CDM

Semigroups and Groups

Klaus Sutner

Carnegie Mellon University

Spring 2025

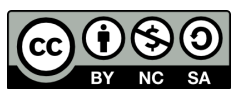

We will start with three fundamental types of structures.

- magmas
- semigroups
- groups

Magmas are borderline algebraic structures.
But they help to clarify the basic concepts.

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. 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 elements of the magma.

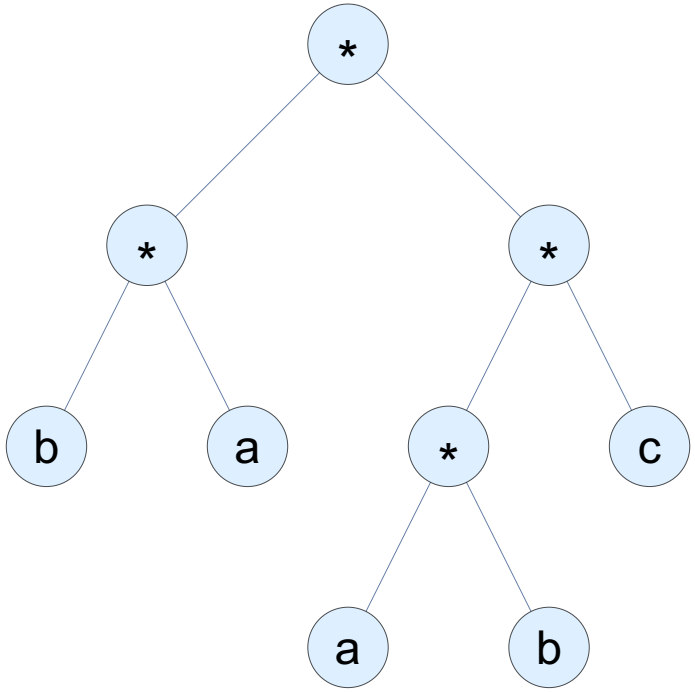

We can evaluate these trees bottom-up eval : $\mathrm{trees} \to M$

- The natural numbers with exponentiation.
- The integers with subtraction.
- The positive rationals with division.
- Lists over some groundset with join as multiplication.

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).

It is often helpful to translate combinatorial structures into algebraic ones.

E.g., suppose $\mathcal{G} = \langle V; E \rangle$ is a digraph and define the magma

$$\mathcal{A}(\mathcal{G}) = \langle V \cup \{\bot\}; * \rangle$$

where $\bot \notin V$ is a new point and

$$u * v = \begin{cases} u & \text{if } u \to v \in E, \\ \bot & \text{otherwise.} \end{cases}$$

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?

First off, some magmas are actually interesting on their own—though it is true that semigroups and groups are much more important.

Second, it is a good idea to try to understand if a concept or result in a semigroup or group also works over magmas.

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

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

- A substructure of $\mathcal{A}$ consists of a set $\emptyset \neq B \subseteq A$ 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

$$x \ \rho \ x' \text{ and } y \ \rho \ y' \quad \text{implies} \quad x * y \ \rho \ x' * y'$$

  a so-called congruence.

- The (Cartesian) product of $\langle A; * \rangle$ and $\langle B; \cdot \rangle$ is the structure

$$\langle A \times B; \otimes \rangle \qquad (x, y) \otimes (x', y') = (x * x', y \cdot y')$$

Another standard problem in algebra is to determine free structures: they satisfy a list of axioms, but have no additional, special properties. [†].

A real definition of freeness requires a bit of category theory[†], we won't go there.

So suppose we have a set X of intended elements of a generic magma. What would the free magma over X look like?

Just the binary trees from above, leaves labeled by the ground set X. The operation is the join of trees: add a new root node, attach the two trees as left children.

[†]Not every reasonable set of algebraic axioms has free structures; e.g. topological groups, ordered groups, fields, integral domains do not.

[†]Essentially we want a left adjoint to the forgetful functor to Set.

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

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

It would not be unreasonable to say that semigroups really are the place where algebra starts, magmas are more combinatorial in nature.

So?

How big is the difference between a magma and a semigroup?

One way to gauge the difference is to compare the free structures (they exist in both cases).

Another is to display specific concepts that are useful for semigroups but not so much for magmas.

Here goes.

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. We might as well think of them as a list.

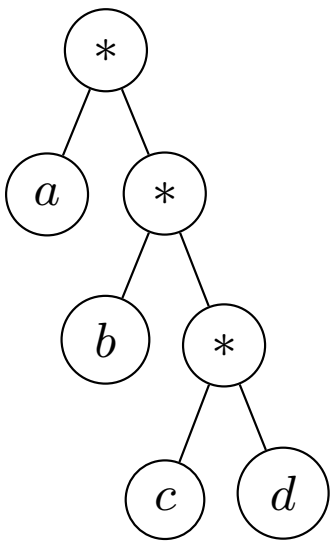

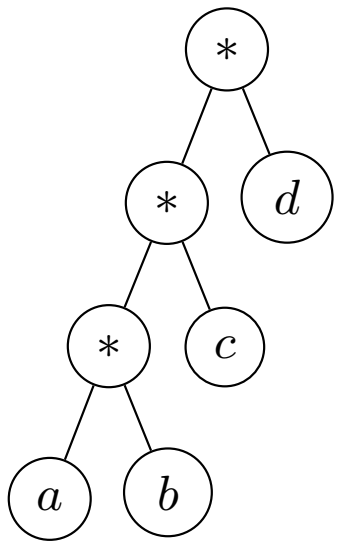

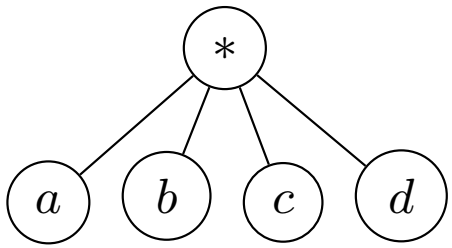

But once we think of the elements as flat lists

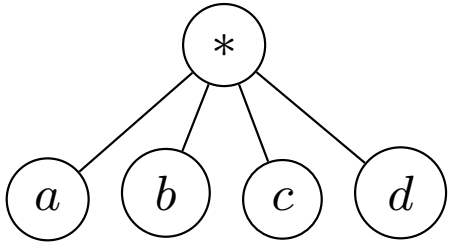

we are basically left with a sequence like $abcd$, a word over the alphabet A.

So the set of non-empty words over A can be thought of as the free semigroup generated by A. This perspective turns out to be very useful in automata theory.

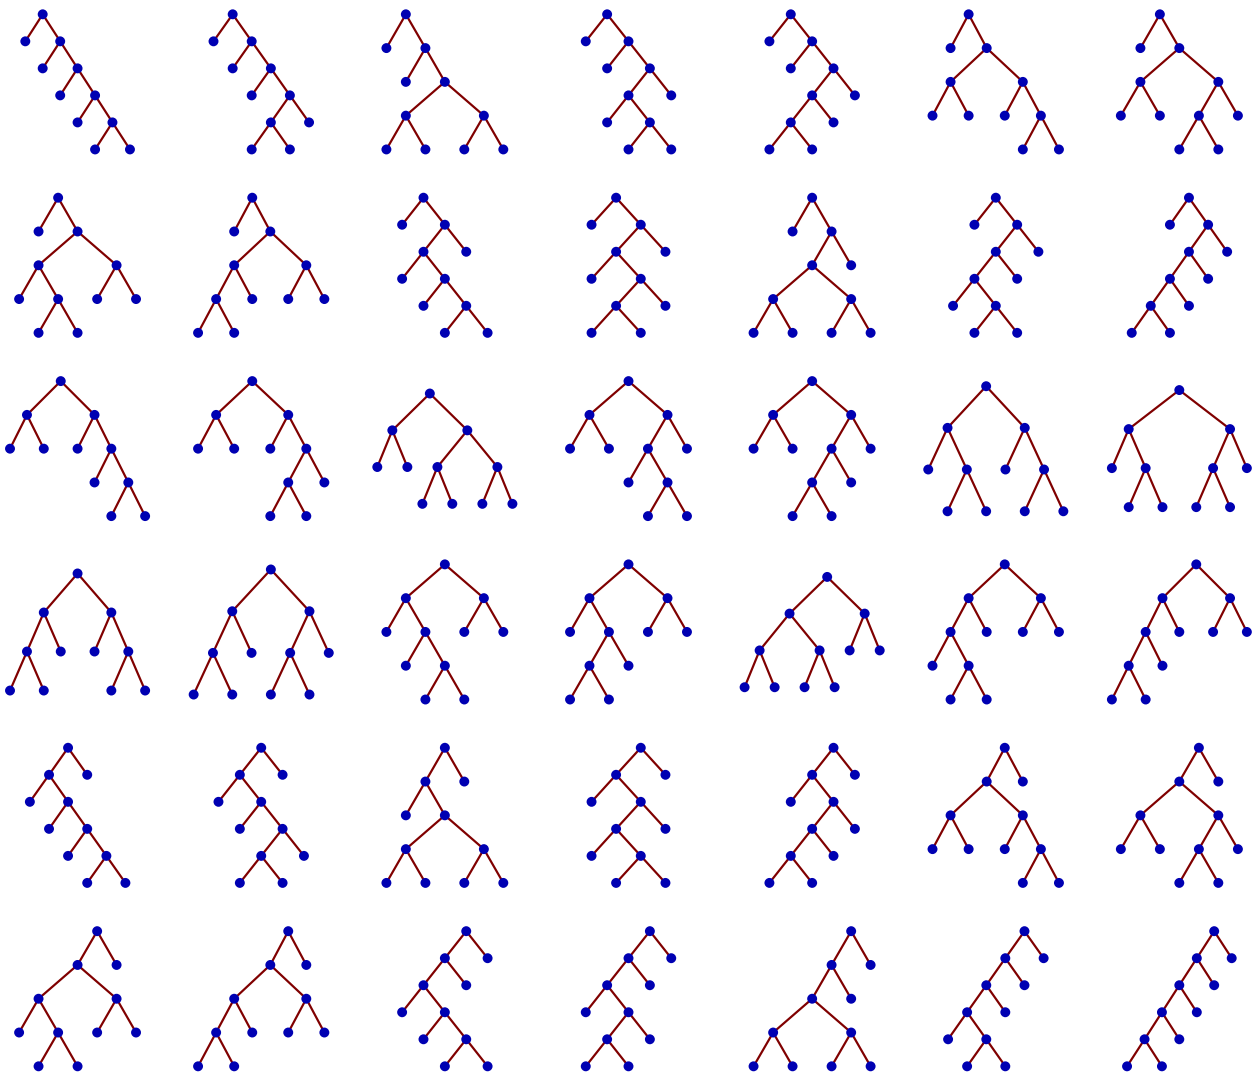

Definition

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

Idempotents are a bit weaker than an identity and turn out to be critical for the study of semigroups. Here is a key lemma.

Lemma

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

The lemma fails in the infinite case, think about a^+.

Definition

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

One usually writes monoids with signature $(2, 0)$ as

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

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 natural numbers with addition form a monoid; 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

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.

When one tries to axiomatize some sort of algebraic structure, one always has a number of civilized examples in mind. The first task is to make sure that the axioms capture all those examples.

It may very well happen, though, that the axioms have strange, unintended models that one does not really anticipate. It is well worthwhile to look around for weird examples (and possibly adjust the axioms if need be).

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).

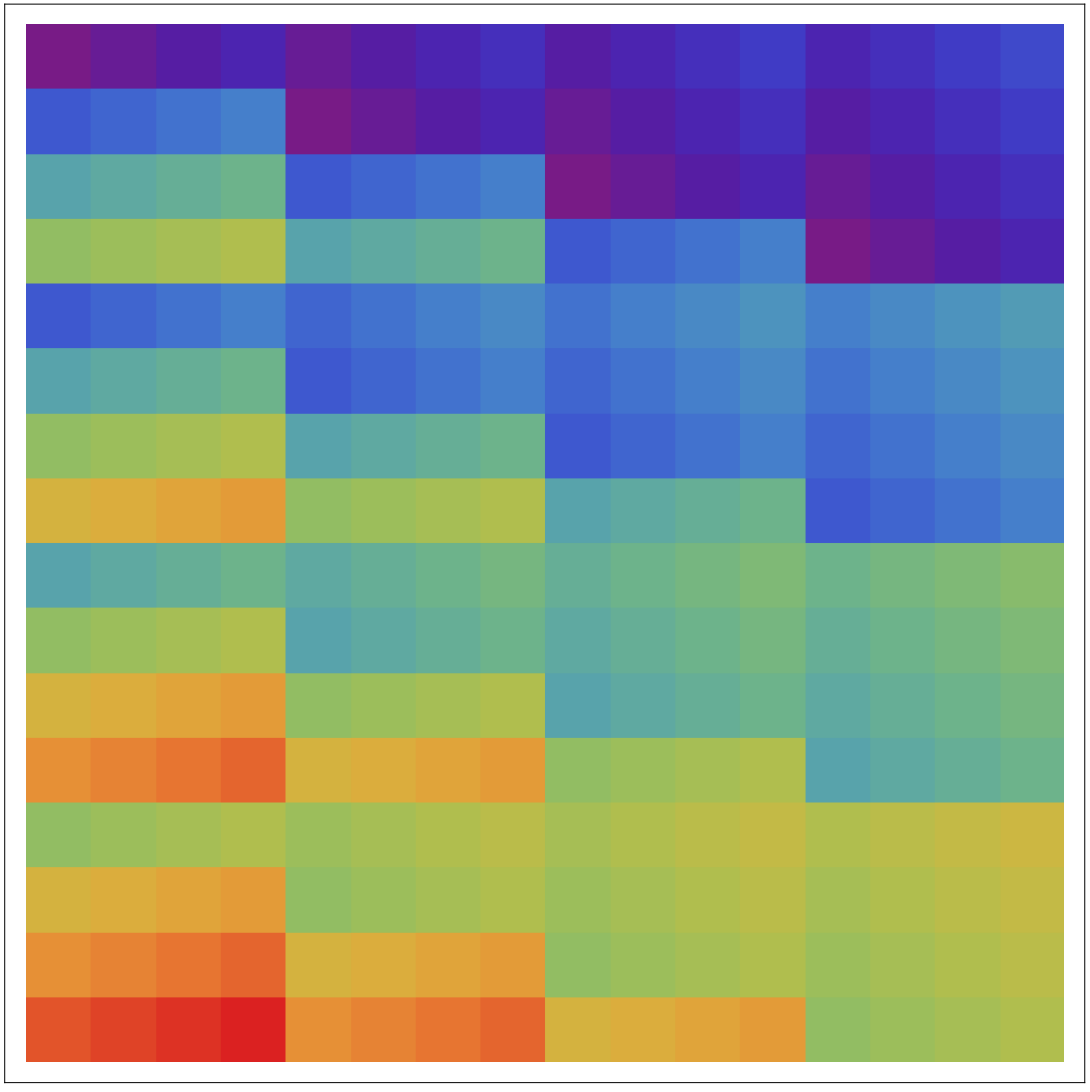

Consider a semigroup $\mathcal{A} = \langle A; * \rangle$.

- A subsemigroup of $\mathcal{A}$ consists of a set $\emptyset \neq B \subseteq A$ that is closed under $*$.
- A semigroup morphism 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 semigroup congruence of $\langle A; *, e \rangle$ is an equivalence relation ρ such that $x \ \rho \ u$ and $y \ \rho \ v$ implies $x * y \ \rho \ u * v$.
- The semigroup product of $\langle A; * \rangle$ and $\langle B; \cdot \rangle$ is the structure $\langle A \times B; \otimes \rangle$ where $(x, y) \otimes (u, v) = (x * u, y \cdot v)$.

The reason congruences are important is that they make it possible to form quotients.

Suppose ρ is a congruence and define the semigroup $\mathcal{A}/\rho$ by

- carrier set: $A/\rho = \{\, [x]_\rho \mid x \in A \,\}$
- operation $[x] * [y] = [x * y]$

One can check that this really produces a semigroup.

Warning: One really needs a congruence here, this does not work for an arbitrary equivalence relation (the operation is not well-defined in general).

Proposition

Suppose $f : \mathcal{A} \to \mathcal{B}$ is a monoid morphism.
Then the kernel relation $x \ \rho \ y \Leftrightarrow f(x) = f(y)$ is a congruence.

Proof.

Any kernel relation is trivially an equivalence relation. Suppose $x \ \rho \ x'$ and $y \ \rho \ y'$. Then

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

so that $x * y \ \rho \ x' * y'$. □

If you still find the congruence condition strange, here is another way to think about it: we have an equivalence relation $\rho \subseteq A \times A$ that is also a submonoid of $A \times A$ (construed as a monoid product).

Consider a monoid $\mathcal{A} = \langle A; *, e \rangle$.

- A submonoid of $\mathcal{A}$ consists of a set $\emptyset \neq B \subseteq A$ that is closed under $*$ and contains e.
- A monoid morphism from $\langle A; *, e \rangle$ to $\langle B; \cdot, e' \rangle$ is a map $\varphi : A \to B$ such that $\varphi(x * y) = \varphi(x) \cdot \varphi(y)$ and $\varphi(e) = e'$.

The only difference is that we have additional conditions regarding the neutral element. Warning: a monoid may have a subsemigroup that is not a submonoid: $\{0\}$ in $\langle \mathbb{Z}; \cdot, 1 \rangle$.

For congruences and products, nothing changes. The neutral element in $A \times B$ is (e_A, e_B).

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

$$a * x = b$$

As an example, consider the monoid $\langle \mathbb{N}; +, 0 \rangle$.

Then $a + x = b$ has exactly one solution whenever $a \leq b$, no solutions otherwise.

Now switch to the multiplicative monoid $\langle \mathbb{N}; \cdot, 1 \rangle$. Then $0 \cdot x = 0$ has infinitely many solution, $2 \cdot x = 42$ has exactly one, and $2 \cdot x = 41$ has none.

Uniqueness can be ensured via the following (left) cancellation property:

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

Note that this property is not equational, we need an implication between equations to express cancellation.

Exercise

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

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 every $x \in G$ has an inverse $y \in G$:

$$x \cdot y = y \cdot x = e$$

The inverse y is uniquely determined by x.

Hence 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 group axioms are then equational, we don't need any pesky quantifiers:

$$x \cdot (y \cdot z) = (x \cdot y) \cdot z$$
$$x \cdot e = e \cdot x = x$$
$$x \cdot x^{-1} = x^{-1} \cdot x = e$$

Actually, it turns out that we can get away with one-sided versions of the axioms as in

$$x \cdot (y \cdot z) = (x \cdot y) \cdot z$$
$$x \cdot e = x$$
$$x \cdot x^{-1} = e$$

Proof

Assume only the right identity and right inverses as in the one-sided axioms.

A right inverse is also a left inverse:

$$x^{-1} \cdot x = (x^{-1} \cdot x) \cdot e = \left(x^{-1} \cdot x\right)\left(x^{-1} \cdot (x^{-1})^{-1}\right) = e$$

Hence, a right identity is also a left identity:

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

In a situation like this, one often chooses the apparently weaker, if slightly cryptic, axioms.

At any rate, 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.

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

$$\langle G; +, -, 0 \rangle \qquad \langle G; +, 0 \rangle \qquad \langle G; + \rangle$$

and general groups multiplicatively as

$$\langle G; \cdot, {}^{-1}, 1 \rangle \qquad \langle G; \cdot, 1 \rangle \qquad \langle G; \cdot \rangle$$

Often one omits the multiplication operator and uses concatenation instead. xy is easier on the eye than $x \cdot y$, but this gets tricky when dealing with multiple operations.

Consider a group $\mathcal{A} = \langle A; *, {}^{-1}, e \rangle$.

- A subgroup consists of a set $\emptyset \neq B \subseteq A$ that is closed under $*$ and ${}^{-1}$.
- A group morphism from $\langle A; *, e \rangle$ to $\langle B; \cdot, e' \rangle$ is a map $\varphi : A \to B$ such that $\varphi(x * y) = \varphi(x) \cdot \varphi(y)$.

For quotients and product nothing really changes from semigroups and monoids.

Note that we dropped the condition that morphisms preserve the neutral element: this can be derived from the group axioms.

Subgroups have a simple but critical property: they make it possible to partition the group into disjoint sets of equal size. To this end, define the cosets of a subgroup $\emptyset \neq B \subseteq A$ by

$$aB = \{\, ax \mid x \in B \,\}$$

Proposition

Let $B \subseteq A$ be a subgroup.
Then the cosets aB partition A and have the same size.

Proof. Suppose $aB \cap a'B \neq \emptyset$, so $ab = a'b'$. But then $a = a'b'b^{-1} = a'b''$ and therefore $aB \subseteq a'B$. Done by symmetry.

There is a simple bijection $\widehat{a} : B \to aB : \widehat{a}(x) = ax$.

□

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.

Often one faces a large, ambient group G and a small collection $X \subseteq G$ of group elements that are of particular interest. One would like to work in a small subgroup $H \subseteq G$ such that $X \subseteq H$. In fact, we want the smallest such subgroup.

Definition

Let G be a group and $X \subseteq G$. The group generated by X is the least subgroup of G that contains X. X is a set of generators for this subgroup.

In symbols: $\langle X \rangle$.

Example: in the group $\langle \mathbb{Z}; + \rangle$, $X = \{a\}$ generates the subgroup $2\mathbb{Z}$, and $X = \{a, b\}$ generates the group $\gcd(a, b)\,\mathbb{Z}$.

In set theory it is easy to show that $\langle X \rangle$ always exists:

$$\langle X \rangle = \bigcap \{ H \subseteq G \mid X \subseteq H \text{ subgroup} \}$$

This works since the set on the right is non-empty, $H = G$ always works.

Alas, computationally this is pretty useless: $\langle X \rangle$ is one of the subgroups on the right, our characterization is mildly circular[†]. Also, even for G finite, we run into exponentially many candidate subsets.

Real Problem:

How do we actually compute $\langle X \rangle$ from X?
In particular when G is finite?

[†] This is called an impredicative definition and is nowadays generally accepted as a perfectly good way of defining mathematical objects.

Bottom-Up Construction

We can construct $\langle X \rangle$ from a given set of generators $X \subseteq G$ by induction.

- $H_0 = X \cup \{1\}$
- H_{n+1} is the closure of H_n with respect to multiplication and inversion
- $H = \bigcup_{n \geq 0} H_n$

When G or at least H is finite, this is a perfectly good algorithm, at least disregarding efficiency.

In the infinite case, this constuction is still logically correct, but now it needs to run for infinitely many steps.

Let G be a group.

Definition

The order of G is the cardinality of G.

The order of an element $a \in G$ is the order of the subgroup $\langle a \rangle$.

In symbols: $|G|$ or $\mathrm{ord}(G)$; $\mathrm{ord}(a)$.

The concept of order is most important when the cardinality is finite.
In the infinite case one writes $\mathrm{ord}(a) = \infty$ (though $\aleph_0$ would be better):

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

If the order k of a is finite, we even have

$$\langle a \rangle = \{\, a^i \mid 0 \leq i < k \,\}$$

Felix Klein's "Erlanger Programm" of 1872 proposed to characterize geometries by studying groups of linear transformations.

Symmetries are a critical application and, in fact, the historical origin of groups.

Consider the square with four vertices in positions $(\pm 1, \pm 1)$.

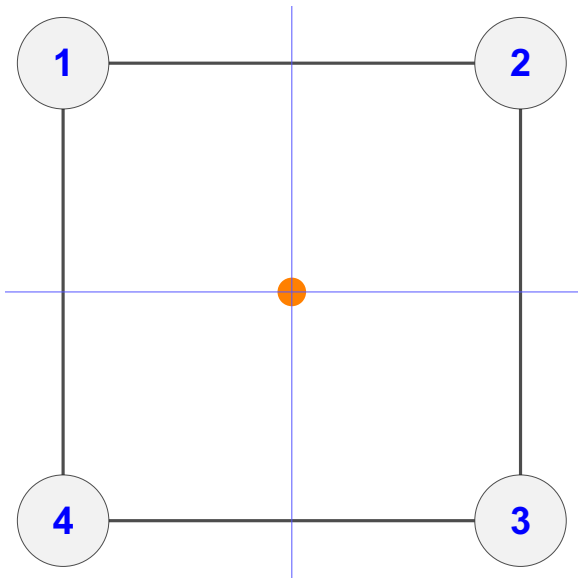

What are the rigid motions of the plane that leave the square unchanged in the sense that they place the square on top of itself?

Your geometric intuition should immediately lead to the following admissible operations.

- Rotations around the origin by multiples of $\pi/2$.
  There are essentially only 4 of these, including the trivial one.
- Reflections along the axes and diagonals.
  There are 4 of these.

There might be more, though, say a combination of a rotation and a reflection.

Let's write α for clockwise rotation by $\pi/2$ and β for reflection along the counter-diagonal.

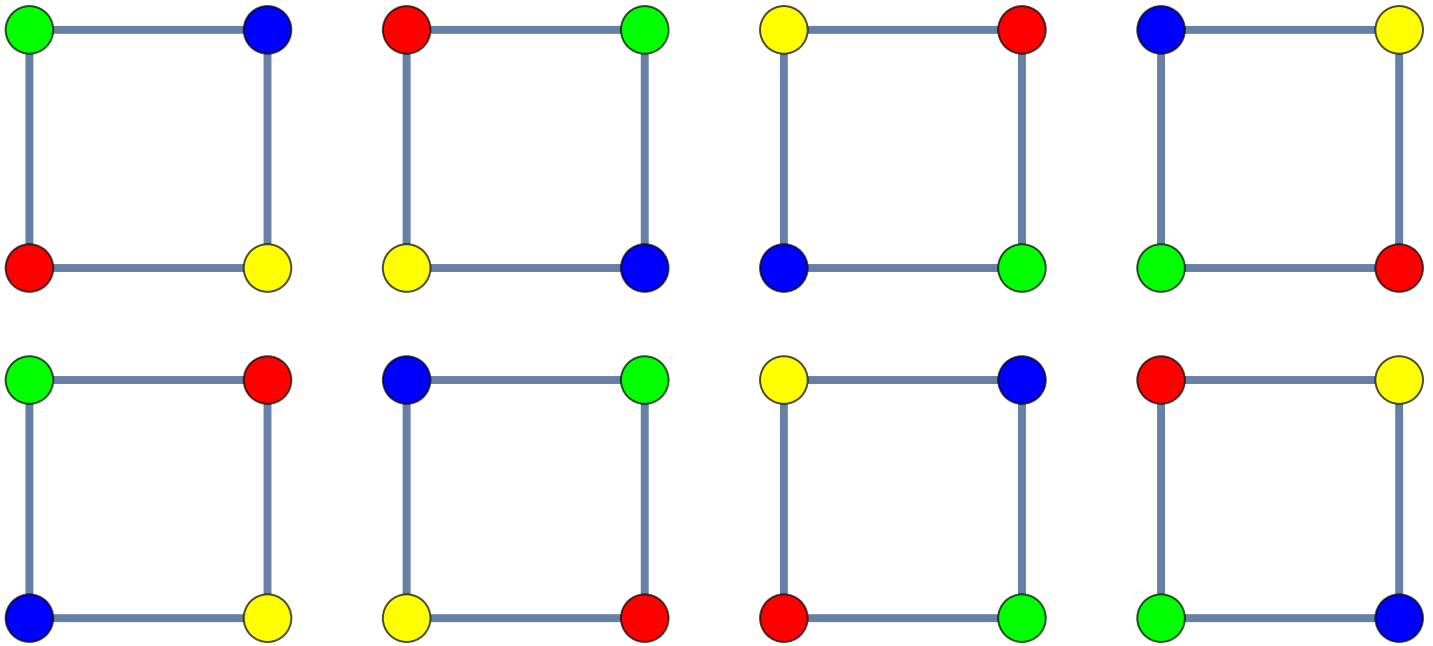

The first row is obtained by repeated application of α.

The second row is obtained by applying β, and then α repeatedly.
Note that all other reflections appear in this row, too.

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.
It has size 8 and is generated by α and β, elements of order 4 and 2, respectively.

The important identities in this group are

$$\alpha^4 = \beta^2 = 1 \qquad \beta\alpha = \alpha^3\beta$$

If we start with a regular n-gon instead, we get the dihedral group D_n.

For example, the dihedral group D_3 represents the symmetries of an equilateral triangle.

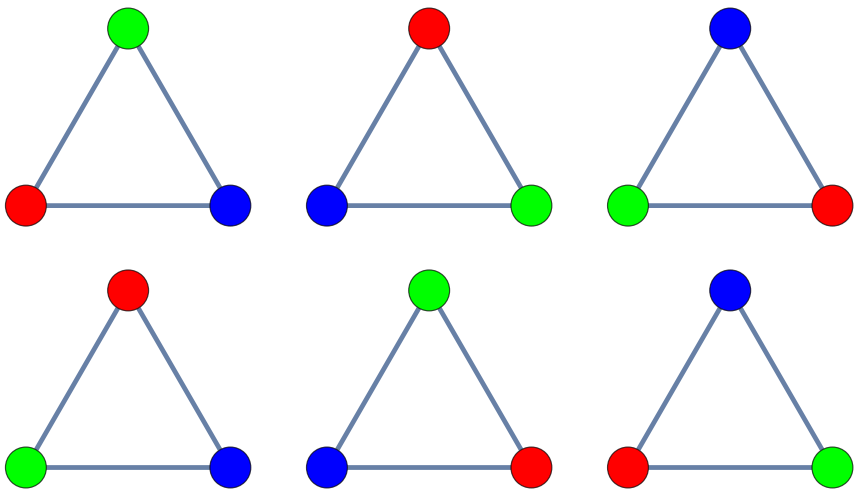

We can apply D_3 to more complicated sets in geometry that are somewhat similar to a plain equilateral triangle.

Klein's program worked perfectly well in the context of the kinds of geometric objects studied in the 19th century, lines, circles, planes, surfaces and so on.

But in the 20th century new objects like fractals were discovered that are somewhat better aligned with aspcets of physical reality. Think about the coastline of Britain or clouds.

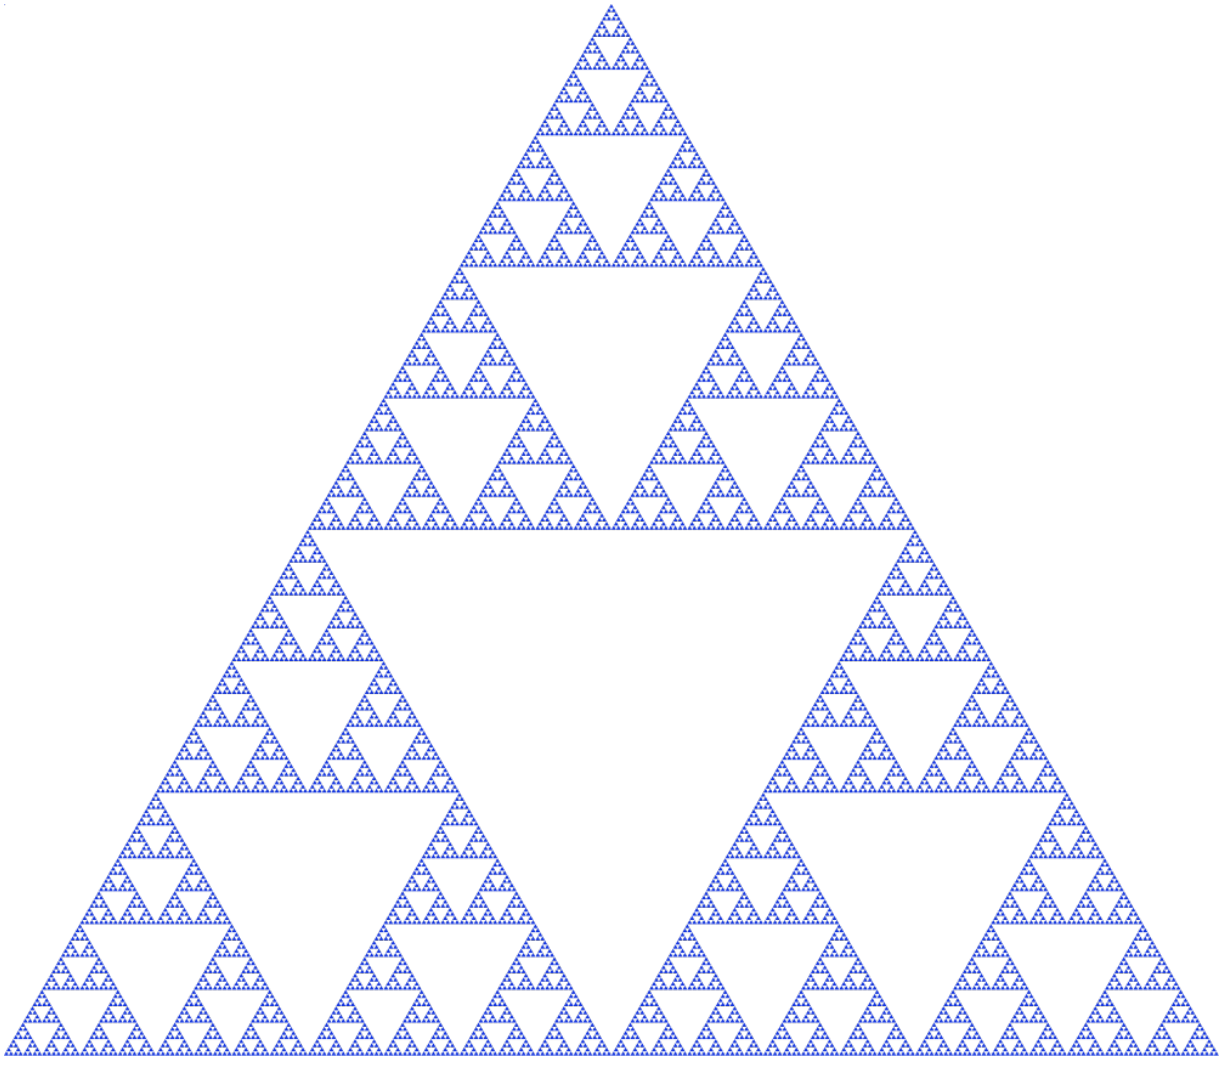

For our purposes the most important examples of groups are those comprised of permutations.

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 $A = [n]$. There are two standard ways to write permutations: two-line representation and one-line representation.

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

$$\big[f(1),\ f(2),\ f(r),\ \dots,\ f(n{-}1),\ f(n)\big]$$

We have written $[a_1, \dots, a_n]$ to specifically indicate a map from $[n] \to [n]$. This is slightly less dangerous than generic list notattion $(a_1, \dots, a_n)$, which could mean a great many things.

Suppose $f : [n] \to [n]$ is some permutation. We can think of f as a directed graph G_f, the functional digraph of f:

$$V = \{1, 2, \ldots, n\}$$
$$E = \{\, x \to f(x) \mid x \in V \,\}$$

f is a permutation iff G_f consists of a collection of disjoint cycles.

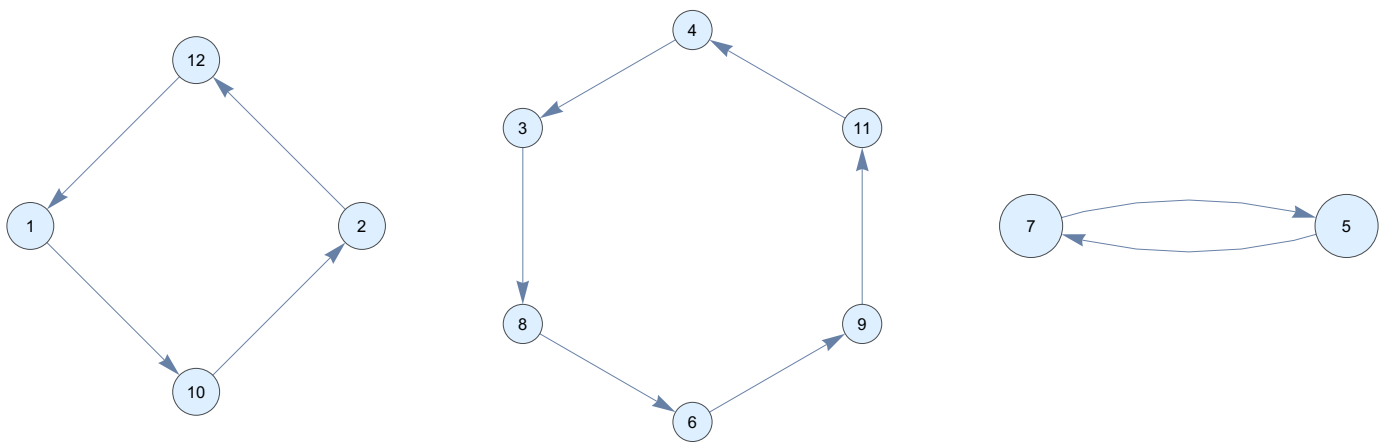

The cycle decomposition of f looks like

$$f = (v_{1,0}, \ldots, v_{1,q_1-1})\,(v_{2,0}, \ldots, v_{2,q_2-1}) \ldots (v_{k,0}, \ldots, v_{k,q_k-1})$$

where all the $v_{i,j}$ are all distinct, $\sum q_i = n$ and $f(v_{i,j}) = v_{i,j+1 \bmod q_i}$.

We write

$$\mathsf{cc}_i(f) = \text{ number of cycles of length } i \text{ in } f$$

$$\mathsf{cc}(f) = \sum \mathsf{cc}_i(f)$$

The list $\mathsf{cc}_1(f), \mathsf{cc}_2(f), \ldots, \mathsf{cc}_n(f)$ is the cycle shape of f.

Thus, $\mathsf{cc}(f)$ is the total number of cycles in f, $1 \leq \mathsf{cc}(f) \leq n$.

Also note that $\sum_i i\ \mathsf{cc}_i(f) = n$.

Example

The functional digraph above represents the permutation

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

In cycle notation, f looks like so:

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

Careful, even if we have just a single cycle of length n, a cyclic shift $(a_1, a_2, \ldots, a_n)$, this is **not** one-line notation.

Omitting Fixed Points

Another trap: it is completely standard in algebra texts to omit fixed points.

The cycle decomposition of

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

would be written as

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

leaving out the fixed points 10 and 12.

Here

$$\mathsf{cc}(f) = 5$$
$$\mathsf{cs}(f) = (2, 2, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0)$$

Computing the cycle decomposition comes down to computing the strongly connected components in the functional digraph and clearly can be handled in linear time and space[†].

Lemma

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

Careful, though, there really is no such thing as **the** cycle decomposition:

- Since the cycles are disjoint we can rearrange them aribitrarily.
- We can rotate each cycle without changing the permutation.

Example:

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

[†]Actually, logarithmic space if we don't need to store the cycles.

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).

RCCD may seem more natural from the implementor's point of view (we usually work from 1 to n), but CCD has better combinatorial properties.

$D =$ nil
for $x = 1, \ldots, n$ **do**
    if x is new
        $C = (x)$
        while $f(x)$ is new
            $x = f(x)$
            append x to C
    append C to D
return D

Most people would write this snippent when asked to implement cycle decomposition.

This algorithm produces RCCD.

Here are the CCDs for all elements of $\mathfrak{S}_4$, enumerated in lex order.

$(1)\,(2)\,(3)\,(4)$	$(1)\,(2)\,(4,3)$	$(1)\,(3,2)\,(4)$	$(1)\,(4,2,3)$
$(1)\,(4,3,2)$	$(1)\,(3)\,(4,2)$	$(2,1)\,(3)\,(4)$	$(2,1)\,(4,3)$
$(3,1,2)\,(4)$	$(4,1,2,3)$	$(4,3,1,2)$	$(3)\,(4,1,2)$
$(3,2,1)\,(4)$	$(4,2,1,3)$	$(2)\,(3,1)\,(4)$	$(2)\,(4,1,3)$
$(3,1)\,(4,2)$	$(4,1,3,2)$	$(4,3,2,1)$	$(3)\,(4,2,1)$
$(2)\,(4,3,1)$	$(2)\,(3)\,(4,1)$	$(4,2,3,1)$	$(3,2)\,(4,1)$

We have writtent the fixed points explicitly.

Flattening CCDs

Here are these CCDs flattened out.

$1,2,3,4$	$1,2,4,3$	$1,3,2,4$	$1,4,2,3$
$1,4,3,2$	$1,3,4,2$	$2,1,3,4$	$2,1,4,3$
$3,1,2,4$	$4,1,2,3$	$4,3,1,2$	$3,4,1,2$
$3,2,1,4$	$4,2,1,3$	$2,3,1,4$	$2,4,1,3$
$3,1,4,2$	$4,1,3,2$	$4,3,2,1$	$3,4,2,1$
$2,4,3,1$	$2,3,4,1$	$4,2,3,1$	$3,2,4,1$

We get all permutations in one-line notation.
Could this be coincidence?

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

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

This corresponds to the natural way one reads a diagram:

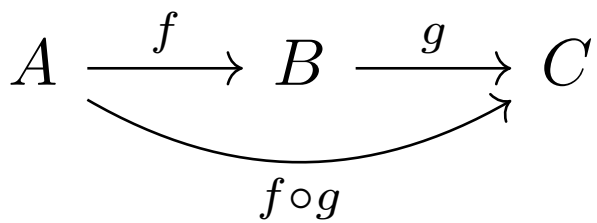

Some (misguided) texts use the opposite convention.
Unfortunately, they are currently the vast majority.

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:

$$\mathsf{flat} : \mathsf{List}(\mathsf{List}(\mathbb{N})) \to \mathsf{List}(\mathbb{N})$$

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 $\mathsf{CCD} \circ \mathsf{flat}$ *is a bijection on* $\mathfrak{S}_n$.

Here are two decomposition questions of the kind an algebraist would be interested in.

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 $a \neq b$.

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 transpositions is not commutative.

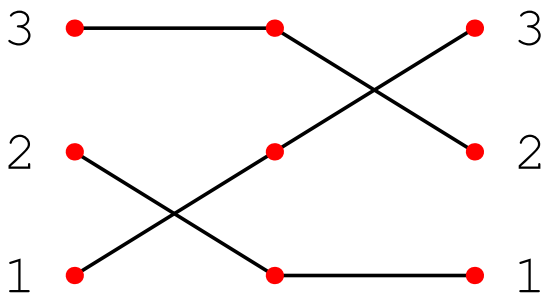

$(1,2) \circ (2,3) = [3,1,2]$

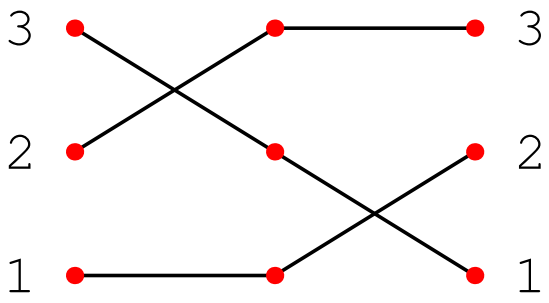

$(2,3) \circ (1,2) = [2,3,1]$

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.

Sketch of proof.

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

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

□

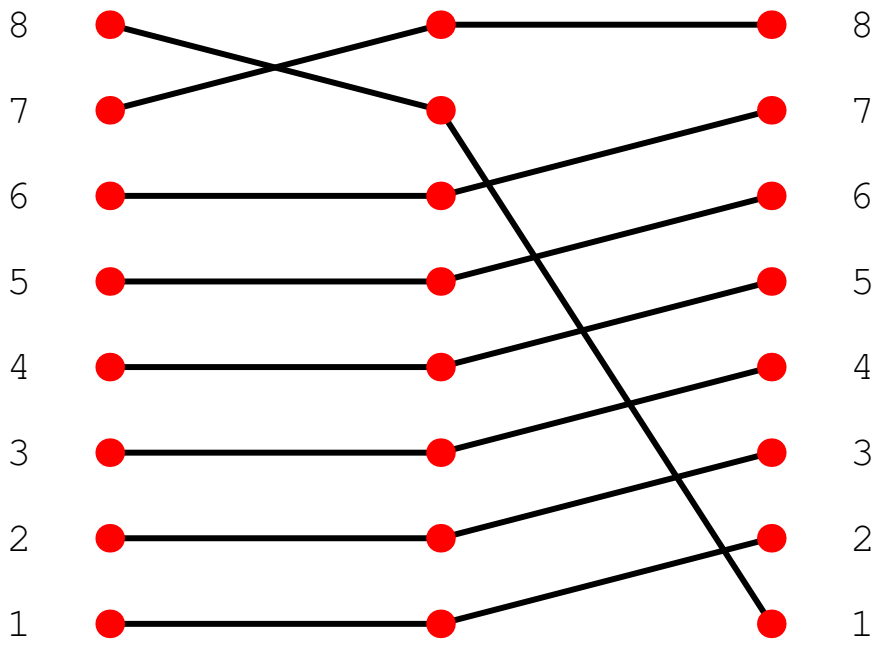
8
7
6
5
4
3
2
1
8
7
6
5
4
3
2
1

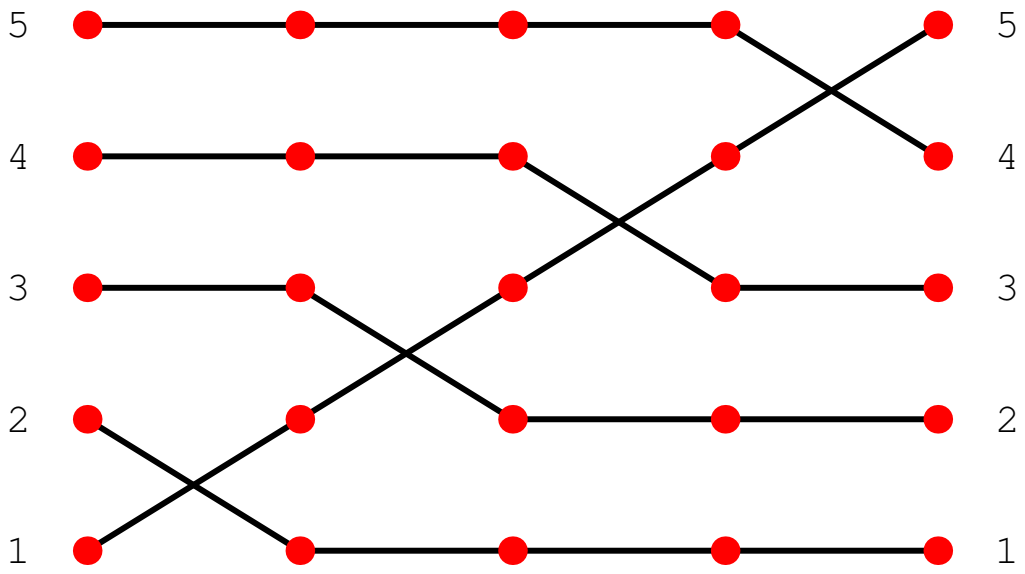

A full decomposition of a cycle intro transpositions.
So $(1,2)(2,3)(3,4)(4,5) = (5,4,3,2,1)$.

More Pictures

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

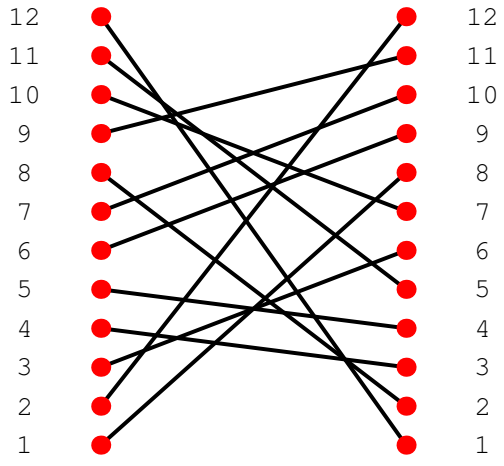

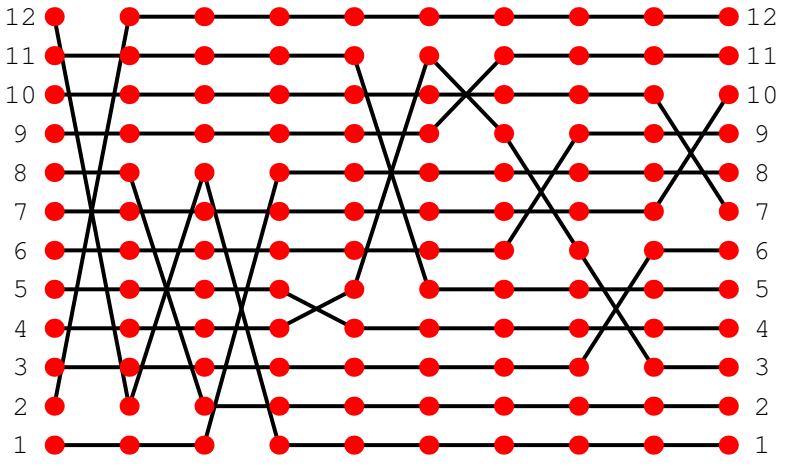

Exercise

Find an algorithm to generate the picture on the right.

Proposition

$$(a, b) \circ (b, c) \circ (a, b) = (a, c)$$

$$(1, \ldots, n)^k \circ (1, 2) \circ (n, \ldots, 1)^k = (k{+}1, k{+}2)$$

This works for $0 \leq k \leq n - 2$.

Make sure to prove these identities.

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

Definition

A permutation is even (or odd) if it can be written as the product of on even (or odd) number of transpositions.

Note the cautious wording: this does not say that every permutation is either even or odd, some might be both. But that cannot happen.

Lemma

No permutation is even and odd.

Proof Sketch

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

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

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

We are using σ to permute the variables, whence $P = \pm P_\sigma$.
But then $P = +P_\sigma$ iff σ is even, and $P = -P_\sigma$ iff σ is odd.

□

The composition of even permutations is again even, so we can assemble them into a new group of size $n!/2$.

Definition

The collection of all even permutations of A, an n-element set, is the alternating group on n points.

Notation: $\mathfrak{A}_n \subseteq \mathfrak{S}_n$.

In one-line notation without brackets, $\mathfrak{A}_4$ has the following elements:

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

Alternating groups are important since, for $n \geq 5$, each alternating group $\mathfrak{A}_n$ is simple: it has only trivial normal subgroups.

So how do we actually compute in a group? In 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 $\{1\}$
- $n = 2$: $\mathbb{Z}_2$

1	a
a	1

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

1	a	b
a	b	1
b	1	a

- $\mathbb{Z}_4$

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 $n = 48$.

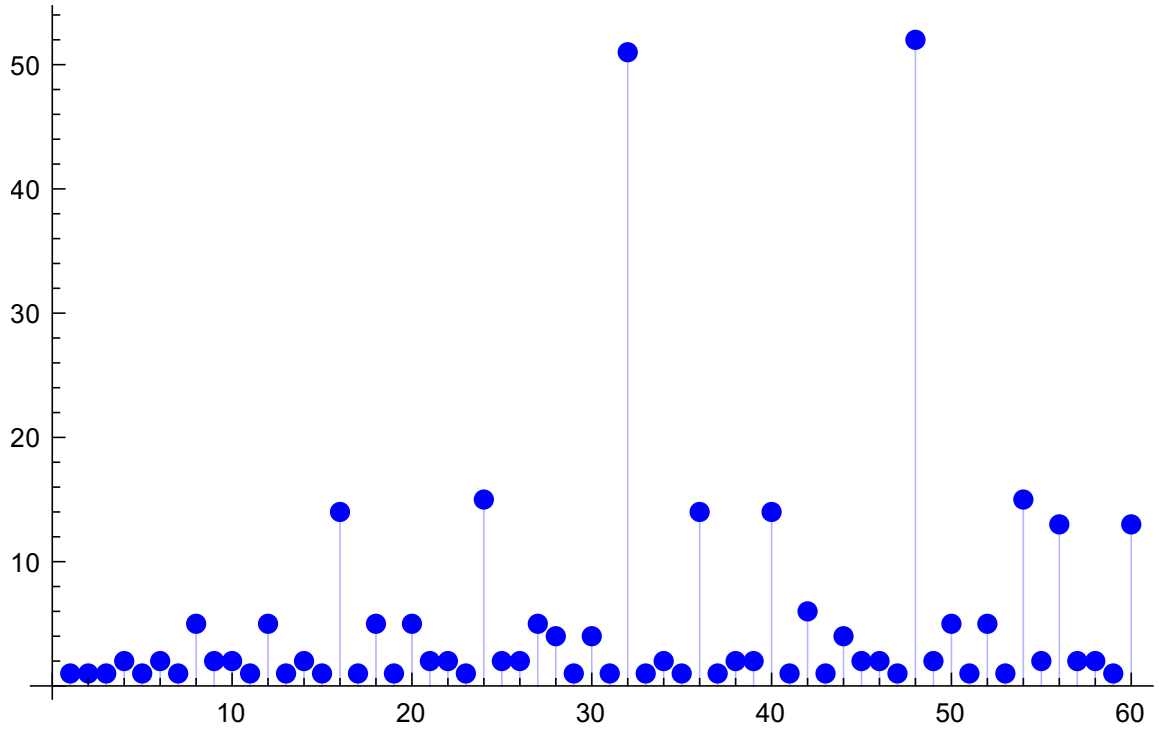
50
40
30
20
10
10
20
30
40
50
60

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

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

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 $\{ a^i \mid 0 \leq i < k \}$ is cyclic (with generator a).

Cyclic Groups are Boring

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.

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

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.
How about the top-right corner?

As we have already seen, the symmetries of a regular n-gon give rise to the dihedral groups D_n.

These groups have order $2n$ and are generated by a rotation a and a reflection b.

The basic identities are

$$a^n = b^2 = 1 \qquad ab = ba^{n-1}$$

Proposition

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

Proof.

Both groups have size 6, so there is a chance the claim might be correct.

In cycle notattion, the permutations $f = (1, 2)$ and $g = (1, 2, 3)$ 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.

□

Exercise

Check the details in the last argument.
Why does this argument not show that $\mathfrak{S}_n$ is isomorphic to D_n in general?

How about the symmetries of a rectangle?

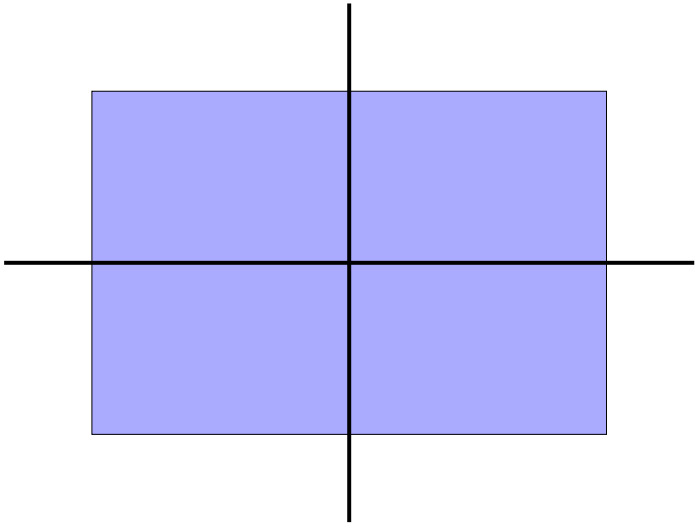

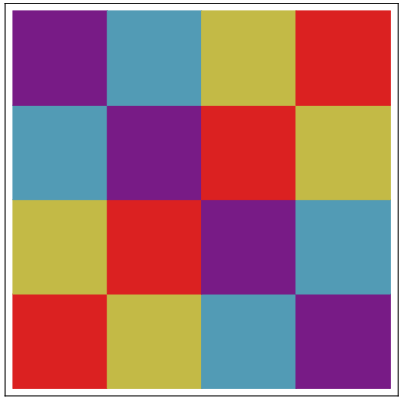

By visual inspection, there are only two reflections, say, a and b. Clearly, $a^2 = b^2 = 1$ and $ab = ba$, so the whole group is just

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

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):

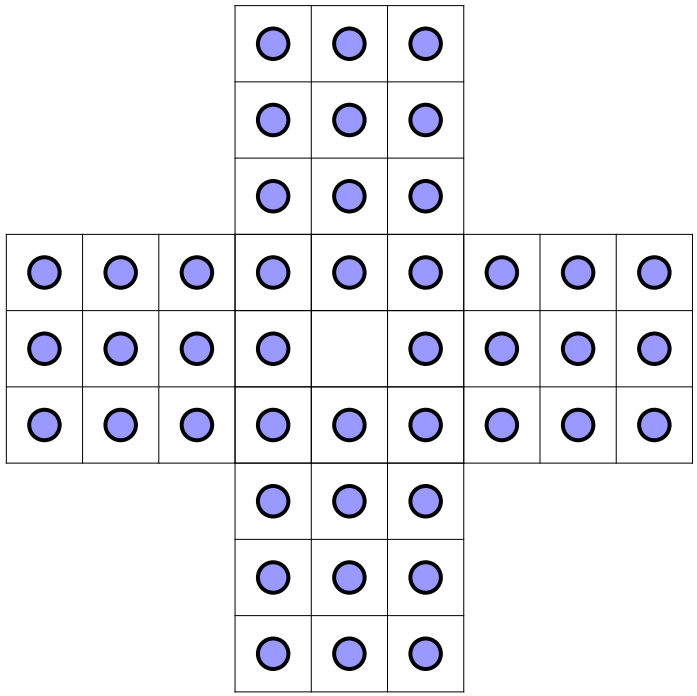

The goal is to "jump-over-and-remove-pebbles" until only one remains.

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; a strong solution leaves the last pebble in the center.

> **Challenge:** Show that any weak solution can be turned into a strong solution by changing at most 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 non-unit elements of the Kleinsche Vierergruppe, $v_s \in \{01, 10, 11\}$, as follows:

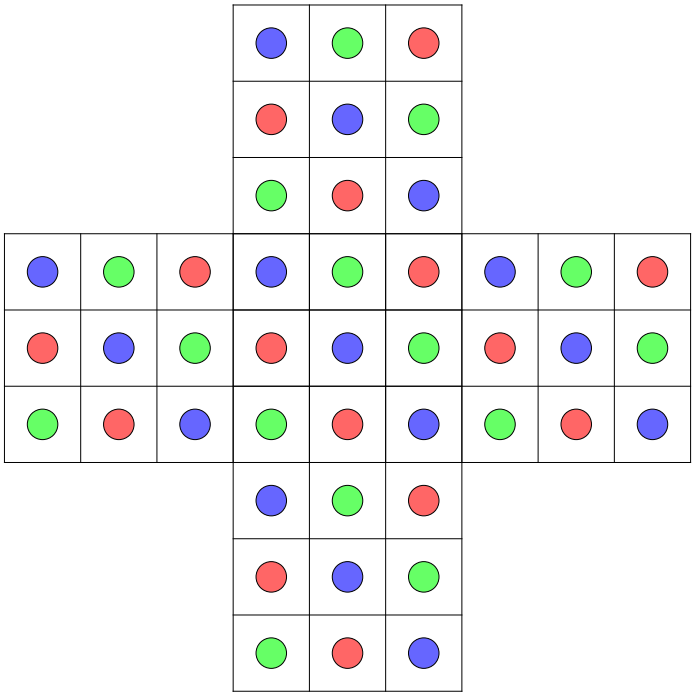

green $\rightsquigarrow 01$

red $\rightsquigarrow 10$

blue $\rightsquigarrow 11$

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

So What?

Indicate presence or absence of a pebble on square s by a Boolean variable $b_s \in \mathbf{2}$ and define the value of the corresponding configuration to be

$$v = \sum_s b_s v_s$$

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

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

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

Done!

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:

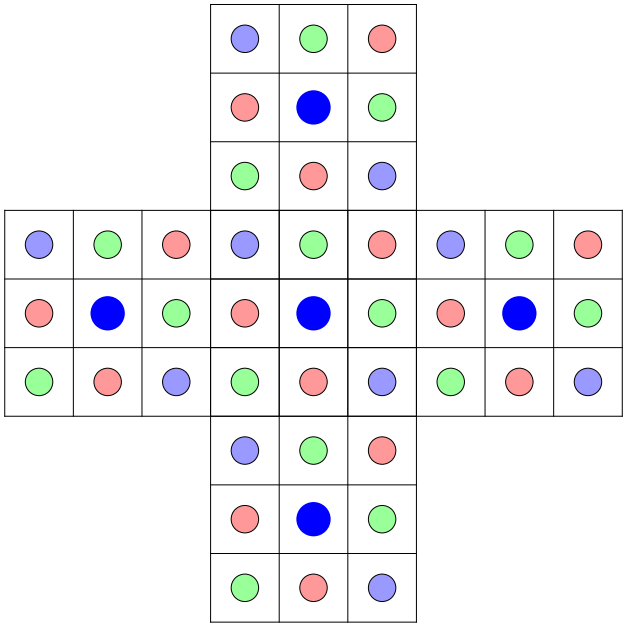

But each one of these can be reached iff all the others can: use symmetry, or change the last move.