CDM

Semigroups and Groups

Klaus Sutner

Carnegie Mellon University
Spring 2023

Recall Bourbaki's central idea: organize math into structures and study their properties and relationships.

As we now know, this idea works particularly well in algebra.

The axiomatic/structural style was originally pioneered by D. Hilbert;
specifically for algebra the major breakthrough publication is van der Waerden's
(1903-1996) classical texts that cemented the notion of algebraic structure (a
first-order structure) as the fundamental concept in algebra:

> B. L. van der Waerden
> Moderne Algebra, Teil I
> Springer Verlag, Berlin, 1930

> B. L. van der Waerden
> Moderne Algebra, Teil II
> Springer Verlag, Berlin, 1931

It is completely clear to me which conditions caused the gradual decadence of mathematics, from its high level some 100 years ago, down to the present hopeless nadir. Degeneration of mathematics begins with the ideas of Riemann, Dedekind and Cantor which progressively repressed the reliable genius of Euler, Lagrange and Gauss. Through the influence of textbooks like those of Hasse, Schreier and van der Waerden, the new generation was seriously harmed, and the work of Bourbaki finally dealt the fatal blow.

C. L. Siegel, letter to A. Weil, 1959.

This is kind of funny, since Bourbaki was in part an attempt to create a counterweight to Hilbert's school.

Meanwhile, I was in a mathematics department, and this style of mathematics was not at all in fashion. Bourbaki was king: The more abstract you could be, expressing everything in terms of morphisms and categories, the better. Highly abstract methods were in favor in all the best mathematical schools. In more and more of the lectures that I was hearing at Caltech, I would find myself sitting in the audience saying to myself, "So what? So what?"

Eventually I switched fields and became a professor of computer science.

D. E. Knuth, 2014

To be clear: Bourbaki won, hands down.

In a sense, that's just too bad: as far as Bourbaki is concerned, there is no computational universe, just some weird, set-theory based, logic-deprived, picture-less, non-applicable, entirely un-algorithmic wasteland.

I think this will change in the foreseeable future, mostly thanks to those pesky CS dudes and all the computing hardware/software everywhere, but right now we all still live in Bourbaki's paradise (actually, hell).

So, we will follow the axiomatic approach very closely, with a little computation added in.

The following general ideas are crucial when dealing with algebraic structures:

Substructure A structure that is obtained by shrinking the carrier set and the algebraic operations.

Morphism A map from one structure to another that preserves the relevant operations and relations. Aka homomorphism.

Quotient A structure obtained by identifying some of the elements of a given structure (via a congruence, an equivalence relation that is compatible with the algebraic operations).

Product A structure that is defined over the Cartesian product of other structures, with appropriately defined operations.

Definition

A magma is a structure with a single binary operation $*$:

$$\mathcal{G} = \langle G, * \rangle$$

where $* : G \times G \to G$.

The operation is often referred to as multiplication.

In a magma, there are no further restrictions on the operation. This is really the logic perspective, a first-order structure with a single binary function. Algebraists are typically not much interested in this level of generality, they want additional constraints on the operation (such as associativity).

The terms over a magma can be construed as full binary trees: the interior nodes correspond to "multiplications" and the leaves are variables or constants.
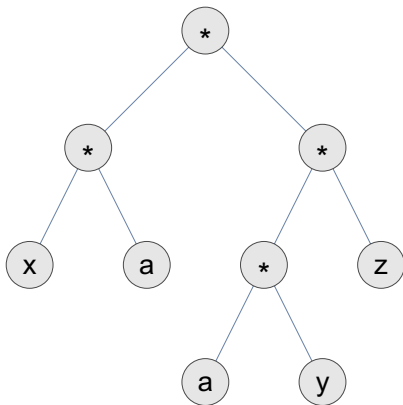
Unfortunately, magmas are also referred to as "groupoids."

This is a bad idea, since groupoids are more generally defined as structures $\langle A, *, {}^{-1} \rangle$ that are, roughly speaking, groups with a partial multiplication. More precisely, we have

- $*$ is a partial binary operation,
- ${}^{-1}$ is a total unary operation,
- subject to the following laws:
  - Partial associativity: if all terms are defined, then $a * (b * c) = (a * b) * c$.
  - Inverse: for all a, $a * a^{-1}$ and $a^{-1} * a$ are defined.
  - Identity: if $a * b$ is defined, so is $a^{-1} * a * b = b$ and $a * b * b^{-1} = a$.

We won't discuss groupoids here.

Example

The natural numbers with exponentiation form a magma.

Example

The integers with subtraction form a magma.

Example

The positive rationals with division form a magma.

Example

Rooted binary trees can be considered as magma

$$\langle \mathcal{T}, * \rangle$$

where $\mathcal{T}$ is the collection of all rooted binary trees and $*$ denotes the operation of attaching two trees to a new root. Note that this operation is highly non-associative:

$$r * (s * t) \neq (r * s) * t$$

no matter what r, s and t are (at least in the finite case).

Example

Likewise we could consider lists over some groundset A as a magma

$$\langle \mathsf{List}(A), * \rangle$$

where $*$ is interpreted as join (concatenation). This operation is associative.

It is often helpful to translate combinatorial structures into algebraic ones. For example, suppose $\mathcal{G} = \langle V, E \rangle$ is a digraph. We can translate the graph into a magma

$$\mathcal{A}(\mathcal{G}) = \langle V_\perp, * \rangle$$

by setting $V_\perp = V \cup \{\perp\}$ where $\perp \notin V$ is a new point and

$$u * v = \begin{cases} u & \text{if } (u, v) \in E, \\ \perp & \text{otherwise.} \end{cases}$$

This operation is not associative in general.

Exercise

Figure out what left (or right) parenthesized products mean in $\mathcal{A}(\mathcal{G})$. Is such a graph algebra commutative?

Magmas are also (mildly) helpful when dealing with more complicated structures: it is always a good idea to try to understand if a result (or even a definition) also works over magmas or whether it really requires the additional assumptions.

All the fundamental notion of sub-structure, congruence, homomorphism and so on already make sense for magmas and are perhaps a bit easier to understand there since there are no other properties lying around that can obscure the view.

Exercise

Rewrite all the definitions below in the context of magmas.

In algebra, it is interesting to understand a structure that satisfies certain specifications, but has no other, special properties. These structures are called free.

So suppose we have a carrier set A. What would the free magma over A look like?

Just like the term trees from above, except that there are no free variables this time, all the leaves are labeled by elements of the ground set A.

Consider a magma $\mathcal{A} = \langle A, * \rangle$.

- A substructure of $\mathcal{A}$ consists of a set $B \subseteq A$ that is closed under $*$.

- A homomorphism from $\langle A, * \rangle$ to $\langle B, \cdot \rangle$ is a map $\varphi : A \to B$ such that $\varphi(x * y) = \varphi(x) \cdot \varphi(y)$.

- A quotient of $\langle A, * \rangle$ is given by an equivalence relation ρ such that $x \, \rho \, x'$ and $y \, \rho \, y'$ implies $x * y \, \rho \, x' * y'$ (a congruence).

- The product of $\langle A, * \rangle$ and $\langle B, \cdot \rangle$ is the structure $\langle A \times B, \otimes \rangle$ where $(x, y) \otimes (x', y') = (x * x', y \cdot y')$.

Definition

A semigroup is a magma with an associative operation.

Thus, for all x, y, z in a semigroup $\mathcal{G} = \langle G, * \rangle$ we have

$$x * (y * z) = (x * y) * z.$$

Many natural algebraic operations have this property, but not all:

- Exponentiation is not associative.

- Subtraction is not associative (but the underlying addition is).

- Graph algebras are generally not associative.

Definition

An idempotent in a semigroup is an element e such that $e * e = e$.

So an idempotent is a bit weaker than an identity.

Lemma

*Let S be a finite semigroup. Then S contains an idempotent.*

Exercise

Prove the idempotent element lemma. Think Floyd.

We have seen that the free magma over A is the collection of ground terms, essentially binary trees with leaves labeled in A.

In a semigroup we have one additional specification: associativity. Hence we can identify all trees with same frontier: they correspond to the same semigroup element. Hence, we might as well think of them as a list.
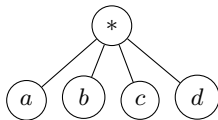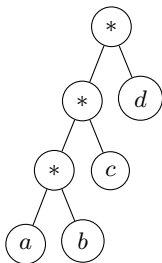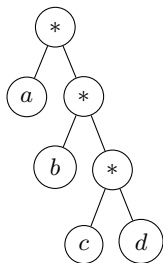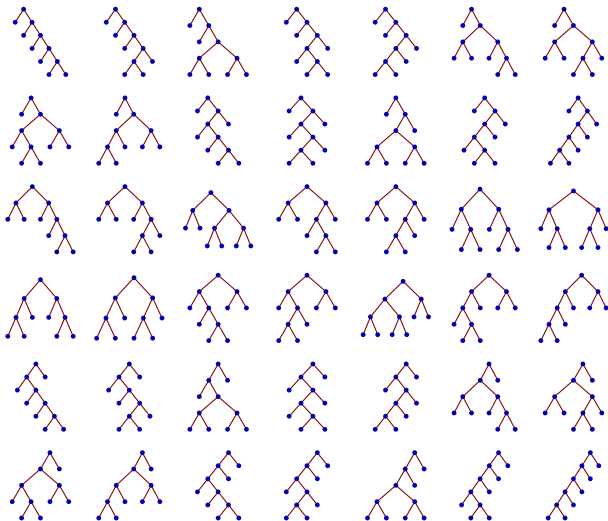
But once we think of the elements as flat lists
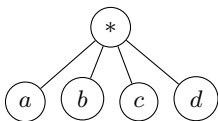
we might as well just use the sequence $abcd$, a word over the alphabet A.

So the set of non-empty words over A is the free semigroup generated by A.

Definition

A monoid is a semigroup with an identity element e: $e * x = x * e = x$.

One usually writes monoids in the form

$$\mathcal{A} = \langle A, *, e \rangle$$

to indicate the neutral element. Of course, the neutral element is idempotent.

Proposition

The neutral element in a monoid is unique.

Proof. $e = e * e' = e'$. □

Example

The set of all words over a fixed alphabet forms a monoid with concatenation as operation. The neutral element is the empty word.

Example

The set of all lists over some fixed ground set forms a monoid with join as operation. The neutral element is the empty list.

Example

The set of all functions $f : A \to A$ for some arbitrary set A forms a monoid with functional composition as operation. The neutral element is the identity function.

Example

The set of all binary relations on A, for some arbitrary ground set A, forms a monoid with relational composition as operation. The neutral element is the identity relation.

Example

The set of natural numbers with addition forms a monoid; the neutral element is 0.
Ditto for integers, rationals, algebraic numbers, reals, complex numbers.

Example

The set of positive natural numbers with multiplication forms a monoid; the neutral element is 1.

Example

The set of all n by n matrices of, say, integers, with matrix multiplications forms a monoid; the neutral element is the identity matrix.

Example

A band is a semigroup defined on the Cartesian product $A \times B$ where A and B are two arbitrary sets (non-empty). The operation is

$$(a, b) * (c, d) = (a, d)$$

It is obvious that this operation is associative. Note that a band is idempotent: $x * x = x$ for all x.

Example

The bicyclic semigroup is defined on $\mathbb{N} \times \mathbb{N}$ by the operation

$$(a, b) * (c, d) = (a - b + \max(b, c), d - c + \max(b, c))$$

Associativity requires a little argument here. This may look strange, but it is just the free semigroup on two generators r and s subject to $sr = 1$. The idempotents of this semigroup are exactly the elements (a, a).

Monoids appear quite frequently, but have one crucial flaw from the point of view of solving equations: in general we cannot solve the equation

$$a * x = b.$$

To make sure solutions always exist we need more assumptions. One step in the right direction is the (left) cancellation property:

$$a * x = a * y \qquad \text{implies} \qquad x = y$$

This guarantees that a solution, if it exist, is unique.

Exercise

Check which of the monoids from above have the cancellation property. What restrictions on the left multiplier a are necessary to guarantee cancellation?

To solve equations, one needs some kind of inverse element. There are several classes of semigroups that are closer to groups then general ones.

Definition

An element a of a semigroup is regular if there is an element b such that $aba = a$. A semigroup is regular if all its elements are.
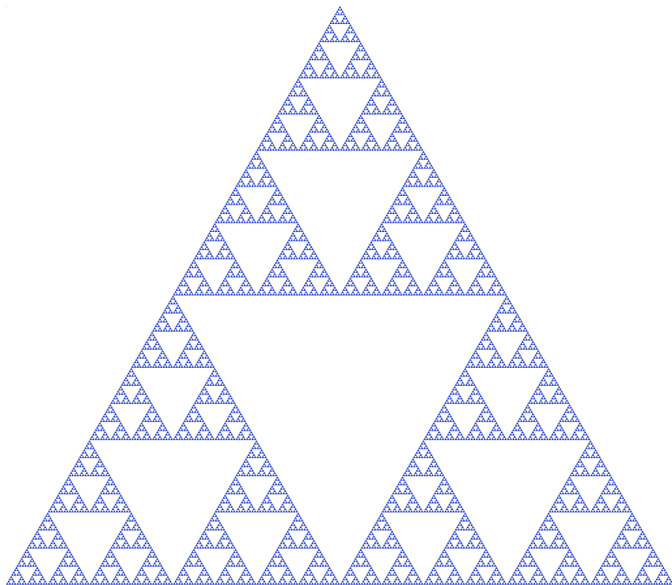
Element a is said to have a generalized inverse if there is a b such that $aba = a$ and $bab = b$. A semigroup is inverse if every element has exactly one generalized inverse.

Writing a' for the generalized inverse in an inverse semigroup we have

$$a'' = a \qquad \text{and} \qquad (ab)' = b'a'$$

Lemma

A semigroup is inverse iff it is regular and all its idempotents commute.

Felix Klein's "Erlanger Programm" of 1872 proposed to characterize geometries by studying the invariants of linear transformation groups.
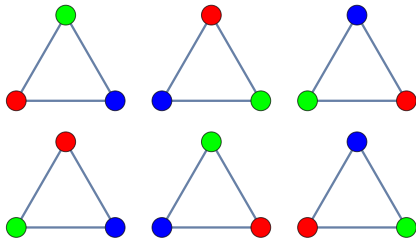
Algebra is the big problem solver.

We can rotate and reflect the Sierpinski triangle, the underlying group being the dihedral group D_3 (symmetries of an isosceles triangle).

Write a for rotation by $2\pi/3$ and b for reflection along the vertical axis. Then

$$aaa = 1 \qquad bb = 1 \qquad ba = aab$$

As a consequence, there are 6 ways to move an isosceles triangle back to itself.

But for the Sierpinski triangle that's not really the whole story, it is much more complicated than just a isosceles triangle. For example, we could shift one of the 3 big component triangles to another one.

Or we could map the whole figure into the upper big triangle; alternatively, the lower left or right triangle.

Or into one of the smaller triangles that appear in various places in the figure.

Or one of the smaller triangles into yet another one.

And so on, and so forth . . .

One reason inverse semigroups are very important is that they generalize symmetry groups.

For any set X consider the symmetric monoid

$$\mathcal{I}(X) = \{\, f : X \rightarrowtail X \mid f \text{ partial, injective} \,\}$$

Then $\mathcal{I}(X)$ is an inverse semigroup.

For the Sierpinski triangle T the symmetric monoid of T has many maps that are absent in the corresponding group.

Exercise

Explain why $\mathcal{I}(X)$ fails to be a group (for X non-empty).

At last, here is the kind of structure that guarantees existence and uniqueness of solutions of linear equations.

Definition

A group is a monoid $\mathcal{G} = \langle G, \cdot, e \rangle$ where

$$\forall x \, \exists y \, (x \cdot y = y \cdot x = e)$$

The y above is uniquely determined by x and called the inverse of x.

Since x uniquely determines the inverse one usually switches the signature to (2,1,0) and writes

$$\mathcal{G} = \langle G, \cdot, ^{-1}, e \rangle$$

This has the major advantage that the axioms are then equational:

$$x \cdot x^{-1} = x^{-1} \cdot x = e$$

Definition

A group is commutative or Abelian if $x \cdot y = y \cdot x$ for all x and y.

Notation:

It is customary to write Abelian groups additively as $\langle G, +, -, 0 \rangle$, $\langle G, +, 0 \rangle$ or $\langle G, + \rangle$.

General groups are written multiplicatively as $\langle G, \cdot, ^{-1}, 1 \rangle$, $\langle G, \cdot, 1 \rangle$ or $\langle G, \cdot \rangle$. It is standard to omit the multiplication operator (and use concatenation instead).

xy is easier on the eye than $x \cdot y$, but this gets tricky when dealing with multiple operations.

Example

The set of integers (rationals, reals, complexes) with addition forms a group; the neutral element is 0.

Example

The set of modular numbers relatively prime to modulus m with multiplication forms a group; the neutral element is 1.

Example

The set of non-zero rationals (reals, complexes) with multiplication forms a group; the neutral element is 1.

Example

The set of all regular n by n matrices of reals, with matrix multiplications forms a group; the neutral element is the identity matrix.

Example

The set of all permutations $f : A \to A$ for some arbitrary set A forms a group with functional composition as operation. The neutral element is the identity function.

In a group, the equation

$$a \cdot x = b$$

always has the unique solution

$$x = a^{-1} \cdot b$$

Note that this is easier than standard arithmetic: there is no need to worry about the case $a = 0$.

The systematic study of abstract groups was one of the central accomplishments of 19th century mathematics.

They appear in many, many places and some understanding of the their basic properties is crucial.

Laziness: any property, of, say, groups derived from only the axioms holds in all groups, automatically. You only check three simple properties, and all results apply.

Psychology: it is sometimes easier to argue abstractly than in a concrete situation (you can't see the forest because of all the trees).

This a bit hard to believe, but true. E.g., consider non-singular matrices of reals. Show that

$$(A \cdot B)^{-1} = B^{-1} \cdot A^{-1}$$

One could try to use the properties of matrix multiplication and, say, Gaussian elimination, to prove this. Any such argument would be very hard and technically difficult.

Exercise
Give a simple, abstract proof of this equation in any group whatsoever.

Aka right angular Artinian groups, don't ask. Here graphs are used as a
convenient tool to describe certain groups. Consider a ugraph $G = \langle V, E \rangle$. We
can translate G into a RAAG as follows

$$\mathcal{R}(G) = \langle V \mid uv = 1 \text{ for } u{-}v \in E \rangle$$

Thus V is the set of generators and two generators commute iff they form an
edge. Here some examples:
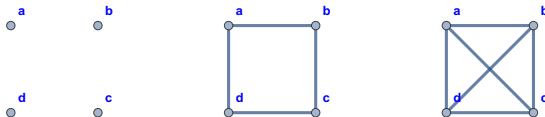
The first corresponds to the free group on 4 generators, the last to the free
Abelian group on 4 generators. The one in the middle is $F(a, c) \times F(b, d)$.

For our purposes the most important examples of groups are those comprised of permutations.

Definition

A permutation is a bijection $f : A \to A$, in particular when A is a finite set. The collection of all permutations on A, an n-element set, under functional composition is the symmetric group (on n letters or points).

Notation: $\mathfrak{S}_n$

As we will see shortly, in most cases the full symmetric group is too large; we need to focus on subgroups of $\mathfrak{S}_n$.

We will focus on the carrier set $A = [n]$. In this case, we can represent f by a $2 \times n$ matrix of the form

$$\begin{pmatrix} 1 & 2 & 3 & \ldots & n-1 & n \\ f(1) & f(2) & f(3) & \ldots & f(n-1) & f(n) \end{pmatrix}$$

This is the so-called two-line representation of f. Needless to say, the first row in this matrix is really redundant, but this redundancy makes it a bit easier to read off specific values. Alternatively, we can use one-line representation:

$$[f(1), f(2), \ldots, f(n-1), f(n)]$$

Note that we have chosen to write $[a_1, \ldots, a_n]$ to specifically indicate a map from $[n] \to [n]$. This is slightly less dangerous than just writing $(a_1, \ldots, a_n)$, which could mean a great many things.

Suppose f is a permutation. The functional digraph of f is particularly simple: it consists only of cycles (all transients are 0).

Example

Here is a permutation on $n = 12$ in two-line notation.

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \\ 10 & 12 & 8 & 3 & 7 & 9 & 5 & 6 & 11 & 2 & 4 & 1 \end{pmatrix}$$

It produces the cycles

$$(1, 10, 2, 12), (3, 8, 6, 9, 11, 4), (5, 7)$$

There are 3 cycles of lengths $4, 6, 2$, respectively.

Note that there could be just a single cycle of length n: $(1, 2, \ldots, n)$ in cycle notation stands for the cyclic shift.

It is customary (and often very useful) to omit fixed points from the list of cycles.

Example

The cycle decomposition of

$$[4, 7, 1, 6, 8, 9, 11, 5, 2, 10, 3, 12, 14, 13]$$

would be written as

$$(1, 4, 6, 9, 2, 7, 11, 3), (5, 8), (13, 14),$$

leaving out the fixed points 10 and 12.

In standard mathematics texts you should expect to find the more compact notation used a lot.

Lemma

We can compute the cycle decomposition of a permutation $f : [n] \to [n]$ in time and space linear in n.

Here we tacitly assume that $f(x)$ can be computed in time $O(1)$ (which is safe since ordinarily f will be given by an explicit array).

There are at least two ways to think about this:

- Compute the strongly connected components in the functional digraph of f.
- Compute the orbits of the function f, exploiting the fact that they are all periodic.

Note that generating the cycle decomposition seems to require linear space (as opposed to Floyd's algorithm for transients and period).

Note that we can rearrange the cycles arbitrarily, and we can rotate each individual cycle without changing the underlying permutation.

For example, the following two cycle decompositions describe the same permutation on $n = 14$.

$$(7, 5), (11, 4, 3, 8, 6, 9), (12, 1, 10, 2)$$
$$(1, 10, 2, 12), (3, 8, 6, 9, 11, 4), (5, 7)$$

The second representation may seem more natural from the implementor's point of view, but it is the first that has better combinatorial properties.

Definition

The canonical cycle decomposition (CCD) of a permutation is obtained by rotating all cycles so that the largest element is up front and the cycles are ordered by first element. If the least element is in the first position we speak of the reverse canonical cycle decomposition (RCCD)

Here is the prototype algorithm that almost everybody would write when asked to implement cycle decomposition.

```
for x = 1,..,n do
    if( x unmarked )
        mark x;
        res = (x);
        while( f(x) unmarked )
                x = f(x);
                mark x;
                append( res, x );
        output res;
```

This program places the least element first in each cycle and returns the cycles sorted by first element. In other words, the straightforward algorithm produces RCCD rather than CCD.

Here are the CCDs for all elements of $\mathfrak{S}_4$, enumerated in lex order. For clarity, we write fixed points.

$$
\begin{array}{llll}
(1),(2),(3),(4) & (1),(2),(4,3) & (1),(3,2),(4) & (1),(4,2,3) \\
(1),(4,3,2) & (1),(3),(4,2) & (2,1),(3),(4) & (2,1),(4,3) \\
(3,1,2),(4) & (4,1,2,3) & (4,3,1,2) & (3),(4,1,2) \\
(3,2,1),(4) & (4,2,1,3) & (2),(3,1),(4) & (2),(4,1,3) \\
(3,1),(4,2) & (4,1,3,2) & (4,3,2,1) & (3),(4,2,1) \\
(2),(4,3,1) & (2),(3),(4,1) & (4,2,3,1) & (3,2),(4,1)
\end{array}
$$

Exercise

Find a good algorithm to compute the CCD of a given permutation. What is the running time of your algorithm?

Here are these CCDs flattened out.

$$
\begin{array}{llll}
1,2,3,4 & 1,2,4,3 & 1,3,2,4 & 1,4,2,3 \\
1,4,3,2 & 1,3,4,2 & 2,1,3,4 & 2,1,4,3 \\
3,1,2,4 & 4,1,2,3 & 4,3,1,2 & 3,4,1,2 \\
3,2,1,4 & 4,2,1,3 & 2,3,1,4 & 2,4,1,3 \\
3,1,4,2 & 4,1,3,2 & 4,3,2,1 & 3,4,2,1 \\
2,4,3,1 & 2,3,4,1 & 4,2,3,1 & 3,2,4,1
\end{array}
$$

We get all permutations. Could this be coincidence?

From the data structure point of view, the cycle decomposition is a list of lists of integers. Hence we can flatten it to obtain a plain list of integers:

$$\text{flat} : \text{List}(\text{List}(\mathbb{N})) \rightarrow \text{List}(\mathbb{N})$$

If we start with the full cycle decomposition (including fixed points) we obtain a permutation (in one-line representation) this way. For arbitrary decompositions this is of little interest, but if we start with the CCD we get the following proposition, which is helpful in enumeration problems related to permutations.

Proposition

The map CCD ∘ flat *is a bijection on* $\mathfrak{S}_n$.

Exercise

Prove that CCD ∘ flat *is indeed a bijection.*

Exercise

What are the fixed points of this bijection?

Exercise

How about RCCD ∘ flat*?*

Since permutations are functions we can compose them by ordinary functional composition $f \circ g$. In this section, we write composition in diagrammatic form:

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

Some (misguided) texts use the opposite convention. Unfortunately, they are currently the vast majority.

Basis Problem:
Find a small and/or simple set of permutations so that all permutations can be written as a product of these.

Decomposition Problem:
Given such a basis B, find a way to decompose a given permutations into a product of permutations in B.

Definition

A transposition is a permutation that consists of a single 2-cycle.

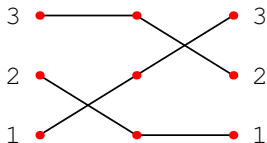In cycle notation, transpositions are exactly the permutations of the form (a, b) for $a \neq b$.

Example

Consider the following transpositions over $[3]$, given in cycle notation.
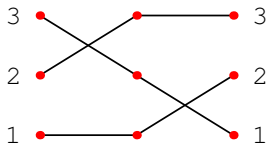
$$(1, 2) \circ (2, 3) = [3, 1, 2]$$
$$(2, 3) \circ (1, 2) = [2, 3, 1]$$

Thus, composition of permutations is not commutative (it is associative, though, since composition of functions is so associative).

$$(1,2) \circ (2,3) = [3,1,2] \qquad\qquad (2,3) \circ (1,2) = [2,3,1]$$

In cycle notation, the two composite permutations are each represented by a 3-cycle: $(1,3,2)$ and $(1,2,3)$.

Lemma

Every permutation can be written as a product of transpositions.

Proof. (sketch)

Since every permutation is composed of disjoint cycles, it suffices to show that every cycle $(a_1, \ldots, a_m)$ is a product of transpositions.

Show this by induction on $m \geq 2$. The crucial step is

$$(a_m, b) \circ (a_1, \ldots, a_m) = (a_1, a_2, \ldots, a_{m-1}, a_m, b)$$
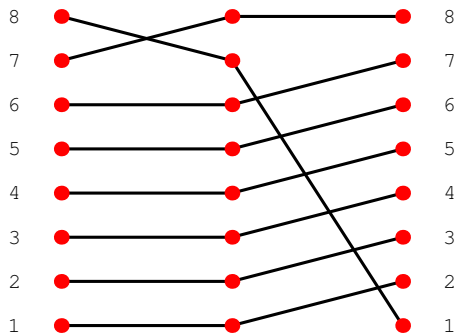
$\square$

Exercise

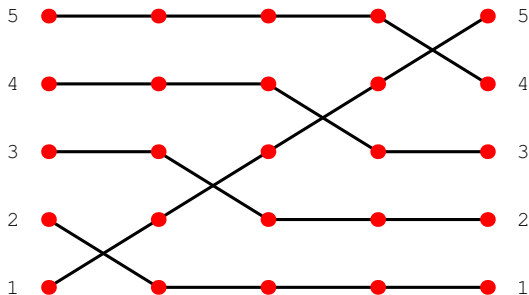*Fill in all the gaps in this argument.*

Exercise

Find a direct decomposition
$(a_1, b_1) \circ (a_2, b_2) \circ \ldots \circ (a_m, b_m) = (c_1, c_2, \ldots, c_m, c_{m+1}).$

For a simple cycle this is easy to see in the composition picture.
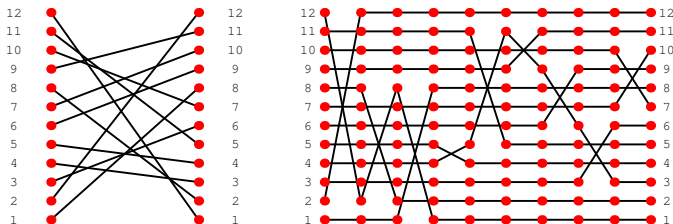
Here is another simple decomposition of a cycle intro transpositions.

So $(1,2)(2,3)(3,4)(4,5) = (5,4,3,2,1)$.

A more complicated permutation on $n = 12$, and its decomposition into transpositions.

Exercise

Find an algorithm to generate the picture on the right.

For the next proposition, we abuse notation and use exponents for permutations given in cycle notation.

Proposition

$$(a, b) \circ (b, c) \circ (a, b) = (a, c)$$
$$(1, \ldots, n)^i \circ (1, 2) \circ (n, \ldots, 1)^i = (i + 1, i + 2)$$

where $0 \leq i \leq n - 2$.

Exercise

Prove these identities.

Needless to say, the decomposition into transpositions is not unique.

Definition

A permutation is even if it can be written as the product of on even number of transpositions, and odd if it can be written as the product of on odd number of transpositions.

Note the cautious wording: this does not say that every permutation is either even or odd. It leaves open the possibility that some permutation could be both even and odd. However, one can show that any permutation is either even or odd, never both.

Lemma

No permutation is even and odd.

Let σ be a permutation of $[n]$. Consider the polynomials

$$P(x_1, \ldots, x_n) = \prod_{i < j} x_i - x_j$$

$$P_\sigma(x_1, \ldots, x_n) = \prod_{i < j} x_{\sigma(i)} - x_{\sigma(j)}$$

Then necessarily $P = \pm P_\sigma$. But then $P = +P_\sigma$ iff σ is even, and $P = -P_\sigma$ iff σ is odd. $\qquad\square$

Note that we are essentially using σ to permute the variables here.

Exercise

Fill in the details of this argument.

The composition of even permutations is again even, so we can assemble them into a new group.

Definition

The collection of all even permutations of A, an n-element set, is the
alternating group on n points.
Notation: $\mathfrak{A}_n \subseteq \mathfrak{S}_n$.

Example

In one-line notation, $\mathfrak{A}_4$ has the following elements:

$$1,2,3,4 \quad 1,3,4,2 \quad 1,4,2,3 \quad 2,1,4,3 \quad 2,3,1,4 \quad 2,4,3,1$$
$$3,1,2,4 \quad 3,2,4,1 \quad 3,4,1,2 \quad 4,1,3,2 \quad 4,2,1,3 \quad 4,3,2,1$$

One can show that $\mathfrak{A}_n$ has size $n!/2$ in general.

Part of the importance of alternating groups comes from the fact that for $n \geq 5$ each alternating group $\mathfrak{A}_n$ is simple: it has only trivial normal subgroups.

Definition

The order of a permutation f is the least $m > 0$ such that $f^m = I$.

Lemma

Let the cycles of permutation f have lengths $l_1, \ldots, l_k$ and let m be the LCM of $l_1, \ldots, l_k$. Then m is the order of f.

This has the consequence that

$$f^{-1} = f^{m-1}.$$

Hence we can compute the inverse by fast iteration when the carrier set is finite. Of course, this does not work in the infinite case.

Needless to say, no one would actually compute the inverse this way.

Here is a computationally better way to get at the inverse. Define g by

$$g(f(i)) = i \text{ for } i = 1, \ldots, n.$$

Then $g \circ f = f \circ g = I$ and thus $g = f^{-1}$. This takes linear time.

Here is another, computationally dubious, way: sort the list of pairs

$$\big((f(1), 1), (f(2), 2), \ldots, (f(n), n)\big)$$

in the usual lexicographic order. Then throw away the first components. The resulting permutation is f^{-1}.

Exercise

Explain how the last method works. What is the running time?

So how do we actually compute in a group? Let's first focus on the finite case, for which there always is a brute-force solution – at least in principle.

Definition

Given a finite group $\mathcal{G} = \langle G, * \rangle$ the Cayley table or multiplication table of $\mathcal{G}$ is an G by G matrix with entries in G: the entry in position (a, b) is $a * b$.

It is usually safe to assume that the group elements are represented by integers, so the size of the Cayley table is $\Theta(n^2)$ using a uniform cost function.

That's OK for small n but not for larger ones.

More importantly, Cayley tables tend to shed little light on the structure of the group, all you have is a pile of data.

- $n = 1$: trivial group $\{1\}$

- $n = 2$: $\mathbb{Z}_2$

$$\begin{array}{cc} 1 & a \\ a & 1 \end{array}$$

- $n = 3$: $\mathbb{Z}_3$

$$\begin{array}{ccc} 1 & a & b \\ a & b & 1 \\ b & 1 & a \end{array}$$

- $\mathbb{Z}_4$

1	a	b	c
a	b	c	1
b	c	1	a
c	1	a	b

- Kleinsche Vierergruppe

1	a	b	c
a	1	c	b
b	c	1	a
c	b	a	1

- $n = 5$: $\mathbb{Z}_5$

- $n = 6$: $\mathbb{Z}_6$, $\mathfrak{S}_3$

- $n = 7$: $\mathbb{Z}_7$

- $n = 8$: 5 groups

It gets to be a bit tedious to write down these Cayley tables. Here is a count of the number of finite groups of size n for $n \leq 60$.

Note that the outliers at $n = 32$ and $n = 48$.
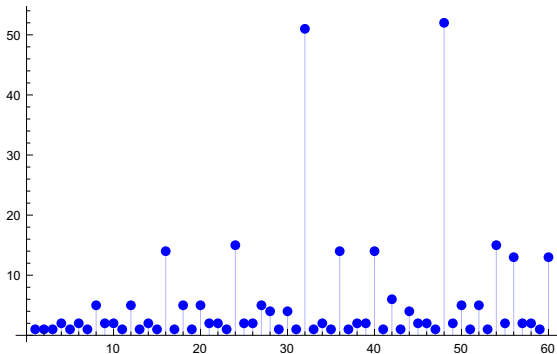
A group G is cyclic if there is some element $a \in G$ such that

$$G = \{\, a^i \mid i \in \mathbb{Z} \,\}$$

In this case a is called a generator.

If G is a finite cyclic group we have

$$G = \{\, a^i \mid 0 \le i < k \,\}$$

where k is the order of a (which is the size of G).

Note that in any finite group G and for any $a \in G$ the subgroup $\{\, a^i \mid 0 \le i < k \,\}$ is cyclic (with generator a).

Up to isomorphism there is only one cyclic group of order k, and it is isomorphic to $\langle \mathbb{Z}_k, +, 0 \rangle$. A generator is 1.

Note that there are other generators, though: ℓ is a generator iff $\gcd(\ell, k) = 1$.

All cyclic groups are commutative.

Recall

$$\mathbb{Z}_m^\star = \{\, x < m \mid \gcd(x, m) = 1 \,\}$$

Example

Here is the Cayley table for $\mathbb{Z}_{20}^\star$.

1	3	7	9	11	13	17	19
3	9	1	7	13	19	11	17
7	1	9	3	17	11	19	13
9	7	3	1	19	17	13	11
11	13	17	19	1	3	7	9
13	19	11	17	3	9	1	7
17	11	19	13	7	1	9	3
19	17	13	11	9	7	3	1

Note the subgroup $\{1, 3, 7, 9\}$ in the top-left corner.

There are important groups resulting from symmetries of regular n-gons. For example, consider the square with four vertices in positions $(\pm 1, \pm 1)$.

What are the rigid motions of the plane that leave the square unchanged in the sense that they place the square on top of itself?
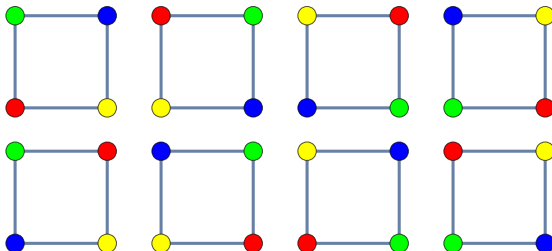
Your geometric intuition should tell you the following:

The motions that leave the square unchanged are precisely:

- Rotations around the origin by multiples of $\pi/2$
  Including the identity, there are 4 of these.

- Reflections along the axes and diagonals.
  There are 4 of these.

But note that the $\pi/2$ rotation has to be applied 4 times to move the square back to its original configuration.

The reflections need to be applied only twice.

These motions naturally form a group: composition of motions is associative, the identity motion is admissible, every motion is reversible, and the composition of two admissible motions is again admissible.

This group is called a dihedral group D_4.

If we replace the square by a regular n-gon we obtain D_n: a group consisting of n rotations and n reflections.

Another tempting generalization is to 3-dimensional space (replace square by cube), but we won't pursue this here.

The symmetries of a pentagon, given by D_5.
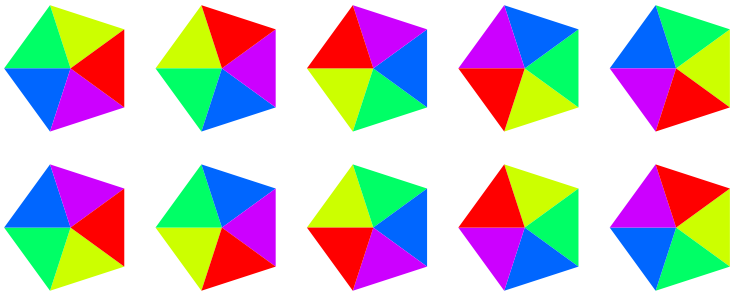
Proposition

The symmetric group on 3 points is isomorphic to the dihedral group D_3.

Proof.

First note that both groups have size 6, so there is a chance the claim might be correct.

The permutations $f = (1, 2)$ and $g = (1, 2, 3)$ (in cycle notation) generate $\mathfrak{S}_3$, so we only need to find their counterparts in D_3.
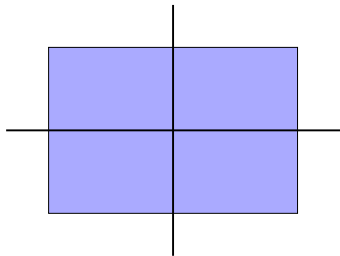
f corresponds to a reflection and g corresponds to a rotation.

$\square$

Exercise

Check the details in the last argument. Why can this line of reasoning not be used to show that $\mathfrak{S}_n$ is isomorphic to D_n in general?

How about the symmetries of a rectangle?

By visual inspection, there are only two reflections, say, a and b. Clearly, $a^2 = b^2 = 1$ and $ab = ba$, so the whole group is just

$$V = \{1, a, b, ab\}$$

A better representation is $\mathbf{2} \times \mathbf{2}$ with addition modulo 2 (or bitwise xor). Since the group is Abelian, we can write $\{00, 01, 10, 11\}$.

The game of (peg) solitaire uses pebbles on a board such as the following one (English version):
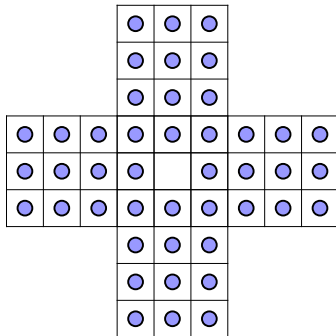
The goal is to "jump-over-and-remove-pebbles" until only one in the middle remains.

It is far from clear that this is even possible–it is for this board, but not for others. A mostly brute-force computational attack succeeds, but does require some cleverness (use of symmetries).

A weak solution is any sequence of moves that leaves just one pebble, somewhere on the board. A strong solution leaves the pebble in the center.

> **Challenge:** Show that any weak solution can be turned into a strong solution by changing just one move.

Think about this a bit, it seems impossibly hard: we have no idea what the space of all weak/strong solutions looks like.

Label the squares s of the board with elements of the Kleinsche Vierergruppe, $v_s \in \{01, 10, 11\}$, as follows:
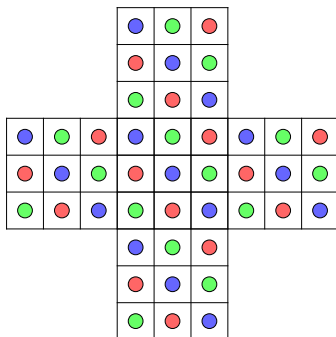
green $\rightsquigarrow 01$

red $\rightsquigarrow 10$

blue $\rightsquigarrow 11$

Note that three consecutive squares, either vertically or horizontally, are always labeled by distinct elements.

Indicate presence or absence of a pebble on square s by a Boolean variable b_s and define the value of the corresponding configuration to be
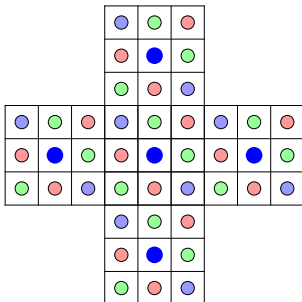
$$\sum_s b_s v_s$$

The value of the whole board is 00, so if we remove the center pebble, the mutilated board has $11 = -11$.

Claim: Any single move does not change the board value.

$$x\, y\, 00 \rightsquigarrow 00\, 00\, z \qquad \text{where } z = x + y$$

In the end, only one pebble in a blue position can be left over. Even better: by symmetry, there can only be 5 possible solutions:

But each one of these can be reached iff all the others can: use symmetry, or change the last move.

Definition

Consider a group $\langle A, \cdot \rangle$. A subgroup of A is a set $B \subseteq A$ such that $\langle B, \odot \rangle$ is a group, where $\odot$ is the restriction of $\cdot$ to B.

This is always written $\langle B, \cdot \rangle$, no one bothers to distinguish between the full group operation $\cdot$ and the restriction $\odot$. Note, though, that the group operation may be much easier to compute in the subgroup.

Example

- $\langle \mathbb{Q}, + \rangle$ is a subgroup of $\langle \mathbb{R}, + \rangle$.

- $\langle \mathbb{Z}, + \rangle$ is a subgroup of $\langle \mathbb{Q}, + \rangle$.

- $\langle 2\mathbb{Z}, + \rangle$ is a subgroup of $\langle \mathbb{Z}, + \rangle$.

- $\{1\}$ is the trivial subgroup of any group (written multiplicatively).

Lemma

Let A be a group and $\emptyset \neq B \subseteq A$.
Then B is a subgroup of A if, and only if, $x, y \in B$ implies $x^{-1} \cdot y \in B$.
If the group is finite than it suffices that $x, y \in B$ implies $x \cdot y \in B$.

Proof.

The first part follows easily from the definition.

For the second part note that B must contain 1: as a finite semigroup B must contain an idempotent, which must be the identity in A. The map

$$B \to B, x \mapsto b \cdot x$$

is a permutation of B for each $b \in B$ (injective implies surjective in the finite case). But then for some $x \in B$: $1 = b \cdot x$ so we have closure under inverses.

$\square$

A map from one group to another is mostly interesting if it preserves structure.

Definition

Suppose G and H are groups. A function $f : G \to H$ is a (group) homomorphism if

$$f(x \cdot y) = f(x) * f(y)$$

Here $\cdot$ is the operation in G, and $*$ the operation in H.

If the function f is in addition injective then it is an monomorphism.
If the function f is in addition surjective then it is an epimorphism.
If the function f is in addition bijective then it is an isomorphism.

Usually one simply writes $f(xy) = f(x)f(y)$ and does not explicitly display the two different group operations.

Proposition

Let $f : G \to H$ be homomorphism.
Then $f(1_G) = 1_H$ and $f(x^{-1}) = f(x)^{-1}$.

Example

- $f : G \to H$, $f(x) = 1$, is a homomorphism.
- $f : G \to G$, $f(x) = x$ is an isomorphism.
- $f : \mathbb{Z} \to \mathbb{Z}_m$, $f(x) = x \bmod m$ is an epimorphism.
- $\log : \mathbb{R}^+ \to \mathbb{R}$ is a isomorphism from $\langle \mathbb{R}^+, \cdot, 1 \rangle$ to $\langle \mathbb{R}, +, 0 \rangle$.

As the last example show, one does not always want to identify isomorphic groups. In fact, the whole purpose of logarithms is to translate multiplication into addition.

Definition

The kernel of a homomorphism $f : G \to H$ is defined as

$$\ker f = \{\, x \in G \mid f(x) = 1 \,\}.$$

Hence

$$f(x) = f(y) \iff y^{-1}x \in \ker f$$

This is slightly different from the kernel relations in combinatorics, but close enough to warrant the same name.

Note that f is injective (a monomorphism) iff the kernel is trivial: $\ker f = 1$.

Proposition

The kernel of a homomorphism is always a subgroup.

We can push this a little bit further based on the last observation:

Definition

For any subgroup $H \subseteq G$ and $a \in G$ define the (left) coset of H by a as

$$aH = \{\, ax \mid x \in H \,\} \subseteq G$$

The number of such cosets is the index of H in G, written $[G : H]$. Right cosets are defined in a similar manner.

Now consider any subgroup $H \subseteq G$ and define a relation

$$x \sim_H y :\Leftrightarrow x^{-1}y \in H$$

We claim that $\sim_H$ is an equivalence relation on G whose equivalence classes are just the cosets aH.

Lemma

$\sim_H$ is an equivalence relation on G, and the equivalence classes of $\sim_H$ all have the same size $|H|$.

Proof.

Reflexivity follows from $1 \in H$.

Symmetry since $x^{-1}y \in H$ implies $(x^{-1}y)^{-1} = y^{-1}x \in H$,

Transitivity since $x^{-1}y, y^{-1}z \in H$ implies $x^{-1}z \in H$.

For the second claim note that $[x]_\sim = xH$.

But $z \mapsto xz$ is a bijection from H to xH. $\qquad\square$

Theorem (Lagrange 1771)

Let G be a finite group, and H any subgroup of G. Then $|G| = |H| \cdot [G : H]$.

In particular, $|H|$ divides $|G|$.

Note how algebra produces a stronger result here: if we look at arbitrary functions $f : A \to B$ then any equivalence relation arises as a kernel relation.

But if we consider groups and homomorphisms we get only very special equivalence relations.

This restriction will turn out to be very helpful to answer various counting problems.

Let $a \in G$. We write $\langle a \rangle$ for the least subgroup of G containing a.

It is not hard to see that
$$\langle a \rangle = \{\, a^i \mid i \in \mathbb{Z} \,\}$$

If G is finite, we have $\langle a \rangle = \{\, a^i \mid i \geq 1 \,\}$.

Definition

The cardinality of $\langle a \rangle$ is the order of a in G.

It follows from Lagrange's theorem that the order of any group element divides the order (cardinality) of the whole group.

Hence for $n = |G|$, $a \in G$ we have $a^n = 1$.

This provides a simple proof for the famous Euler-Fermat theorem.

Recall that $\mathbb{Z}_m^\star$ is the group of elements in $\mathbb{Z}_m$ that have multiplicative inverses.

Also, $\varphi(m)$ is Euler's totient function: $\varphi(m) = |\mathbb{Z}_m^\star|$.

Theorem (Euler-Fermat)

The order of $a \in \mathbb{Z}_m^\star$ divides $\varphi(m)$.

Write G/H for $G/\sim_H$, the collection of H cosets. Wurzelbrunft remembers from algebra lecture that quotients are really only useful if they carry some natural algebraic structure. He proposes to turn G/H into a group as follows:

$$aH * bH := abH$$

and we get the Wurzelbrunft quotient group G/H. An example of this construction are the modular numbers from above.

Since the group structure is inherited from G, this should be quite useful.

Right?

For this to work we need to show that this multiplication is well-defined.

So let $a \sim a'$ and $b \sim b'$. We need

$$abH = a'b'H$$

But all the information we have is that $a' = ah_1$ and $b' = bh_2$, $h_i \in H$.

H is a subgroup, so $h_2 H = H$, which produces

$$a'b'H = ah_1bh_2H = ah_1bH$$

Alas, now we are stuck, we cannot get rid of the pesky h_1.

As it turns out, there is no way to get around this problem: we need more than just a plain subgroup. In fact, in a way, ordinary subgroups are not the right notion of substructure in the case of groups, they don't produce useful quotients.

Definition

A subgroup H of G is normal if for all $x \in H$, $a \in G$: $axa^{-1} \in H$.

In other words, a subgroup is normal if it is invariant under the conjugation maps $x \mapsto axa^{-1}$. Equivalently, $aH = Ha$.

- In a commutative group all subgroups are normal.
- The trivial group 1 and G itself are always normal subgroups (groups that have no other subgroups are called simple, a hugely important concept in the classification of groups).
- There are non-commutative groups where all subgroups are normal, but that is a rare property.
- The group of all translations in the plane is a normal subgroup of the group of all rigid motions (translations plus rotations and reflections).

Now we can fix Wurzelbrunft's argument: assume H is normal. Then

$$abH = aHb = ah_1Hb = a'bH = a'bh_2H = a'b'H$$

Definition

This group is called the quotient group of G modulo (the normal subgroup) H and written G/H.

So where do we get normal subgroups?

Proposition

A subgroup H of G is normal iff it is the kernel of a homomorphism $f : G \to G'$ where G' is some other group.

To hammer this home: let $f : G \to G'$ be a homomorphism and $\sim = \sim_{\ker f}$ the equivalence relation induced by it. We can define a multiplication on the equivalence classes of $\sim$ by setting

$$[x] * [y] := [x\, y]$$

This is well-defined: let $x \sim x'$ and $y \sim y'$, then

$$f(xy) = f(x)f(y) = f(x')f(y') = f(x'y'),$$

so that $[xy] = [x'y']$. It is not hard to see that this produces a group structure on $G/\sim$.

Example 1 101

Let G be the integers under addition and $H = m\mathbb{Z}$. Then

$$x \sim y \iff y - x \in m\mathbb{Z}$$
$$\iff x = y \pmod{m}$$

H is the kernel of the epimorphism $x \mapsto x \bmod m$.

Example 2 102

Let G be the group of all permutations on $[n]$. Define

$$f(x) = \begin{cases} 0 & \text{if } x \text{ is even,} \\ 1 & \text{otherwise.} \end{cases}$$

Then f is homomorphism from G to the additive group $\mathbb{Z}_2$.

The kernel of f is the subgroup

$$H = \{\, x \in G \mid x \text{ even} \,\}$$

Note that $|H| = |G|/2 = n!/2$.

Example 3 103

Consider the multiplicative group

$$G = \mathbb{Z}_{13}^* = \{1, 2, \ldots, 12\}$$

We one can check that $H = \{1, 3, 9\}$ is a subgroup with cosets

$$H = \{1, 3, 9\}, 2H = \{2, 5, 6\}, 4H = \{4, 10, 12\}, 7H = \{7, 8, 11\}$$

The multiplication table for G/H written with canonical representatives is

$$
\begin{array}{cccc}
1 & 2 & 4 & 7 \\
2 & 4 & 7 & 1 \\
4 & 7 & 1 & 2 \\
7 & 1 & 2 & 4 \\
\end{array}
$$

and is isomorphic to the additive group $\mathbb{Z}_4$.

Lemma

Every homomorphism $f : G \to H$ can be written as $f = \nu \circ \iota$ where ν is an epimorphism and ι is a monomorphism.

Proof.

Let $K \subseteq G$ be the kernel of f, a normal subgroup. Define

$$\nu : G \to G/K \qquad x \mapsto [x]$$
$$\iota : G/K \to H \qquad [x] \mapsto f(x)$$

It is easy to check that these functions work. $\qquad\qquad\square$

To obtain the quotient group G/H we need to factor by a special type of equivalence relation.

Definition

Suppose G is a group and $\sim$ an equivalence relation on G. $\sim$ is a congruence if for all $x, y, x', y' \in G$:

$$x \sim x', y \sim y' \quad \text{implies} \quad xy \sim x'y'.$$

Again, congruences are very important since they make it possible to define a group structure on the quotient set $G/\sim$:

$$[x] \cdot [y] = [x \cdot y]$$

Unfortunately, the equivalence relations $\sim_H$ for arbitrary subgroups H are not congruences in general, we need normal subgroups for this to work.

Proposition

If H is a normal subgroup, then $\sim_H$ is a congruence.

Proposition

H is the kernel of a homomorphism $f : G \to G'$ iff H is normal.

Exercise

Prove these propositions.

You know this already. E.g., let p and q be two distinct primes.

$$f : \mathbb{Z} \to \mathbb{Z}_p \times \mathbb{Z}_q$$
$$f(x) = (x \bmod p, x \bmod q)$$

Then $H = pq\,\mathbb{Z}$ and the quotient is $\mathbb{Z}/(pq\mathbb{Z}) = \mathbb{Z}_{pq}$.

One can show that f is an epimorphism (this requires a little argument).
Hence $\mathbb{Z}_{pq}$ is isomorphic to $\mathbb{Z}_p \times \mathbb{Z}_q$.

Hence we can either compute

- with one number modulo pq, or
- with two numbers, one modulo p and the other modulo q.

$\mathbb{Z}_{pq}$ and $\mathbb{Z}_p \times \mathbb{Z}_q$ are isomorphic, but computationally there is a difference. This can be exploited sometimes to fake high-precision computations with small word sizes.

Also note that the correctness proof for RSA more or less requires the product representation.

Theorem (Cayley 1854)

Every group is isomorphic to a subgroup of a permutation group.

Proof. Let $\mathcal{A} = \langle A, \cdot \rangle$ be a group, and let $\mathfrak{S}_A$ be the full permutation group over A. Define a map

$$\varphi : \mathcal{A} \to \mathfrak{S}_A$$
$$\varphi(a)(x) = x \cdot a$$

Then φ is a homomorphism: $\varphi(a \cdot b) = \varphi(a) \circ \varphi(b)$. Moreover, φ is mono: the kernel is just $1 \in A$. Hence, the range of φ is a subgroup of $\mathfrak{S}_A$ that is isomorphic to $\mathcal{A}$. □

Note that this representation is not too helpful computationally: each permutation in $\mathfrak{S}_A$ has the same size as A.

Recall our proof of the fact that no permutation is both even and odd.

One way to explain (and make precise) what is going on there is to consider the sign function from the group of all permutations

$$\mathsf{sg} : \mathfrak{S}_n \to \{+1, -1\}$$
$$\mathsf{sg}(\sigma) = P_\sigma(\boldsymbol{x})/P(\boldsymbol{x})$$

where the operation on the right is ordinary multiplication. It is not hard to see that sg is a homomorphism and the kernel of sg is exactly the collection of all even permutations.

In other words, $\mathfrak{A}_n$ is the kernel of the homomorphism sg.