

CDM

Group Actions and Counting

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

SPRING 2021



1 Some Counting Problems

2 Actions

3 Burnside's Lemma

4 Some Applications


Here is our next **Big Problem**:

We have a finite set A and an equivalence relation $\approx$ on A . We want to count the number of equivalence classes of $\approx$ (also known as the index of the equivalence relation).

Think of an equivalence class as a **pattern**, so we want to count patterns.

Of course, A will be large. In fact, often we have a whole parametrized family A_n and we want an answer in terms of a function of n .

Suppose we want to enumerate hydrocarbon molecules where some hydrogen atoms have been replaced by hydroxyl groups. We could use alkenes (C_2H_5OH), but let's work with carbocycles like benzene instead.



Some counting: $\binom{6}{2} = 15$, $\binom{5}{1} = 5$ or, taking into account symmetry, 3.

Only the last answer makes chemical sense.

G. Pólya

Kombinatorische Anzahlbestimmungen für Gruppen, Graphen
und chemische Verbindungen

Acta Mathematica 68 (1937) 1: 145–254.

Actually, Pólya was scooped:

J. H. Redfield


The Theory of Group-Reduced Distributions

American J. Mathematics 49 (1927) 3, 433–455

How many different ways are there to place 3 crosses and 3 naughts on a (standard 3×3) Tic-Tac-Toe board?


No problem: there are $\binom{9}{3,3,3} = 1680$ possible placements.

OK, but what if we identify boards that can be obtained from rotation, or looking through a mirror?



Clearly this involves the dihedral group D_4 , and a single placement of marks can have as many as 8 variants.

So we should roughly expect $1680/8 = 210$ patterns.



But that is not quite right, either: some placements have fewer than 8 variants, so we are under-counting the number of patterns.

Here is a similar but somewhat less frivolous problem. Suppose we want to implement Boolean functions $f : \mathbf{2}^n \rightarrow \mathbf{2}$ as circuits.

There are 2^{2^n} such functions, but we don't need as many circuits. For example, we may have $f(x, y, z) = g(y, z, x)$ so it suffices to implement either f or g .

In general, we can permute the variables arbitrarily: two functions f and g are equivalent if

$$f(x) = g(\pi(x))$$

for some permutation π .

How many Boolean functions are there modulo input permutations?

Permuting the inputs is an obvious modification of a circuit, but there are other possibilities.

For example, it is straightforward to negate input bits and/or the output bit. Thus, there is another equivalence

$$f(x) = g(x \oplus c) \oplus d$$

where $c \in \mathbf{2}^n$, $d \in \mathbf{2}$ and $\oplus$ denotes exclusive or.

In other words, we allow flipping some input bits, as well as the output bit.

We can combine different kinds of modifications.

Given a permutation π on $[n]$, we can define the modification f of g :

$$\begin{aligned} f(\mathbf{x}) &= g(\pi(\mathbf{x}) \oplus \mathbf{c}) \oplus d \\ &= g(x_{\pi(1)} \oplus c_1, x_{\pi(2)} \oplus c_2, \dots, x_{\pi(n)} \oplus c_n) \oplus d \end{aligned}$$

So how many functions $\mathbf{2}^n \rightarrow \mathbf{2}$ are there modulo this identification?

This gets fairly messy fairly soon. On the face of it, it's not clear how to go about this kind of counting problem.

1 Some Counting Problems

2 **Actions**

3 Burnside's Lemma


4 Some Applications

To tackle this type of problem in style one can use some ideas from classical algebra:

- groups
- subgroups
- homomorphisms
- actions

These algebraic ideas are all 19th century, but the applications are 20th century. And, as we will see, to get real answers requires quite a bit of computation.

Groups describe symmetries, so if we want to count modulo symmetry it is natural to use groups to do so.



Observe that I write functions on the right and functional composition from left to right. This is undoubtedly the *Wave of the Future*. It makes functional diagrams easier to read and corresponds to the natural order of doing things on a pocket calculator.

This is from a 1976 paper “Some Applications of the Wreath Product Construction” by the category theorist Charles Wells. Obviously, Wells was the proud owner of an HP calculator.

Alas, he missed the boat on this one. By about a lightyear.

We will realize half of Wells’ dream: composition is diagrammatic, but we chicken out and write function application on the left. Yes, I know ...

Let's go back to the Tic-Tac-Toe problem from above: we want to count patterns, where two boards are equivalent if one can be moved to the other by rotations and/or reflections.

Clearly, we have to consider the interaction between Tic-Tac-Toe boards and the elements of the dihedral group D_4 .

As we have seen, we can safely assume that the group in question is always a subgroup of a symmetric group, so we are dealing with a class of permutations (but see below for some twists).

What is needed is some glue that connects the permutations with the objects we are interested in (such as the boards, or carbocycles, or circuits, ...)

We are dealing with

- a collection X of objects (configurations),
- a collection S of operations on the carrier set X .

A priori, X is just a flat set with no particular structure. We can think of S as a collection of atomic actions that can be performed on X .

It is entirely natural to consider composite operations that are obtained by applying a whole sequence of operations from S . This can be modeled naturally by the free monoid S^* , so we are dealing with **monoid actions**.

Free monoids Σ^* are the mathematical model for temporal sequences: we interpret aba as: first do a , then b , then a again.

Actually, this should sound eminently familiar: a DFA is a perfect example of an action of the free monoid Σ^* on a finite set Q :

Given “instruction” $a \in \Sigma$, apply δ_a to the current state $p \in Q$.

In our current setting, we are interested in actions that are reversible, so we have to deal with groups rather than just plain monoids. Other than that, it's very much the same idea.

Proviso: We will focus on a collection S of basic, reversible operations. In this case, it is natural to consider the group generated by S rather than the monoid.

The main idea is simply that s^{-1} should undo the effect of s .

As before with monoids, 1 should have no effect whatsoever and things need to be compatible in the right way.

Following the current Bourbaki/Hilbert standard, we need to axiomatize this basic idea.

Definition

Let G be a group and X a set. A **left action of G on X** is a function φ such that

$$\varphi : G \times X \rightarrow X$$

$$\varphi(a * b, x) = \varphi(a, \varphi(b, x))$$

$$\varphi(1, x) = x$$

Here $a, b \in G$ and $x \in X$. X is also called a **G -set**.

Notation:

It is customary to write $a \cdot x$ or even ax instead of $\varphi(a, x)$. Hence

$$(a * b) \cdot x = a \cdot (b \cdot x)$$

$$1 \cdot x = x$$

This is much better notation, albeit slightly dangerous.

We can push our luck a bit and write

$$(ab)x = a(bx)$$



$$1x = x$$

This gets to be particularly interesting if we deal with left actions and right actions (see below) at the same time.

In case of doubt, write out the multiplication operators. Or go anal-retentive and write φ if you want to play it absolutely safe.

Let $G = \text{GL}(2, \mathbb{R})$ be the group of invertible 2-by-2 matrices over the reals, and $X = \mathbb{R}^2$ the 2-dimensional plane.

Then G acts on X via $A \cdot x = Ax$.



The action induced by the rotation matrix $A = \frac{1}{2} \begin{pmatrix} 1 & -\sqrt{3} \\ \sqrt{3} & 1 \end{pmatrix}$.

Recall that $\mathfrak{S}(X)$ or $\mathfrak{S}_X$ denotes the group of all permutations of X under composition.

As usual, we use diagrammatic (l2r) composition for $\mathfrak{S}(X)$.

Definition

A **permutation group over X** is a subgroup G of $\mathfrak{S}(X)$.

The **order** of G is its cardinality and the **degree** of G is the cardinality of X .

By Cayley's Theorem every group G is isomorphic to a permutation group: we can think of the carrier set as being G . Note that in general the degree of G may be much smaller than the order of G , though.

An element of a permutation group can naturally be used to “rearrange” objects.

Consider a list of n objects:

$$\mathbf{x} = (x_1, x_2, \dots, x_n)$$

We can use any permutation f in $\mathfrak{S}_n$ to rearrange the elements of $\mathbf{x}$:

$$\mathbf{x}' = (x_{f(1)}, x_{f(2)}, \dots, x_{f(n)})$$

More generally, a permutation group of degree n can operate on n -vectors.

If we have some natural group G , say, some geometric group of rotations and/or reflections, we can translate it into a permutation group.

We connect G to a symmetry group by a homomorphism $\Phi : G \rightarrow \mathfrak{S}(X)$. Particularly interesting to us is the case where Φ is a monomorphism, but the idea works in general.

Then

$$\varphi(a, x) = \Phi(a^{-1})(x)$$

is a left action of G on X , and all left actions arise in this way.

You may find the a^{-1} a bit peculiar and probably expected a plain a instead. The reason we need the inverse is that we want a left action and we use diagrammatic composition.

$$\begin{aligned}\varphi(a * b, x) &= \Phi((a * b)^{-1})(x) \\ &= \Phi(b^{-1} * a^{-1})(x) \\ &= (\Phi(b^{-1}) \circ \Phi(a^{-1}))(x) \\ &= \Phi(a^{-1})(\Phi(b^{-1})(x)) \\ &= \varphi(a, \varphi(b, x))\end{aligned}$$

Exercise

What if we had failed Wells and used the wrong definition of composition?

Consider a permutation group $G \subseteq \mathfrak{S}_n$ (of degree n and order at most $n!$).

One very useful space of objects here is

$$X = A^n$$

the set of n -vectors over A , an arbitrary set (later A will often have additional structure, but for the time being it's a naked set).

Claim

G acts on X on the left via

$$f \cdot \mathbf{x} = (x_{f(1)}, x_{f(2)}, \dots, x_{f(n)})$$

It is clear that $1 \cdot x = x$.

Consider

$$(f \circ g) \cdot x = y_1$$

versus

$$f \cdot (g \cdot x) = y_2.$$

We need to show that $y_1 = y_2$.

Recall that we compose functions from left to right, so that

$$y_1 = (x_{g(f(1))}, \dots, x_{g(f(n))}).$$

$$\text{But then } y_2 = f \cdot (x_{g(1)}, \dots, x_{g(n)}) = y_1.$$

□

Right?

Think carefully – this looks absolutely wrong, but it's right. Take a good look at the following.

Write $u_i = x_{g(i)}$ by the free-country argument.

$$\begin{aligned}g \cdot \mathbf{x} &= (x_{g(1)}, x_{g(2)}, \dots, x_{g(n)}) \\ &= (u_1, u_2, \dots, u_n)\end{aligned}$$

$$\begin{aligned}f \cdot \mathbf{u} &= (u_{f(1)}, u_{f(2)}, \dots, u_{f(n)}) \\ &= (x_{g(f(1))}, x_{g(f(2))}, \dots, x_{g(f(n))})\end{aligned}$$

Exercise

Make sure you really understand the proof.

What would happen if we did composition the other way around?

Wurzelbrunft Wisdom: Where there's a left, there must be a right ...

Definition

Let G be a group and X a set. A **right action of G on X** is a function φ such that

$$\varphi : X \times G \rightarrow X$$

$$\varphi(x, a * b) = \varphi(\varphi(x, a), b)$$

$$\varphi(x, 1) = x$$

Needless to say, this is often written $x \cdot a$ and xa .

To maintain a semblance of sanity, we always write $a, b, c, \dots$ for group elements and $x, y, z, \dots$ for the elements of X .

As before for left actions, we can use group homomorphisms $\Phi : G \rightarrow \mathfrak{S}(X)$ to obtain right actions.

This time we define

$$\varphi(x, a) = \Phi(a)(x)$$

to get a right action of G on X , and vice versa.

Recall that for a left action we had to use a^{-1} , now the definition is perhaps a bit more natural.

Exercise

Verify that this definition really produces a right action.

Recall the operation of prefix quotients on languages: $w^{-1}L$ is obtained by deleting prefix w from words in L .

$$w^{-1}L = \{v \mid wv \in L\}$$

Since the prefix is deleted at the beginning (left end) of the word this is usually written on the left, as indicated – but these quotients are another example of a right action of the monoid Σ^* on the collection of all languages.

Alas, since we write the action on the left we get the awkward

$$(uv)^{-1}L = v^{-1}u^{-1}L.$$

Exercise

Check carefully that the last two examples really are right actions of the monoid Σ^ . Find better notation for the quotient operation.*

Exercise

Find some natural examples for left monoid actions.

Exercise

Verify that the “standard example” $x \cdot f = (x_{f^{-1}(1)}, x_{f^{-1}(2)}, \dots, x_{f^{-1}(n)})$ really produces a right action.

Again, here is our primary example of a group action: some permutation f from the symmetric group $\mathfrak{S}_n$ rearranging the elements of $x = (x_1, x_2, \dots, x_n)$.

Let $f(i) = j$, so the permutation moves i to j .

type	intuitively		result
left	x_i is replaced by x_j		$(x_{f(1)}, \dots, x_{f(n)})$
right	x_i moves to x_j		$(x_{f^{-1}(1)}, \dots, x_{f^{-1}(n)})$

Having two versions to deal with might seem plain annoying, but there are occasions when left actions are more convenient to work with, and there are occasions when right actions are more convenient. Grin and bear it.

For group actions, we can interchange left and right in the following sense.

Proposition

Consider two maps $\varphi : G \times X \rightarrow X$ and $\psi : X \times G \rightarrow X$ such that

$$\psi(x, a) = \varphi(a^{-1}, x).$$

Then φ is a left action if, and only if, ψ is a right action.

Proof. Suppose φ is a left action.

$$\begin{aligned}\psi(x, a * b) &= \varphi((a * b)^{-1}, x) \\ &= \varphi(b^{-1} * a^{-1}, x) \\ &= \varphi(b^{-1}, \varphi(a^{-1}, x)) \\ &= \psi(\psi(x, a), b)\end{aligned}$$

The other direction is entirely similar.

□

Another way to establish a connection between left and right actions is to reverse the multiplication. Given a group $\mathcal{G} = \langle G, \cdot \rangle$ define a new group

$$\mathcal{G}^{\text{op}} = \langle G, * \rangle \quad a * b = b \cdot a.$$

It is not hard to check that $\mathcal{G}^{\text{op}}$ is in fact a group.

Now any left action φ over $\mathcal{G}$ translates into a right action ψ over $\mathcal{G}^{\text{op}}$ by

$$\psi(x, a) = \varphi(a, x)$$

Exercise

Give a detailed proof of this claim.

From a sufficiently abstract perspective, left and right actions are the same: it doesn't matter much if we replace each group element by its inverse or change the order of multiplication. In fact, there are older texts that just speak about "a group acting on a set". The following is a classic, highly recommended.

N.G. de Bruijn

Pólya's Theory of Counting

E.F. Beckenbach (ed.): Applied Combinatorial Mathematics, Wiley (1964).

That's fine, but when it comes to actual implementation one has to be more careful, the code for both versions is different. More importantly, you must never mix the two versions within the same algorithm.



Recall the homomorphism $\Phi : G \rightarrow \mathfrak{S}(X)$ that associates each group element with the corresponding map on X .

Definition

An action is **faithful** if $\Phi(a) = \Phi(b)$ implies $a = b$.

If an action fails to be faithful, consider the kernel of Φ , the normal subgroup

$$H = \{ a \in G \mid \Phi(a) = I \}$$

Then the quotient G/H acts faithfully in the natural manner.

A left action is

- **transitive** if $\forall x, y \exists a (ax = y)$
- **free** if $\exists x (ax = bx)$ implies $a = b$
- **regular** if it is transitive and free

Free means that $ax = x = 1x$ implies $a = 1$; so free implies faithful.

Analogous definitions can be given for right actions.

Example (Regular Action)

Let G be any group. Then G acts on G via $a \cdot x = ax$.

Example (Conjugation Action)

Let G be any group. Then G acts on G via $a \cdot x = axa^{-1}$.

Example (Subgroup Conjugation Action)

Let G be any group and X the set of subgroups of G .

Then G acts on X via $a \cdot H = aHa^{-1} = \{aha^{-1} \mid h \in H\}$.

Define three word maps $\alpha, \beta, \gamma : \mathbf{2}^* \rightarrow \mathbf{2}^*$ by $\alpha(\varepsilon) = \beta(\varepsilon) = \gamma(\varepsilon) = \varepsilon$ and

$$\alpha(0x) = 1 \gamma(x)$$

$$\alpha(1x) = 0 \beta(x)$$

$$\beta(sx) = s \alpha(x)$$

$$\gamma(sx) = s \beta(x)$$

Claim

The maps α , β and γ are bijections.

Hence we can define an action of F_3 , the free group of rank 3, on $\mathbf{2}^*$.

Question: Is this action faithful?

Alas, this action is not faithful: the word maps commute and there is another somewhat unexpected identity:

$$\begin{aligned}\alpha\beta &= \beta\alpha & \alpha\gamma &= \gamma\alpha & \beta\gamma &= \gamma\beta \\ \alpha^2\beta^2\gamma &= I\end{aligned}$$

The hard part is finding these identities; they are easy to prove by induction.


With a little effort, one can show that these are the “only identities,” and we get a faithful action of $\mathbb{Z}^2$ on 2^* :

$$(a, b)x = \alpha^a\beta^b(x)$$

Exercise


Prove all these claims.

Let's come back to the dihedral group D_4 which we can think of as the symmetries of a square.




Abstractly, D_4 has the presentation

$$\langle \alpha, \beta \mid \alpha^4 = \beta^2 = 1, \beta\alpha = \alpha^3\beta \rangle$$



A 4x4 grid representing the Cayley table for the dihedral group D_4 . The grid contains 16 colored squares in a checkerboard pattern. The colors are purple, blue, green, yellow, orange, and red, arranged in a repeating 2x2 block pattern.



The Cayley table and Cayley graph for D_4 .

We can also represent the rigid motions of the plane in D_4 as 2×2 matrices:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$$
$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$$

This has the advantage that we can directly compute the image of a corner point $(\pm 1, \pm 1)$ of the square that is moved around.

The natural degree of D_4 is 4, and we can identify this group with a subgroup G of $\mathfrak{S}_4$ (the permutations that preserve adjacency):

$$\begin{array}{cccc} [1, 2, 3, 4] & [2, 3, 4, 1] & [4, 1, 2, 3] & [3, 4, 1, 2] \\ [3, 2, 1, 4] & [1, 4, 3, 2] & [4, 3, 2, 1] & [2, 1, 4, 3] \end{array}$$

It is this subgroup G that induces a regular action on the square.

For simplicity, one might express this by casually omitting G and saying something like “ D_4 acts on the square.”

Recall that any permutation can be decomposed into **disjoint cycles**:

$$\rho = (v_{1,1}, v_{1,2}, \dots, v_{1,q_1}), (v_{2,1}, \dots, v_{2,q_2}), \dots, (v_{p,1}, \dots, v_{p,q_p})$$

Note that it is customary to omit 1-cycles, so the notation can be ambiguous.

The **cycle shape** of a permutation is a (sorted) list of the cycle lengths. In the notation above, the cycle shape would be (the sorted version of) $q_1, \dots, q_p$.

Sometimes the cycle shape is expressed more compactly by using shamelessly exploiting polynomial notation. Type theorists should now swallow an aspirin.

ρ	cycles	shape	short shape
1	(1), (2), (3), (4)	(1, 1, 1, 1)	1^4
α	(1, 2, 3, 4)	(4)	4^1
α^2	(1, 3), (2, 4)	(2, 2)	2^2
α^3	(1, 4, 3, 2)	(4)	4^1
β	(1, 4), (2, 3)	(2, 2)	2^2
$\alpha\beta$	(2), (4), (1, 3)	(1, 1, 2)	$1^2 + 2^1$
$\alpha^2\beta$	(1, 2), (3, 4)	(2, 2)	2^2
$\alpha^3\beta$	(1), (3), (2, 4)	(1, 1, 2)	$1^2 + 2^1$

Here α and β are the standard generators of D_4 (a rotation and a reflection).

1^5	1
$1^3 + 2^1$	10
$1^2 + 3^1$	20
$1^1 + 2^2$	15
$1^1 + 4^1$	30
$2^1 + 3^1$	20
5^1	24

In this case, it is easy to check correctness.

$1^3 + 2^1$ corresponds to $\binom{5}{2} = 10$ choices of the points on the 2-cycle.

Similarly, there are 10 choices for the 2 fixed points in $1^2 + 3^1$; for each, the remaining 3 points can be arranged into a 3-cycle in 2 ways.

1^6	1
$1^4 + 2^1$	15
$1^3 + 3^1$	40
$1^2 + 2^2$	45
$1^2 + 4^1$	90
$1^1 + 2^1 + 3^1$	120
$1^1 + 5^1$	144
2^3	15
$2^1 + 4^1$	90
3^2	40
6^1	120

Make sure to check some of these numbers and explain where they come from.

1^7	1
$1^4 + 3^1$	70
$1^3 + 2^2$	105
$1^2 + 5^1$	504
$1^1 + 2^1 + 4^1$	630
$1^1 + 3^2$	280
$2^2 + 3^1$	210
7^1	720

This is quite similar to the last slide, but one also needs to handle the decomposition into transpositions.

While it is natural for D_4 to act on the vertices of a square, there are other options:

- Act on the sides of the square. This is also degree 4, and faithful.
- Act on the diagonals of the square. This is degree 2, and fails to be faithful. After factoring, the “real” group is $\mathbb{Z}_2$.

We can now formally describe patterns by having the whole group act on an element in X .

Definition

Let $\varphi : G \times X \rightarrow X$ be a left action. The **orbit** of $x \in X$ under G is

$$Gx := \{ax \mid a \in G\}.$$

One says that the elements in an orbit are **G -equivalent**.

So a pattern is simply an orbit under G .

Note that our venerable old notion of orbit obtained by iterating an endofunction $f : A \rightarrow A$ is entirely analogous: it's just the special case where we have the additive monoid $\mathbb{N}$ acting on A , rather than a group.

Proposition

Let G be a group. Then the orbits Gx form a partition of X .

Proof. $z \in Gx \cap Gy$ implies $ax = z = by$ for some $a, b \in G$. But G is a group, so $x = (a^{-1}b)y \in Gy$. □

So, our terminology makes sense: the blocks of this partition are exactly the patterns we are interested in.

Note, though, that we really need G to be a group, the argument fails for monoids. Over a monoid, all we have is a basin of attraction.

If the action is transitive, then there is only one orbit.

1 Some Counting Problems

2 Actions

3 **Burnside's Lemma**

4 Some Applications

Now comes the important idea: using subgroups and fixed points to count. Let $\varphi : G \times X \rightarrow X$ be a left group action.

Definition

The **stabilizer** of G at $x \in X$ is

$$G_x := \{ a \in G \mid a x = x \}$$

The **invariant subset** of X at $a \in G$ is

$$X_a := \{ x \in X \mid a x = x \}$$

So both definitions involve fixed points of the action, once from the perspective of the group, and once from the perspective of the G -set.

We want to determine the size of an orbit Gx . The only obvious bounds are

$$1 \leq |Gx| \leq |G|$$

The problem is that we may well have $ax = bx$ for $a \neq b$. But then

$$ax = bx \iff a^{-1}bx = x \iff a^{-1}b \in G_x \iff b \in aG_x$$

So the size of the orbit $|G|/|G_x|$.

Note that the algebraic manipulations are all justified by the definition of action.

Proposition

The stabilizers G_x are subgroups of G .

Proof.

Let $a, b \in G_x$. Then

$$(a^{-1}b)x = a^{-1}(bx) = a^{-1}(x) = a^{-1}(ax) = x$$

Hence $a^{-1}b \in G_x$ and we are done.



Proposition

The index $[G : G_x]$ is the size of the orbit of x .

Proof. As already pointed out

$$a x = b x \iff b \in a G_x$$

Hence $|G_x|$ many elements in G produce the same element in the orbit.

So $|Gx| = [G : G_x]$.

□

Recall that by Lagrange's theorem, $[G : H] = |G|/|H|$ is integral for finite groups.

So we can write the partition of G into blocks as

$$G/G_x = \{g_1 G_x, g_2 G_x, \dots, g_r G_x\}$$

where $r = [G : G_x]$ then the orbit has the form

$$Gx = \{g_1 x, g_2 x, \dots, g_r x\}$$

Hence, if we know representatives for the cosets, then we can enumerate the orbit of x directly without repetitions.

The bad news: We can always choose $g_1 = 1$, but other than that it may not be so easy to get at the g_i (the problem of finding a complete set of representatives). This is a standard problem in computational group theory.

Double counting is a very simple but sometimes surprisingly powerful idea.

Suppose R (for rows) and C (for columns) are two finite sets and M is an R by C matrix with 0/1 entries.

Let $\text{row}(r)$ be the number of 1's in row $r \in R$, and let $\text{col}(c)$ be the number of 1's in columns $c \in C$.

Then

$$\sum_{r \in R} \text{row}(r) = \sum_{c \in C} \text{col}(c).$$

Yes, yes, that's trivial. But ...

Lemma

$$\sum_{a \in G} |X_a| = \sum_{x \in X} |G_x|.$$

Proof.

Consider the **action matrix**: a G by X matrix M defined by


$$M(a, x) = \begin{cases} 1 & \text{if } ax = x, \\ 0 & \text{otherwise.} \end{cases}$$

The rows are bitvectors for the invariant sets, and the columns are bitvectors for the stabilizers.



For a positive integer n let $d(n)$ be the number of divisors of n .

$d(n)$ is fairly complicated.




How about the average

$$\widehat{d}(n) := 1/n \sum_{x \leq n} d(x).$$

Just as hopeless? Even more hopeless?

Let M be the n by n binary matrix with 1 in position (x, y) iff x divides y .



The number of 1's in column y is just $d(y)$, and difficult.

But the number of 1's in row x is simply $\lfloor n/x \rfloor$.

So the total number of 1's is

$$\sum_{x \leq n} \text{row}(x) = \sum_{x \leq n} \lfloor n/x \rfloor \leq \sum_{x \leq n} n/x = n \cdot H_n.$$

and the error is at most n .

Hence $\widehat{d}(n)$ is about $\log n$.

Theorem

Let N be the number of distinct orbits of G acting on X . Then

$$N = \frac{1}{|G|} \sum_{a \in G} |X_a|.$$

Proof.

$$\frac{1}{|G|} \sum_{a \in G} |X_a| = \frac{1}{|G|} \sum_{x \in X} |G_x| = \sum_{x \in X} \frac{1}{[G : G_x]} = \sum_{x \in X} \frac{1}{|G_x|} = N.$$

□

So the number of orbits is the average of the number of fixed points.

Burnside published this lemma in 1900.

Unfortunately, Frobenius already published the same result in 1887: *Über die Congruenz nach einem aus zwei endlichen Gruppen gebildeten Doppelmodul*.

And Frobenius was scooped by Cauchy in 1835: *Mémoire sur diverses propriétés remarquables des substitutions régulières ou irrégulières, et des systèmes de substitutiones conjuguées*.

Note how papers used to have long, descriptive names.

And, a searchable web really is a blessing (of course, this assumes proper semantic markup, currently a pipedream).

In practice, this means that one has to

- Determine the group of actions G .
- Compute the (sizes of the) invariant sets X_a for all group elements a .

Usually G is clear from the given atomic actions, but sometimes even this step requires a bit of work.

Counting fixed points can be problematic when the group is large, or when the action is very complicated. In the 21st century, one can use computer algebra to take care of the more painful computations.

1 Some Counting Problems

2 Actions

3 Burnside's Lemma

4 **Some Applications**

Here is a trivial counting problem, but it provides an opportunity to use the new machinery.

We consider binary lists (bitvectors) of length n .

We want to identify two lists when one is obtained from the other by flipping all bits.

To apply Burnside let

$$X = 2^n$$

$$G = \{1, s\}$$

where s means flip all bits. Hence $s^2 = 1$ and G really is a group.

In fact, G is isomorphic to $\mathbb{Z}_2$ but let's use the multiplicative notation.

We need to calculate the invariant sets, which is easy in this case.

$$X_1 = X$$

$$X_s = \emptyset$$

Hence

$$N = \frac{1}{2}(2^n + 0) = 2^{n-1}$$

So each orbit has the form $\{x, sx\}$ and has size 2.

Not very exciting, but at least it's correct.

In the context of cellular automata one encounters the following problem: we have binary lists of length $n = 2^k$ encoding the local functions.

Two lists L and K are equivalent if K can be obtained from L by any combination of reversing the list, or flipping all bits.

To apply Burnside let

$$X = 2^n$$

$$G = \langle r, s \mid r^2 = s^2 = 1, rs = sr \rangle = \{1, r, s, rs\}$$

where r means reversal, s means flip all bits. Note that G is isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2$, the Kleinsche Vierergruppe.

Since n is even we have

$$X_1 = X$$

$$X_r = \{ uu^{\text{op}} \mid u \in \mathbf{2}^{n/2} \}$$

$$X_s = \emptyset$$

$$X_{rs} = \{ u \bar{u}^{\text{op}} \mid u \in \mathbf{2}^{n/2} \}$$


Hence

$$N = \frac{1}{4}(2^n + 2^{n/2} + 0 + 2^{n/2}) = 2^{n-2} + 2^{n/2-1}$$

Color the corners of a square red and blue. Obviously there are 2^4 colorings (our configurations).

Now suppose we do not wish to distinguish between colorings that can be obtained from each other by rotations and reflections (our patterns).

For 2 colors and 4 vertices we can easily compute this to death, but think about the analogous problem with c colors and n vertices.



So there are 6 patterns.

Let X be the set of all colorings. It is convenient to think of X as 2^4 (read off the colors in clockwise order).





The patterns are exactly the orbits of $x \in X$ under D_4 .

Hence

$$N = \frac{1}{8} \sum_{a \in D_4} |X_a|.$$

So what are the invariant sets?

Recall $D_4 = \{1, \alpha, \alpha^2, \alpha^3, \beta, \alpha\beta, \alpha^2\beta, \alpha^3\beta\}$



$$X_1 = X$$

$$X_\alpha = X_{\alpha^3} = \{0000, 1111\}$$

$$X_{\alpha^2} = \{0000, 0101, 1010, 1111\}$$

$$X_\beta = \{0000, 0011, 1100, 1111\}$$

$$X_{\alpha\beta} = \{ijk \mid i, j, k \in \mathbf{2}\}$$


$$X_{\alpha^2\beta} = \{0000, 0101, 1010, 1111\}$$

$$X_{\alpha^3\beta} = \{jiki \mid i, j, k \in \mathbf{2}\}$$

Hence

$$N = \frac{1}{8}(16 + 2 + 4 + 2 + 4 + 8 + 4 + 8) = 6$$

Let's return to the historical root: counting chemical compounds. For example, suppose we want to enumerate carbocycles like benzene, where 2 H atoms have been replaced by OH groups.



Clearly we can model this again using something like a dihedral group.

A k -ary bracelet is a circular string of beads in k different colors: do not distinguish between variants obtained by rotation or reflection.

It is customary to represent each equivalence class by its lexicographically first element.

Example: all 21 ternary bracelets of length 4.

(1, 1, 1, 1)	(1, 1, 1, 2)	(1, 1, 1, 3)	(1, 1, 2, 2)
(1, 1, 2, 3)	(1, 1, 3, 3)	(1, 2, 1, 2)	(1, 2, 1, 3)
(1, 2, 2, 2)	(1, 2, 2, 3)	(1, 2, 3, 2)	(1, 2, 3, 3)
(1, 3, 1, 3)	(1, 3, 2, 3)	(1, 3, 3, 3)	(2, 2, 2, 2)
(2, 2, 2, 3)	(2, 2, 3, 3)	(2, 3, 2, 3)	(2, 3, 3, 3)
(3, 3, 3, 3)			

To apply Burnside note that the group acting on X is the same as for a regular n -gon, so we can use

$$\begin{aligned}X &= [k]^n \\ G &= D_n\end{aligned}$$

What are the invariant sets?

We need the cardinality of

$$\begin{aligned}X_{\alpha^s} &= \{x \in X \mid \alpha^s x = x\} \\ X_{\alpha^s \beta} &= \{x \in X \mid \alpha^s \beta x = x\}\end{aligned}$$

For the rotations this means $x_0 = x_s = x_{2s} = \dots$ and things wrap around modulo n .

Remember the circulant function



$$\begin{array}{ccc} f_s : & \mathbb{Z}_n & \longrightarrow \mathbb{Z}_n \\ & z & \longmapsto z + s \bmod n \end{array}$$

We are walking around a circle of length n using stride s . f_s has $\gcd(n, s)$ distinct orbits, each of length $n / \gcd(n, s)$.

An action fixed point $\alpha^s x = x$ means that the list elements on each orbit of f_s are the same.

Hence there are $k^{\gcd(n, s)}$ invariant lists for the rotation α^s .

How about the reflections $\alpha^s \beta$? A few pictures help a lot in this case.



The pictures are for $n = 12$ and $s = 2, 3$.

It looks like there 2-cycles and possibly fixed points, nothing else.

Remember that the motions in D_n are either rotations or reflections with respect to a properly chosen axis, nothing else can happen.

As a consequence, for even n , there are either $n/2$ many 2-cycles or $(n/2 - 1)$ many 2-cycles and two fixed points: depending on whether the axis of the reflection passes through vertices or the center of the sides of the n -gon.

For odd n , there are always $(n - 1)/2$ many 2-cycles plus one fixed point: the reflection axis always has to pass through a vertex and the center of one side.

Exercise

Draw pictures to confirm these assertions.

For simplicity, assume n is odd. Then the number of k -ary necklaces of length n is

$$\begin{aligned} \frac{1}{2n} \left(\sum_{s < n} |X_{\alpha^s}| + \sum_{s < n} |X_{\alpha^s \beta}| \right) &= \\ \frac{1}{2n} \left(\sum_{s < n} k^{\gcd(n,s)} + \sum_{s < n} k^{(n+1)/2} \right) &= \\ \frac{1}{2n} \sum_{d|n} \varphi(n/d) k^d + k^{(n+1)/2}/2 \end{aligned}$$

where φ is Euler's totient function: $\varphi(m) = |\mathbb{Z}_m^*|$.

For even n the counting result is very similar.

$$\frac{1}{2n} \sum_{d|n} \varphi(n/d) k^d + (k+1)k^{(n+1)/2}/4$$

Not too pretty, but no nice simple closed form exists. Enough, though, to compute some values. $k = 2, \dots, 6$ and $n = 1, \dots, 8$:

2	3	4	6	8	13	18	30
3	6	10	21	39	92	198	498
4	10	20	55	136	430	1300	4435
5	15	35	120	377	1505	5895	25395
6	21	56	231	888	4291	20646	107331

For our original problem, there are 13 possible molecules obtained from substituting some H atoms by OH groups.

Let's check.

(1, 1, 1, 1, 1, 1)		
(1, 1, 1, 1, 1, 2)		
(1, 1, 1, 1, 2, 2)	(1, 1, 1, 2, 1, 2)	(1, 1, 2, 1, 1, 2)
(1, 1, 1, 2, 2, 2)	(1, 1, 2, 1, 2, 2)	(1, 2, 1, 2, 1, 2)
(1, 1, 2, 2, 2, 2)	(1, 2, 1, 2, 2, 2)	(1, 2, 2, 1, 2, 2)
(1, 2, 2, 2, 2, 2)		
(2, 2, 2, 2, 2, 2)		

Can we answer the old Tic-Tac-Toe question at this point? We need to count the number of patterns of type $(3, 3, 3)$.

So we have

- configuration space X consisting of all $(3, 3, 3)$ placements,
- dihedral group D_4 acting on X .

More formally, $X \subseteq \{0, 1, 2\}^{3,3}$ is determined by the condition that the number of 0's, 1's and 2's in a matrix is exactly 3 each.

Dire Warning: The group here is not D_4 , but an isomorphic subgroup of $\mathfrak{S}_9$: the board has 9 squares. We need to boost D_4 to a permutation group of degree 9.

Let us number the squares in row-major order:

1	2	3
4	5	6
7	8	9

Then clockwise rotation and reflection along the horizontal axis in cycle notation correspond to the two permutations


$$((1, 3, 9, 7), (2, 6, 8, 4))$$

$$((1, 7), (2, 8), (3, 9))$$

These two duly generate a subgroup of $\mathfrak{S}_9$ isomorphic to D_4 .

Next we need to determine the cardinalities of the invariant sets X_a for all 8 group elements.

We claim that, for rotations other than the identity, all invariant sets are empty. We only consider rotation by $\pi/2$, the other cases are entirely similar.






To see that $X_\alpha = \emptyset$ note that since we only have 3 marks of each kind no 4-cycle can be invariant.

Note that this argument is a bit frail: if we were to look at different kinds of configurations we would need to start all over again – more on this later.

Exercise

Carry out the same argument for the invariant subset of rotation α^2 . Also argue that $X_{\alpha^2} = \emptyset$ implies that $X_\alpha = \emptyset$.

For reflections things are more interesting. We only consider reflection along the vertical axis, the other cases are entirely similar: the invariant set has size $6 \times 6 = 36$.



Hence the number of $(3, 3, 3)$ Tic-Tac-Toe configurations is

$$1/8 (1680 + 36 + 36 + 36 + 36) = 228$$

Not bad, but as already mentioned, this method becomes tedious if we ask about other types of boards. E.g., we have no idea how many patterns there are for 2 crosses and 2 naughts. It would be nice to have a global tool to handle all possible cases (a, b, c) where $a + b + c = 9$.

Exercise

Determine the largest invariant set (other than X_1 of course) for all possible configurations.