Permutation Groups and Transformation Semigroups
Lecture 1: Permutation groups and group actions

Peter J. Cameron
Permutation Groups summer school, Marienheide
18–22 September 2017

In this first introductory lecture, I introduce some basic ideas about permutation groups: their connection with group actions; orbit decomposition; primitivity and multiple transitivity.

These ideas will be expanded by the other lecturers. My intention is to include just what I need for the remainder of my lectures.

1 Permutation groups and group actions

Let Ω be a set, which may be finite or infinite (but will usually be finite). We denote by $\mathrm{Sym}(\Omega)$ the *symmetric group* on Ω, the group whose elements are all the *permutations* of Ω (the bijective maps from Ω to itself), with the operation of composition.

If Ω is finite, say $|\Omega| = n$, we often write $\mathrm{Sym}(\Omega)$ as S_n.

Remark We compose permutations from left to right, so that g_1g_2 means "apply first g_1, then g_2". This goes naturally with writing a permutation on the right of its argument:

$$\alpha(g_1g_2) = (\alpha g_1)g_2.$$

Now a *permutation group* on Ω is simply a subgroup of $\mathrm{Sym}(\Omega)$; that is, a permutation group G is a set of permutations of Ω which is closed under composition, contains the identity permutation, and contains the inverse of each of its elements.

Remark Let S be a mathematical structure of virtually any type built on the set Ω. Then the automorphism group of S is usually a permutation group on Ω. (A little care is required: if S is a topology, then taking "automorphism" to mean "continuous bijection" does not work; we should take "automorphism" to be "homeomorphism" in this case.)

There is a related concept, that of a *group action*.

Let G be a group (in the abstract sense of group theory, a set with a binary operation). Then an *action* of G on Ω is a homomorphism from G to $\mathrm{Sym}(\Omega)$; in other words, it associates a permutation with each element of G. The image of a group action is a permutation group; the extra generality is that the action may have a kernel. The extra flexibility is important, since the same group may act on several different sets.

Example As a running example, let G be the group of symmetries of a cube (Figure 1).

Figure 1: A cube

Let Ω be the set of size 26 consisting of the 8 vertices, 12 edges, and 6 faces of the cube. Then G acts on Ω; the action is faithful (no symmetry can fix all the vertices except the identity), so we can regard G as a permutation group on Ω.

It is often the case, as in the examples below, that when we say "Let G be a permutation group on Ω", we could as well say "Let the group G act on Ω". For example, any permutation group property immediately translates to group actions.

2 Orbits and transitivity

In our example, the group G contains permutations which map any vertex to another vertex; we cannot map a vertex to an edge. We formalise this by the notion of orbits.

Let G be a permutation group on Ω. Define a relation $\sim$ on Ω by the rule

$$\alpha \sim \beta \text{ if and only if there exists } g \in G \text{ such that } \alpha g = \beta.$$

Question Show that $\sim$ is an equivalence relation on Ω. (You will find that the reflexive, symmetric and transitive laws correspond to the identity, inverse, and closure properties of G.)

Defend the thesis "Most equivalence relations arising in practice come from groups in the way just described."

Since $\sim$ is an equivalence relation, Ω decomposes as a disjoint union of its equivalence classes. These classes are called *orbits*.

In our running example, the sets of vertices, edges and faces form the three orbits of G.

Question Take a golf ball; calculate the group of rotational symmetries, and count its orbits on the set of dimples on the ball.

If a permutation group has just a single orbit, we say that it is *transitive*.

This can be put into group-theoretic terms. For $\alpha \in \Omega$, we define the *stabiliser* of α in G to be the subgroup

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

of G. [Check that it is a subgroup!] We write G_α for the stabiliser of α in G.

In the other direction, let H be an arbitrary subgroup of G. Let $G : H$ denote the set of right cosets of H in G. (This is sometimes written as $H\backslash G$.) Then there is an action of G on $G : H$, defined by the rule that the group element g induces the permutation π_g of $G : H$, where

$$(Hx)\pi_g = Hxg.$$

[Check that, in this action, the stabiliser of the element H is the subgroup H, while the stabiliser of Hx is the conjugate $x^{-1}Hx$.] This is the action of G by *right multiplication* on the coset space $G : H$.

Now there is a notion of *isomorphism* of group actions, and the following theorem is true:

Theorem 2.1 *Let G act transitively on Ω. For $\alpha \in \Omega$, let H be the stabiliser of α. Then the given action of G on Ω is isomorphic to the action of G on the set $G : H$ of right cosets of H by right multiplication.*

Moreover, the actions of G on coset spaces $G : H$ and $G : K$ are isomorphic if and only if H and K are conjugate subgroups of G.

Remark I have used the notation $H\backslash G$ for the right coset space; the companion notation for the left coset space is G/H. This notation is commonly used by geometers. The disadvantage is that group theorists are made unhappy by seeing G/H when H is not a normal subgroup of G. Another notation for the right coset space is $G : H$, as was used by Csaba in his talks. This has the advantage that the index of H in G (the number of cosets) is $|G : H|$. The small disadvantage is that there is no companion notatiaon for the left coset space.

We have given the conventional definition of transitivity. I will now give a different definition which can be used for all the other concepts I need.

Let S be a mathematical structure on the set Ω. I will say that S is *trivial* if it is preserved by the symmetric group $\mathrm{Sym}(\Omega)$, and *non-trivial* otherwise.

Thus, a subset A of Ω is trivial if and only if either $A = \emptyset$ or $A = \Omega$. Hence we can say,

> The permutation group G on Ω is transitive if and only if the only G-invariant subsets of Ω are the trivial ones.

Other examples we will meet later include the following:

- A partition of Ω is trivial if and only if it is either the partition into sets of size 1 or the partition with a single part.
- A graph on the vertex set Ω is trivial if and only if it is either the null graph or the complete graph.

A permutation group G on Ω is *regular* if it is transitive and the stabiliser of any point is the identity. [**Question:** Why are the order and degree of a regular permutation group equal?] *Cayley's Theorem* says that every group is isomorphic to a regular permutation group. So every group of order n is isomorphic to a subgroup of S_n; but the theorem works in the infinite case too.

3 Primitivity

I will treat the remaining concepts more briefly; these will reappear in the other lectures.

Let G be a transitive permutation group on Ω. We say that G is *primitive* if the only G-invariant partitions of Ω are the trivial ones. Thus G is *imprimitive* if it preserves some non-trivial partition of Ω.

An equivalence class B of a G-invariant equivalence relation has the property that, for all $g \in G$, either $Bg = B$, or $B \cap Bg = \emptyset$. A set with this property is called a *block (of imprimitivity)* for G. Thus, G is primitive if and only if the only blocks are the empty set, singletons, and Ω.

In our example, let G be the symmetry group of the cube, and let Ω_0 be the G-orbit consisting of the vertices of the cube. The action of G on Ω_0 is imprimitive. In fact, there are two non-trivial partitions preserved by G:

- the vertices of the cube fall into two interlocking regular tetrahedra, which are preserved or interchanged by all symmetries;
- there is a partition into four pairs of antipodal vertices, which is also preserved.

Theorem 3.1 *(a) Let G be a transitive permutation group on Ω, where $|\Omega| > 1$. Then G is primitive if and only if the stabiliser of a point of Ω is a maximal proper subgroup of G.*

(b) Let G be primitive on Ω. Then every non-trivial normal subgroup of G is transitive.

(c) Let G be primitive on Ω. Then G has at most two minimal normal subgroups; if there are two, then they are isomorphic and non-abelian, and each of them acts regularly.

We saw that every group is isomorphic to a transitive permutation group (Cayley's Theorem). The last part of the theorem above shows that not every group is isomorphic to a primitive permutation group.

4 Basic groups and O'Nan–Scott

In this section we specialise to finite groups.

A *Cartesian structure* on Ω is an identification of Ω with A^d, where A is some set. We can regard A as an "alphabet", and A^d as the set of all words of length d over the alphabet A. Then A^d is a metric space, with the *Hamming metric* (used in the theory of error-correcting codes): the distance between two words is the number of positions in which they differ.

A Cartesian structure is non-trivial if $|A| > 1$ and $d > 1$.

Let G be a primitive permutation group on Ω. We say that G is *basic* if it preserves no non-trivial Cartesian structure on Ω. (Although this concept is only defined for primitive groups, we see that the imprimitive group we met earlier, the symmetry group of the cube acting on the vertices, does preserve a Cartesian structure. Note that the automorphism group of a Cartesian structure over an alphabet of size 2 is necessarily imprimitive – generalise our argument for the cube to see this.)

The group-theoretic structure of basic groups is even more restricted. Part of the celebrated O'Nan–Scott Theorem asserts the following. In this theorem, a permutation group G is called *affine* if it acts on a vector space V and its elements are products of translations and invertible linear transformations of V, so that G contains all the translations. It is *almost simple* if $T \leq G \leq \mathrm{Aut}(T)$, where T is a finite simple group and $\mathrm{Aut}(T)$ its automorphism group (T embeds into $\mathrm{Aut}(T)$ as the group of inner automorphisms or conjugations). I will not define *diagonal* groups, but simply give an example. Let T be a finite simple group. Then $T \times T$, acting on T by the rule

$$x(g,h) = g^{-1}xh \text{ for all } x,g,h \in G,$$

is a diagonal group. (The name comes from the fact that the stabiliser of the identity is the *diagonal subgroup* $\{(g,g) : g \in G\}$ of $G \times G$.)

Theorem 4.1 *Let G be a finite basic primitive permutation group. Then G is affine, diagonal, or almost simple.*

See Pablo's lectures for much more detail on this.

5 Multiple transitivity

For any permutation group G on Ω, there is an induced action of G on the set of t-element subsets of Ω, or of t-tuples of elements of Ω, for any natural number t. This is defined in the obvious way:

$$\begin{aligned} \{\alpha_1, \ldots, \alpha_t\}g &= \{\alpha_1 g, \ldots, \alpha_t g\}, \\ (\alpha_1, \ldots, \alpha_t)g &= (\alpha_1 g, \ldots, \alpha_t)g. \end{aligned}$$

We say that G is *t-homogeneous* (or t-set transitive) if it acts transitively on the set of t-element subsets; and G is *t-transitive* if it acts transitively on the set of t-tuples of *distinct* elements of Ω. (The word "distinct" is necessary here; for example, no permutation can carry the pair (α, α) to (β, γ) if $\beta \neq \gamma$.)

It is clear that, for $t \leq |\Omega|$, a t-transitive group is t-homogeneous.

A consequence of the *Classification of Finite Simple Groups* (CFSG) is that all finite 2-transitive groups are known: indeed:

Theorem 5.1 *For $t \geq 6$, the only finite t-transitive groups are the symmetric and alternating groups.*

For $t = 2$, we can put this in terms of non-trivial structures:

- G is 2-homogeneous if and only if it preserves no non-trivial graph on the vertex set Ω.
- G is 2-transitive if and only if it preserves no non-trivial directed graph on Ω.

For larger t, we could formulate these notions in terms of "hypergraphs", but I will not be concerned with this.

Later in the lectures I will say more about non-trivial G-invariant graphs, which will also be treated by other lecturers.

To summarise some of this in a table. Going down the table, the conditions are meant to become stronger; so we assume that primitive groups are transitive, basic groups are primitive, and so forth. Sometimes these implications hold without being asserted.

Condition	Preserves no notrivial
Transitive	subset
Quasiprimitive	—
Primitive	partition disjoint union of complete graphs
Basic	Cartesian structure Hamming graph
...	...
2-homogeneous	undirected graph
2-transitive	directed graph

Cheryl's talks concern quasiprimitive groups, which have no simple characterisation in this sense. Csaba and Pablo discuss things around the primitive and basic borderlines.

In my third lecture, time permitting, I will insert two more properties between basic and 2-homogeneous (called *synchronizing* and *separating*), and in the fifth lecture I will insert one more (the *road closure property*).

There is no need to go further, since the 2-transitive finite groups are known.