Left Actions by Groups

10/31/06 (revision of 10/01/06) Radford

The notion of group is based on functions $G\times G \longrightarrow G$. Suppose that G is a group and A is a non-empty set. To study groups it is very convenient to consider more general functions $G\times A \longrightarrow A$ which satisfy "monoid type" axioms. For a set S, which could be infinite, we let $|S|$ denote the cardinality of S.

Let $G\times A \longrightarrow A$ be a function which we describe by $(g, a) \mapsto g{\cdot}a$. For $g \in G$ we define

$$\sigma_g : A \longrightarrow A$$

by

$$\sigma_g(a) = g{\cdot}a$$

for all $a \in A$. Then

$$\sigma_e(a) = e{\cdot}a, \quad (\sigma_g{\circ}\sigma_h)(a) = \sigma_g(\sigma_h(a)) = g{\cdot}(h{\cdot}a), \quad \text{and} \quad \sigma_{gh}(a) = (gh){\cdot}a$$

for all $a \in A$ and $g, h \in G$. The function $G\times A \longrightarrow A$ is a *left action of* G *on* A if

$$e{\cdot}a = a \qquad \text{and} \qquad g{\cdot}(h{\cdot}a) = (gh){\cdot}a$$

for all $a \in A$ and $g, h \in A$; that is

$$\sigma_e = \mathrm{Id}_A \qquad \text{and} \qquad \sigma_g{\circ}\sigma_h = \sigma_{gh}$$

for all $g, h \in G$.

Suppose the map $G\times A \longrightarrow A$ is a left action. We will say that G *acts on* A (*on the left*). Let $g \in G$. Then $\sigma_g{\circ}\sigma_{g^{-1}} = \sigma_{gg^{-1}} = \sigma_e = \mathrm{Id}_A$. Consequently $\sigma_{g^{-1}}{\circ}\sigma_g = \sigma_{g^{-1}}{\circ}\sigma_{(g^{-1})^{-1}} = \mathrm{Id}_A$. We have shown that σ_g and $\sigma_{g^{-1}}$ are function inverses; in particular $\sigma_g \in S_A$. Let

$$\pi : G \longrightarrow S_A$$

be the function defined by $\pi(g) = \sigma_g$ for all $g \in G$. The calculation

$$\pi(g)\circ\pi(h) = \sigma_g\circ\sigma_h = \sigma_{gh} = \pi(gh)$$

for all $g, h \in G$ shows that π is a homomorphism. The map π is called a *permutation representation of* G. Note that

$$g{\cdot}a = \pi(g)(a) \tag{1}$$

for all $g \in G$ and $a \in A$.

Conversely, suppose that $\pi : G \longrightarrow S_A$ is a homomorphism. Then $\pi(e) = \mathrm{Id}_A$. Define a function $G{\times}A \longrightarrow A$ by (1) and set $\sigma_g = \pi(g)$ for all $g \in G$. Then $\sigma_e = \mathrm{Id}_A$ and $\sigma_g\circ\sigma_h = \sigma_{gh}$ for all $g, h \in G$. Our function $G{\times}A \longrightarrow A$, which is defined by $(g, a) \mapsto \pi(g)(a)$, is a left action of G on A and π is the associated permutation representation. Thus the left actions of G on A are in bijective correspondence with the homomorphisms $\pi : G \longrightarrow S_A$.

Suppose that $G{\times}A \longrightarrow A$ is a left action of G on A. There are two basic types of associated actions which arise from restriction.

Let $H \leq G$. Then the action of G on A restricts to a left action of H on A. Suppose that B is a non-empty subset of A such that $g{\cdot}b \in B$ for all $g \in G$ and $b \in B$. Then the action on A restricts to a left G-action on B.

1 Orbits and Stabilizers

Throughout this section $G{\times}A \longrightarrow A$ is a left action of G on a non-empty set A. We continue with the notation above.

Let $a \in A$. Then

$$G{\cdot}a = \{g{\cdot}a \,|\, g \in G\}$$

is the *G-orbit of a*. The relation on A defined by $a \sim b$ if and only if $b = g{\cdot}a$ for some $g \in G$ is an equivalence relation on A. Observe that

$$[a] = G{\cdot}a;$$

that is the equivalence class containing a and the G-orbit of a are one in the same. Since equivalence classes partition:

$$\text{The } G\text{-orbits of } A \text{ partition } A. \tag{2}$$

The subset of G defined by

$$G_a = \{g \in G \mid g{\cdot}a = a\}$$

is called the *stabilizer of* a. It is easy to see that $G_a \leq G$.

Consider the function $f : G \longrightarrow G{\cdot}a$ defined by $f(g) = g{\cdot}a$ for all $g \in G$. Since f is surjective, $g{\cdot}a \mapsto f^{-1}(g{\cdot}a)$ defines a bijection between the orbit $G{\cdot}a$ and the set of fibers of f. We show that

$$f^{-1}(g{\cdot}a) = gG_a \tag{3}$$

for all $g \in G$. To see this, first suppose that $x \in gG_a$. Then $x = gh$ for some $h \in G_a$. Thus

$$f(gh) = (gh){\cdot}a = g{\cdot}(h{\cdot}a) = g{\cdot}(a) = g{\cdot}a$$

which shows that $gG_a \subseteq f^{-1}(g{\cdot}a)$. To complete the proof we need only show that $f^{-1}(g{\cdot}a) \subseteq gG_a$.

Suppose that $x \in f^{-1}(g{\cdot}a)$. Then $f(x) = g{\cdot}a$. Since $f(x) = x{\cdot}a$, from $g{\cdot}a = x{\cdot}a$ we deduce that $a = (g^{-1}x){\cdot}a$. Therefore $g^{-1}x \in G_a$ which means $x = g(g^{-1}x) \in gG_a$. We have shown $f^{-1}(g{\cdot}a) \subseteq gG_a$.

By (3) the elements of $G{\cdot}a$ are in one-one correspondence with the set of left cosets of G_a in G. Therefore

$$|G{\cdot}a| = |G : G_a| \tag{4}$$

for all $a \in A$. In particular $|G{\cdot}a|$ divides $|G|$ for all $a \in A$ when G is finite.

Let $\pi : G \longrightarrow S_A$ be the permutation representation associated with the left action. Then

$$\begin{aligned} \operatorname{Ker}\pi &= \{g \in G \mid \pi(g) = \mathrm{Id}_A\} \\ &= \{g \in G \mid \pi(g)(a) = a \ \ \forall a \in A\} \\ &= \{g \in G \mid g{\cdot}a = a \ \ \forall a \in A\} \end{aligned}$$

which means that

$$\operatorname{Ker}\pi = \bigcap_{a \in A} G_a, \tag{5}$$

the intersection of the stabilizers of all of the elements of A.

2 The Transitive Case

Throughout this section $G{\times}A \longrightarrow A$ is a left action of G on a non-empty set A. By (2) the G-orbits of A partition A. The action is called *transitive* if there is only one orbit; that is the partition has one cell.

Lemma 1 *Suppose that G acts on a non-empty set A transitively and write $A = G{\cdot}a$, where $a \in A$. Let $\pi : G \longrightarrow S_A$ be the associated permutation representation. Then:*

(a) $\operatorname{Ker}\pi \trianglelefteq G$ *and* $\operatorname{Ker}\pi \leq G_a$.

(a) *If $N \trianglelefteq G$ and $N \leq G_a$ then $N \leq \operatorname{Ker}\pi$.*

PROOF: Part (a) follows from the fact that kernels of homomorphisms are normal subgroups and (5). To show part (b), suppose that $N \trianglelefteq G$ and $N \leq G_a$. To show that $N \leq \operatorname{Ker}\pi$ we need only show that $n{\cdot}x = x$ for all $x \in A$; that is $n{\cdot}(g{\cdot}a) = g{\cdot}a$ for all $g \in G$.

Let $g \in G$. Then $g^{-1}ng = g^{-1}n(g^{-1})^{-1} \in N$ since $N \trianglelefteq G$. Thus

$$n{\cdot}(g{\cdot}a) = g{\cdot}((g^{-1}ng){\cdot}a) = g{\cdot}(a) = g{\cdot}a$$

and we are done. $\Box$

We may paraphrase the conclusion of the lemma by saying that $\operatorname{Ker}\pi$ is the largest normal subgroup of G contained in G_a.

3 The Case when G is Finite Cyclic

Proposition 1 *Suppose that $G = <g>$ is cyclic of order n and acts on A on the left. Let $a \in A$ and $|G{\cdot}a| = m$. Then:*

(a) *m divides n.*

(b) *The m-element set $G{\cdot}a = \{a, g{\cdot}a, \ldots, g^{m-1}{\cdot}a\}$ and $g^m{\cdot}a = a$.*

PROOF: Part (a) follows by (4) since the index of a subgroup of a finite group divides the order of the group. As for part (b), note that the list

$$a = e{\cdot}a = g^0{\cdot}a, g{\cdot}a = g^1{\cdot}a, g^2{\cdot}a, g^3{\cdot}a, \ldots$$

has a repetition since $G{\cdot}a$ is finite and mimic the steps in the analysis of the cyclic group $G = <g>$ which starts with the list $e = g^0, g^1, g^2, g^3, \ldots$. $\Box$

4 Application to Permutations

Let $n \geq 1$ and $\mathcal{G} = S_n$. Then $\mathcal{G}$ acts on $A = \{1, 2, \ldots, n\}$ by function evaluation:

$$\sigma{\cdot}\ell = \sigma(\ell)$$

for all $\sigma \in \mathcal{G}$ and $1 \leq \ell \leq n$. Let $\tau \in \mathcal{G}$ and set $G = <\tau>$. Then G acts on A by restriction. Let $\ell \in A$, let $m = |G{\cdot}\ell|$, and let $n = |G|$ which is the order of τ. Then m divides n by part (a) of Proposition 1. By part (b) of the same $G{\cdot}\ell = \{\ell, \tau(\ell), \tau^2(\ell), \ldots, \tau^{m-1}(\ell)\}$ and $\tau^m(\ell) = \ell$. The effect of τ on $G{\cdot}\ell$ is the same as the m-cycle

$$(\ell \; \tau(\ell) \; \cdots \; \tau^{m-1}(\ell)).$$

Observe that the order of the m-cycle is its length m. Since the G-orbits of A partition A we conclude that τ is the product of disjoint cycles and their orders (lengths) divide the order of τ by part (b) of Proposition 1. Usually 1-cycles are omitted from the product since they are the identity map. If τ is written as the product of disjoint cycles then each cycle accounts for a G-orbit of A. We have essentially shown:

Proposition 2 *Suppose that $n > 1$ and $\mathrm{Id} \neq \tau \in S_n$. Then:*

(a) *τ is the product of disjoint cycles of length greater than one. The cycles commute and this decomposition is unique up to reordering factors.*

(b) *The order of τ is the least common multiple of the orders (lengths) of the non-trivial cycles of part (a).*

□

We refer to $G{\cdot}\ell$ as a τ-orbit. Let $(a\ b)$ be a transposition and consider the product $\tau' = \tau(a\ b)$. We will show that the τ'-orbits are the τ-orbits with one exception: either two of the τ-orbits combine to give one τ'-orbit or one of the τ-orbits splits into two τ'-orbits. Observe that if a τ-orbit contains neither a nor b then it is a τ'-orbit.

Case 1: a and b are in different τ-orbits.

By part (b) of Proposition 1 we may write these orbits as

$$\{a, \tau(a), \ldots, \tau^{r-1}(a)\} \qquad \{b, \tau(b), \ldots, \tau^{s-1}(b)\}$$

where $1 \leq r, s$ and $\tau^r(a) = a$, $\tau^s(b) = b$. Observe that the τ'-orbit of a is

$$\{a, \tau(b), \ldots, \tau^{s-1}(b), b, \tau(a), \ldots, \tau^{r-1}(a)\}$$

which is the union of the two τ-orbits. *Thus τ' combines these two τ-orbits into a single τ'-orbit.*

Case 2: a and b are in the same τ-orbit.

We may write this orbit as

$$\{a, \tau(a), \ldots, \tau^r(a), \ldots, \tau^{s-1}(a)\},$$

where $s \geq 2$, $\tau^s(a) = a$, $1 \leq r \leq s-1$, and $\tau^r(a) = b$. Observe that this orbit splits into two τ'-orbits which are

$$\{\tau(a), \ldots, \tau^r(a)\} \qquad \text{and} \qquad \{a, \widehat{\tau(a)}, \ldots, \widehat{\tau^r(a)}, \ldots, \tau^{s-1}(a)\},$$

where the "hat" symbol means omission. *Thus τ' splits this τ-orbit into two τ'-orbits.*

Lemma 2 *Let $\tau_1, \ldots, \tau_r \in S_n$ be transpositions and suppose $\tau_1 \cdots \tau_r = \mathrm{Id}$. Then r is even.*

PROOF: Consider the sequence

$$\mathrm{Id}, \mathrm{Id}\tau_1, \mathrm{Id}\tau_1\tau_2, \ldots, \mathrm{Id}\tau_1 \cdots \tau_r.$$

Let c be the number of times the orbits of a term in the sequence are formed by combining two orbits of its predecessor and let s be the number of times they are formed by splitting an orbit of its predecessor. Then $r = c + s$. Now n is the number of orbits of Id. Thus $\mathrm{Id}\tau_1 \cdots \tau_r$ has $n + s - c$ orbits. Since this permutation is Id it follows that $n + s - c = n$. Therefore $s = c$ and r is even. $\Box$

Corollary 1 *Suppose that $\tau_1, \ldots, \tau_r, \tau'_1, \ldots, \tau'_{r'} \in S_n$ are transpositions and $\tau_1 \cdots \tau_r = \tau'_1 \cdots \tau'_{r'}$. Then r and r' are both even or they are both odd.*

PROOF: We build on the proof of the previous Lemma. Since $\mathrm{Id}\tau_1 \cdots \tau_r = \mathrm{Id}\tau'_1 \cdots \tau'_{r'}$ we have the equation $n+s-c = n+s'-c'$ from which $s-c = s'-c'$ follows. Thus

$$r' = c' + s' = c + s + (c' - c) + (s' - s) = c + s + 2(c' - c) = r + 2(c' - c)$$

which completes our proof. $\square$

Let $n \geq 2$ and $\sigma = (a_1\ a_2\ \ldots\ a_r) \in S_n$. When $r > 2$ then σ is the product of transpositions in various ways, for example

$$\begin{aligned} \sigma &= (a_1\, a_2\, \ldots\, a_r) \\ &= (a_r\, a_1) \cdots (a_3\, a_1)(a_2\, a_1) \\ &= (a_1\, a_2)(a_r\, a_2) \cdots (a_4\, a_2)(a_3\, a_2) \\ &\vdots \end{aligned}$$

since $\sigma = (a_1\ a_2\ \ldots\ a_r) = (a_2\ a_3\ \ldots\ a_r\ a_1) = \cdots$. Thus by part (b) of Proposition 2 every permutation is the product of transpositions.

A permutation is called *even* it it can be written as a product of an even number of transpositions and is called *odd* otherwise. Thus, by definition, if an odd permutation is written as a product transpositions the number of transpositions must be odd. By virtue of the preceding corollary, if an even permutation is written as a product of transpositions the number of transpositions must be even.

Define $\varsigma : S_n \longrightarrow \{-1, 1\}$ by

$$\varsigma(\tau) = \begin{cases} 1 & : \tau \text{ is even} \\ -1 & : \tau \text{ is odd} \end{cases}$$

Let $\sigma, \tau \in S_n$. If σ, τ are even, or they are odd, then $\sigma\tau$ is even. If one of σ, τ is even and one is odd then $\sigma\tau$ is is odd. Thus ς is a homomorphism to the multiplicative subgroup $\{-1, 1\}$ of the non-zero real numbers. Note that $A_n = \mathrm{Ker}\, \varsigma$ is a set of even permutations of S_n. It is easy to see that

$$A_n \trianglelefteq S_n \qquad \text{and} \qquad |S_n : A_n| = 2.$$

5 Cayley's Theorem

Let G be any group and let $\mathcal{A}$ be the set of all non-empty subsets of G. Then G acts on $\mathcal{A}$ by

$$s{\cdot}S = gS$$

for all $g \in G$ and $S \in \mathcal{A}$. For a subset $S \in \mathcal{A}$ observe that

$$G{\cdot}S = \{gS \mid g \in G\}. \tag{6}$$

Now suppose that $H \leq G$. Then with $S = H$ we see by (6) that $G{\cdot}H$ is the set of left cosets of H in G. The action of G on $\mathcal{A}$ restricts to an action of G on the set of left cosets $A = G{\cdot}H$ of H in G. Observe that the stabilizer of H is

$$G_H = \{g \in G \mid gH = H\} = H.$$

Let $\pi : G \longrightarrow S_A$ be the corresponding permutation representation. Then $\mathrm{Kef}\,\pi$ is the largest normal subgroup of G contained in H by part (b) of Lemma 1. If the only normal subgroup of G contained in H is (e) then π is injective. This is the case when $H = (e)$; here we may identify the set of left cosets of H with G with since $gH = \{ge\} = \{g\}$ for all $g \in G$.

Theorem 1 *Let G be a group. Then G is isomorphic to a subgroup of the permutation group S_G.* $\square$

When G if finite $S_G \simeq S_{|G|}$.

Corollary 2 (Cayley's Theorem) *Every finite group is isomorphic to a subgroup of S_n for some positive integer n.* $\square$

6 The Class Equation and a Generalization

As in the previous section, let G be any group and let $\mathcal{A}$ be the set of all non-empty subsets of G. Then G acts on $\mathcal{A}$ by

$$g{\cdot}S = gSg^{-1}$$

for all $g \in G$ and $S \in \mathcal{A}$. For a element $S \in \mathcal{A}$ observe that

$$G{\cdot}S = \{gSg^{-1} \mid g \in G\} \tag{7}$$

is the set of conjugates of S in G and the stabilizer

$$G_S = \{g \in G \mid gSg^{-1} = S\} = \mathrm{N}_G(S)$$

is the normalizer of S in S. Thus

$$|G : \mathrm{N}_G(S)| = |G{\cdot}S| \tag{8}$$

by (4). As a consequence, when G is finite the number of conjugates of a non-empty subset of G divides the order of G.

Suppose that $S = \{s\}$ is a singleton set and let $g \in G$. since $g\{s\}g^{-1} = \{s\}$ if and only if $gsg^{-1} = s$, or equivalently $gs = sg$,

$$\mathrm{N}_G(\{s\}) = \mathrm{C}_G(\{s\}) = \mathrm{C}_G(s). \tag{9}$$

Since $g{\cdot}\{s\} = \{s\}$ it follows that G acts on the set of all singleton subsets of G. Identifying s with $\{s\}$ gives us the left action of G on itself by conjugation; that is

$$g{\cdot}s = gsg^{-1}.$$

The class equation is derived from an analysis if the conjugation action of G on itself.

For $g \in G$ the element gsg^{-1} is called a *conjugate* of s. The orbit

$$G{\cdot}s = \{g{\cdot}s \mid g \in G\} = \{gsg^{-1} \mid g \in G\}$$

is thus the set of conjugates of s and is called *the conjugacy class of* s. Since $s \in G{\cdot}s$, note that

$$|G{\cdot}s| = 1 \quad \text{if and only if} \quad s \in \mathrm{Z}(G). \tag{10}$$

Now suppose that G is finite and let $G{\cdot}s_1, \ldots G{\cdot}s_r$ be a listing of the distinct orbits with more than one element. As $|G{\cdot}s_i| = |G : C_G(s_i)|$ by (8) and (9), we have the class equation:

$$|G| = |\mathrm{Z}(G)| + \sum_{i=1}^{r} |G : C_G(s_i)|, \tag{11}$$

where $|G : C_G(s_i)| > 1$ for al $1 \leq i \leq r$.

A finite group G is a *p-group* if p is a prime integer and $|G| = p^m$ for some $m \geq 1$. Such a group is not simple as:

Proposition 3 *A finite p-group has a non-trivial center.*

PROOF: Let G be a finite p-group and consider the class equation. Since the index of a subgroup of a finite group divides the order of the group, p divides $|G : C_G(s_i)|$ for all $1 \leq i \leq r$. Since p divides $|G|$, by the class equation p divides $|\mathrm{Z}(G)|$. Therefore $\mathrm{Z}(G) \neq (e)$. $\square$

There is a generalization of the class equation for left actions of a group G on a finite set A. Let $z(A)$ be the set of elements $a \in A$ such that $G{\cdot}a = \{a\}$

and suppose that $G\cdot a_1, \ldots, G\cdot a_r$ is a list of the distinct orbits of A with more than one element. Then

$$|A| = |z(A)| + \sum_{i=1}^{r} |G : G_{a_i}| \tag{12}$$

since $|G\cdot a_i| = |G : G_{a_i}|$ by (2). Part (b) of the following proposition generalizes Proposition 3.

Proposition 4 *Suppose that G is a finite p-group.*

(a) *Let A be a finite set on which G acts on the left. Then $|A| = |z(A)| + pk$ for some non-negative integer k. In particular p divides $|z(A)|$ if and only if p divides $|A|$.*

(b) *Let $(e) \neq N \trianglelefteq G$. Then $N \cap Z(G) \neq (e)$.*

Proof: Since G is finite $|G\cdot a| = |G : G_a|$ divides $|G|$ for all $a \in A$. Thus part (a) follows by (12). As for part (b) we note that G acts on N by conjugation. Since $|N|$ divides $|G|$, we conclude from part (a) that p divides the order of $z(N) = N \cap Z(G)$. $\square$