g^{-1}).$$

Let b_n be the number of G-orbits on functions of weight n, and define the *function-counting series* to be the generating function

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

Now an application of the Orbit-counting Lemma gives the Cycle Index Theorem, which asserts that

$$B(x) = Z(G; s_i \leftarrow A(x^i)),$$

where $F(s_i \leftarrow t_i)$ denotes the result of substituting t_i for s_i in F for $i = 1, \ldots, n$.

For a simple example, if we want to count orbits on colourings of Ω with k colours, we take each colour to be a figure of weight zero, that is, $A(x) = k$. If we are interested in the number of times that a distinguished colour occurs, let this colour have weight 1, and take $A(x) = x + k - 1$.

In particular, the choice $A(x) = x + 1$ enumerates orbits on subsets of Ω by cardinality of the subset. Thus, for example, we can count the t-designs admitting a given t-transitive group of automorphisms, and in particular those for which the group acts block-transitively. However, calculating the sizes of the orbits (and hence the parameters of the designs) requires further analysis, as we now describe.

8 Möbius function

Let P be a partially ordered set. The *zeta-function* of P is the function from $P \times P$ to $\mathbb{Z}$ given by

$$\zeta(x,y) = \begin{cases} 1 & \text{if } x \leq y, \\ 0 & \text{otherwise} \end{cases}.$$

We can regard this as a matrix of size $|P|$; for a suitable ordering of P (extending the partial order), it is upper-triangular with diagonal entries 1, and so is invertible. Its inverse is the *Möbius function* of P; it is also an integral upper triangular matrix with diagonal entries 1.

See the topic essay on partially ordered sets for more detail, and especially for a discussion of Möbius inversion.

Now suppose that G is a permutation group on Ω. We suppose that we know the Möbius function of the partially ordered set $L(G)$ of subgroups of G, and have some information about how these subgroups act.

If we know the orbit lengths of a subgroup H, then we can compute the number $a_k(H)$ of k-sets fixed by H for all k. The generating function is

$$\sum_{k\geq 0} a_k x^k = \prod_{i\in I}(1+x^{m_i}),$$

where I indexes the orbits of H, and m_i is the size of the ith orbit.

Now the number b_k of k-sets whose stabiliser is precisely H is given by Möbius inversion: we have

$$a_k(H) = \sum_{H\leq K\leq G} b_k(K),$$

since a set is fixed by H if and only if its stabiliser K satisfies $H \leq K \leq G$; and so

$$b_k(H) = \sum_{H\leq K\leq G} \mu(H,K)a_k(K).$$

Now a set B with stabiliser H lies in an orbit of size $|G|/|H|$ by Lagrange's Theorem. If G is t-transitive, then the sets in this orbit form a block-transitive t-design on Ω whose parameters can now be calculated.

By taking unions of orbits we obtain all the G-invariant t-designs.

References

[1] Peter J. Cameron, *Permutation Groups*, London Math. Soc. Student Texts **45**, Cambridge University Press, Cambridge, 1999.

[2] John D. Dixon and Brian Mortimer, *Permutation Groups*, Springer, 1996.

Peter J. Cameron
July 6, 2004