CDM

Group Actions and Counting

Klaus Sutner

Carnegie Mellon University

Spring 2021

1 **Some Counting Problems**

2 **Actions**

3 **Burnside's Lemma**

Counting Patterns

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.

Classical Killer App: Chemistry

Suppose we want to enumerate hydrocarbon molecules where some hydrogen atoms have been replaced by hydroxyl groups.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.

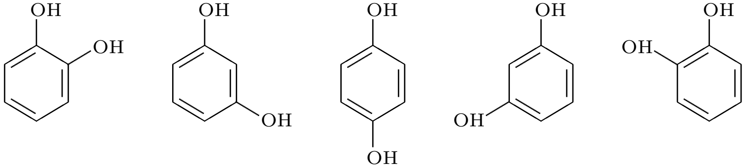

Some counting: $\binom{6}{2} = 15$, $\binom{5}{1} = 5$ or, taking into account symmetry, 3.

Only the last answer makes chemical sense.

Pólya-Redfield Counting

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

Problem 1: Tic-Tac-Toe

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?

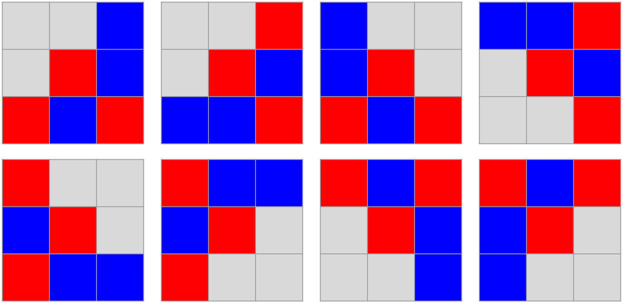

Eyeballing It

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.

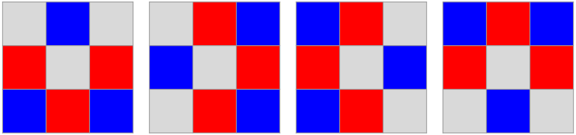

But that is not quite right, either: some placements have fewer than 8 variants, so we are under-counting the number of patterns.

Problem 2: Boolean Circuits

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

$$f(\boldsymbol{x}) = g(\pi(\boldsymbol{x}))$$

for some permutation π.

How many Boolean functions are there modulo input permutations?

Another Version

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

$$f(\boldsymbol{x}) = g(\boldsymbol{x} \oplus \boldsymbol{c}) \oplus d$$

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

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

And More . . .

We can combine different kinds of modifications.

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

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

So how many functions $\mathbf{2}^n \to \mathbf{2}$ 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

Burnside, Pólya and Redfield

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.

Central Idea

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

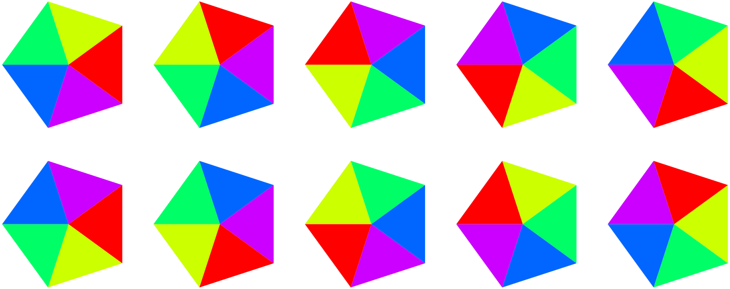

Before we start . . .

> 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, . . .)

Déjà Vu, All Over Again

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 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 $S^\star$, so we are dealing with monoid actions.

nil novis sub solem

Free monoids $\Sigma^\star$ 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 $\Sigma^\star$ 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.

Reversibility

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.

Left Actions

Definition

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

$$\begin{aligned}&\varphi : G \times X \to X\\&\varphi(a * b, x) = \varphi(a, \varphi(b, x))\\&\varphi(1, x) = x\end{aligned}$$

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 $a\,x$ instead of $\varphi(a, x)$. Hence

$$\begin{aligned}&(a * b) \cdot x = a \cdot (b \cdot x)\\&1 \cdot x = x\end{aligned}$$

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 $G = \mathsf{GL}(2, \mathbb{R})$ 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 $A \cdot x = Ax$.

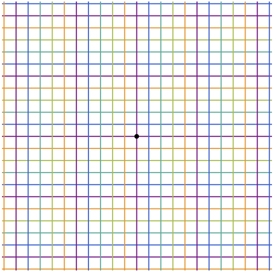

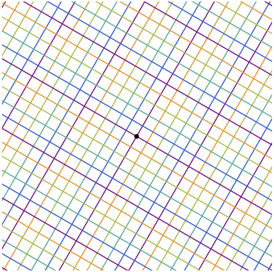

The action induced by the rotation matrix $A = 1/2 \begin{pmatrix} 1 & -\sqrt{3} \\ \sqrt{3} & 1 \end{pmatrix}$.

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.

Moving Things Around

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

Consider a list of n objects:

$$\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$$

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

$$\boldsymbol{x}' = (x_{f(1)}, x_{f(2)}, \ldots, x_{f(n)})$$

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 rations and/or reflections, we can translate it into a permutation group.

Actions and Homomorphisms

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

Then

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

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.

Check the Axioms

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

Exercise

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

The Standard Example

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

$$f \cdot \boldsymbol{x} = (x_{f(1)}, x_{f(2)}, \ldots, x_{f(n)})$$

Proof

It is clear that $1 \cdot \boldsymbol{x} = \boldsymbol{x}$.

Consider

$$(f \circ g) \cdot \boldsymbol{x} = \boldsymbol{y}_1$$

versus

$$f \cdot (g \cdot \boldsymbol{x}) = \boldsymbol{y}_2.$$

We need to show that $\boldsymbol{y}_1 = \boldsymbol{y}_2$.

Recall that we compose functions from left to right, so that
$\boldsymbol{y}_1 = (x_{g(f(1))}, \ldots, x_{g(f(n))})$.
But then $\boldsymbol{y}_2 = f \cdot (x_{g(1)}, \ldots, x_{g(n)}) = \boldsymbol{y}_1$.

□

Right?

Stop! Mental Health Alert

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.

$$
\begin{aligned}
g \cdot \boldsymbol{x} &= \big(x_{g(1)}, x_{g(2)}, \ldots, x_{g(n)}\big) \\
&= \big(u_1, u_2, \ldots, u_n\big) \\[1ex]
f \cdot \boldsymbol{u} &= \big(u_{f(1)}, u_{f(2)}, \ldots, u_{f(n)}\big) \\
&= \big(x_{g(f(1))}, x_{g(f(2))}, \ldots, x_{g(f(n))}\big)
\end{aligned}
$$

Exercise

Make sure you really understand the proof.
What would happen if we did composition the other way around?

And Right Actions

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

$$\begin{aligned}&\varphi : X \times G \to X\\&\varphi(x, a * b) = \varphi(\varphi(x, a), b)\\&\varphi(x, 1) = x\end{aligned}$$

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

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

Again: Actions and Homomorphisms

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

This time we define

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

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.

Yet Another 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 $\Sigma^\star$ on the collection of all languages.

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

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

Exercises

Exercise

Check carefully that the last two examples really are right actions of the monoid $\Sigma^\star$. Find better notation for the quotient operation.

Exercise

Find some natural examples for left monoid actions.

Exercise

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

Left versus Right

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

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

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

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.

Left is Right

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

Proposition

Consider two maps $\varphi : G \times X \to X$ *and* $\psi : X \times G \to X$ *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. □

Right is Left

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

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

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

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

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

Exercise

Give a detailed proof of this claim.

Mathematics versus Implementation

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.

☠ ☠ ☠ ☠ ☠

Faithful Actions

Recall the homomorphism $\Phi : G \to \mathfrak{S}(X)$ 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.

More Properties of Actions

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.

Groups Acting on Groups

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 $a \cdot H = aHa^{-1} = \{\, aha^{-1} \mid h \in H \,\}$.

Flipping Bits

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

$$\begin{aligned} \alpha(0x) &= 1\,\gamma(x) \\ \alpha(1x) &= 0\,\beta(x) \\ \beta(sx) &= s\,\alpha(x) \\ \gamma(sx) &= s\,\beta(x) \end{aligned}$$

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 $\mathbf{2}^\star$.

Question: Is this action faithful?

Identities

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

$$\alpha\beta = \beta\alpha \qquad \alpha\gamma = \gamma\alpha \qquad \beta\gamma = \gamma\beta$$
$$\alpha^2\beta^2\gamma = I$$

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 $\mathbf{2}^\star$:

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

Exercise

Prove all these claims.

Dihedral Groups

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

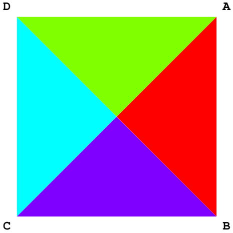

Abstractly, D_4 has the presentation

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

D_4

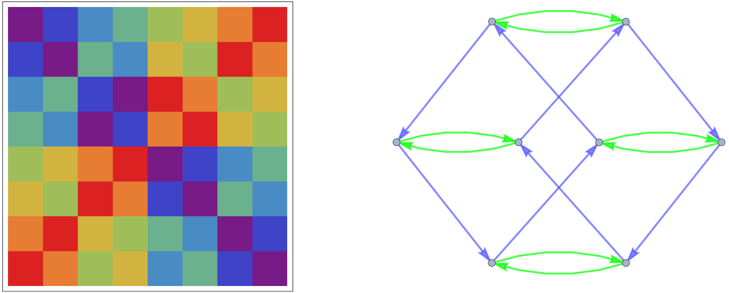

The Cayley table and Cayley graph for D_4.

Matrix Form

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.

Embedding into $\mathfrak{S}_4$

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."

Cycle Decomposition

Recall that any permutation can be decomposed into disjoint cycles:

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

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) $q_1, \ldots, q_p$.

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

D_4 Cycle Shapes

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).

$\mathfrak{S}_5$ Cycle Shapes

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.

$\mathfrak{S}_6$ Cycle Shapes

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.

$\mathbb{A}_6$ Cycle Shapes

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.

Degree

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 $\mathbb{Z}_2$.

Patterns and Orbits

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

Definition

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

$$G\,x := \{\, a\,x \mid a \in G \,\}.$$

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 $f : A \to A$ is entirely analogous: it's just the special case where we have the additive monoid $\mathbb{N}$ acting on A, rather than a group.

Chopping Things Up

Proposition

Let G be a group. Then the orbits $G\,x$ form a partition of X.

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

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**

Stabilizers and Invariants

Now comes the important idea: using subgroups and fixed points to count. Let $\varphi : G \times X \to X$ 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.

How to Count

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

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

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

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

Stabilizers

Proposition

The stabilizers G_x are subgroups of G.

Proof.

Let $a, b \in G_x$. Then

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

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

Orbit Sizes

Proposition

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

Proof. As already pointed out

$$a\,x = b\,x \iff b \in aG_x$$

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

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

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

$$G\,x = \{g_1\,x, g_2\,x, \ldots, g_r\,x\}$$

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.

Digression: Double Counting

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

Then

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

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

Application

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

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

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

□

Sub-Digression: Number of Divisors

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

$d(n)$ is fairly complicated.

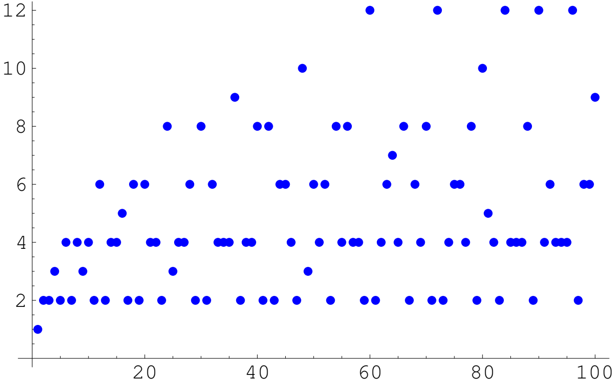

Average Number of Divisors

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.

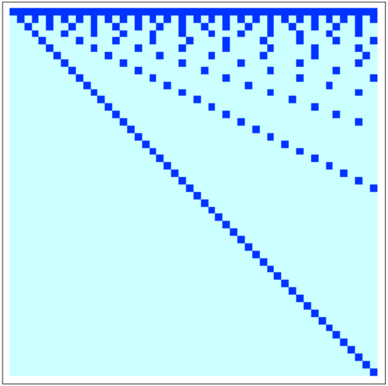

The number of 1's in column y is just $d(y)$, and difficult.

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

So the total number of 1's is

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

and the error is at most n.

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

Burnside's Lemma

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.

Full Disclosure

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 systémes de substitutiones conjugées.*

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).

Computing Burnside

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.