

Section 1: Groups, intuitively

Matthew Macauley

Department of Mathematical Sciences
Clemson University
<http://www.math.clemson.edu/~macaule/>

Math 4120, Modern Algebra

A famous toy

Our introduction to group theory will begin by discussing the famous **Rubik's Cube**.



It was invented in 1974 by Ernő Rubik of Budapest, Hungary.

Ernő Rubik is a Hungarian inventor, sculptor and professor of architecture.

According to his Wikipedia entry:

He is known to be a very introverted and hardly accessible person, almost impossible to contact or get for autographs.

A famous toy

Not impossible ... just **almost** impossible.





Figure: June 2010, in Budapest, Hungary

A famous toy

- The cube comes out of the box in the **solved position**:



- But then we can scramble it up by consecutively rotating one of its 6 faces:



A famous toy

- The result might look something like this:



- The goal is to return the cube to its original solved position, again by consecutively rotating one of the 6 faces.

Since Rubik's Cube does not seem to require any skill with numbers to solve it, you may be inclined to think that this puzzle is not mathematical.

Big idea

Group theory is not primarily about numbers, but rather about **patterns** and **symmetry**; something the Rubik's Cube possesses in abundance.

A famous toy

Let's explore the Rubik's Cube in more detail.

In particular, let's identify some key features that will be recurring themes in our study of **patterns** and **symmetry**.

First, some questions to ponder:

- How did we scramble up the cube in the first place? How do we go about unscrambling the cube?
- In particular, what actions, or moves, do we *need* in order to scramble and unscramble the cube? (There are many correct answers.)
- How is Rubik's Cube different from checkers?
- How is Rubik's Cube different from poker?

Four key observations

Observation 1

There is a predefined list of moves that never changes.

Observation 2

Every move is reversible.

Observation 3

Every move is deterministic.

Observation 4

Moves can be combined in any sequence.

In this setting, a *move* is a twist of one of the six faces, by 0° , 90° , 180° , or 270° .

We could add more to our list, but as we shall see, these 4 observations are sufficient to describe the aspects of the mathematical objects that we wish to study.

What does group theory have to do with this?

Group theory studies the mathematical consequences of these 4 observations, which in turn will help us answer interesting questions about symmetrical objects.

Group theory arises everywhere! In puzzles, visual arts, music, nature, the physical and life sciences, computer science, cryptography, and of course, all throughout mathematics.

Group theory is one of the most beautiful subjects in all of mathematics!

Instead of considering our 4 observations as descriptions of Rubik's Cube, let's rephrase them as rules (axioms) that will define the boundaries of our objects of study.

Advantages of our endeavor:

1. We make it clear what it is we want to explore.
2. It helps us speak the same language, so that we know we are discussing the same ideas and common themes, though they may appear in vastly different settings.
3. The rules provide the groundwork for making logical deductions, so that we can discover new facts (many of which are surprising!).

Our informal definition of a group

Rule 1

There is a predefined list of **actions** that never changes.

Rule 2

Every action is reversible.

Rule 3

Every action is deterministic.

Rule 4

Any sequence of consecutive actions is also an action.

Definition (informal)

A **group** is a set of actions satisfying Rules 1–4.

Observations about the “Rubik’s Cube group”

Frequently, two sequences of moves will be “indistinguishable.” We will say that two such moves are *the same*. For example, rotating a face (by 90°) once has the same effect as rotating it five times.

Fact

There are 43,252,003,274,489,856,000 distinct configurations of the Rubik’s cube.

While there are infinitely many possible sequences of moves, starting from the solved position, there are 43,252,003,274,489,856,000 “truly distinct” moves.

All 4.3×10^{19} moves are **generated** by just 6 moves: a 90° clockwise twist of one of the 6 faces.

Let’s call these generators a , b , c , d , e , and f . Every **word** over the alphabet $\{a, b, c, d, e, f\}$ describes a unique configuration of the cube (starting from the solved position).

Summary of the big ideas

Loosing speaking a **group** is a **set of actions** satisfying some mild properties: deterministic, reversibility, and closure.

A **generating set** for a group is a **subcollection of actions** that together can produce all actions in the group – like a **spanning set** in a vector space.

Usually, a generating set is *much smaller* than the whole group.

Given a generating set, the individual actions are called **generators**.

The set of all possible ways to scramble a Rubik's cube is an example of a group. Two actions are the same if they have the *same “net effect”*, e.g., twisting a face 1 time vs. twisting a face 5 times.

Note that the group is the set of **actions** one can perform, not the set of configurations of the cube. However, there is a bijection between these two sets.

The Rubik's cube group has 4.3×10^{19} **actions** but we can find a **generating set** of size 6.

A road map for the Rubik's Cube

There are many solution techniques for the Rubik's Cube. If you do a Google search, you'll find several methods for solving the puzzle.

These methods describe a sequence of moves to apply relative to some starting position. In many situations, there may be a shorter sequence of moves that would get you to the solution.

In fact, it was shown in July 2010 that every configuration is **at most 20 moves** away from the solved position!

Let's pretend for a moment that we were interested in writing a complete solutions manual for the Rubik's Cube.

Let me be more specific about what I mean.

A road map for the Rubik's Cube

We'd like our solutions manual to have the following properties:

1. Given any scrambled configuration of the cube, there is a unique page in the manual corresponding to that configuration.
2. There is a method for looking up any particular configuration. (The details of how to do this are unimportant.)
3. Along with each configuration, a list of available moves is included. In each case, the page number for the outcome of each move is included, along with information about whether the corresponding move takes us closer to or farther from the solution.

Let's call our solutions manual the *Big Book*.

You are **15 steps** from the solution.

Cube front



Cube back

Face	Direction	Destination page	Progress
Front	Clockwise	36,131,793,510,312,058,964	Closer to solved
Front	Counterclockwise	12,374,790,983,135,543,959	Farther to solved
Back	Clockwise	26,852,265,690,987,257,727	Closer to solved
Back	Counterclockwise	41,528,397,002,624,663,056	Farther to solved
Left	Clockwise	6,250,961,334,888,779,935	Closer to solved
Left	Counterclockwise	10,986,196,967,552,472,974	Farther to solved
Right	Clockwise	26,342,598,151,967,155,423	Farther to solved
Right	Counterclockwise	40,126,637,877,673,696,987	Closer to solved
Top	Clockwise	35,131,793,510,312,058,964	Closer to solved
Top	Counterclockwise	33,478,478,689,143,786,858	Farther to solved
Bottom	Clockwise	20,625,256,145,628,342,363	Farther to solved
Bottom	Counterclockwise	7,978,947,168,773,308,005	Closer to solved

A road map for the Rubik's Cube

We can think of the *Big Book* as a road map for the Rubik's Cube. Each page says, “you are here” and “if you follow this road, you'll end up over there.”




Figure: Potential cover and alternative title for the *Big Book*

A road map for the Rubik's Cube

Unlike a vintage *Choose Your Own Adventure* book, you'll additionally know whether “over there” is where you want to go or not.

Pros of the *Big Book*:

- We can solve any scrambled Rubik's Cube.
- Given any configuration, every possible sequence of moves for solving the cube is listed in the book (long sequences and short sequences).
- The *Big Book* contains complete data on the moves in the Rubik's Cube universe and how they combine.

Cons of the *Big Book*:

- We just took all the fun out of the Rubik's Cube.
- If we had such a book, using it would be fairly cumbersome.
- We can't actually make such a book. Rubik's Cube has more than 4.3×10^{19} configurations. The paper required to write the book would cover the Earth many times over. The book would require over a billion terabytes of data to store electronically, and no computer in existence can store that much data.

What have we learned?

Despite the *Big Book*'s apparent shortcomings, it made for a good thought experiment.

The most important thing to get out of this discussion is that the *Big Book* is a **map of a group**.

We shall not abandon the mapmaking ideas introduced by our discussion of the *Big Book* simply because the map is too large.

We can use the same ideas to map out any group. In fact, we shall frequently do exactly that.

Let's try something simpler. . .

The Rectangle Puzzle group

- Consider a clear glass **rectangle** and label it as follows:

1	2
4	3

If you prefer, you can use colors instead of numbers:



We'll use numbers, and call the above configuration the **solved state** of our puzzle.

- The idea of the game is to scramble the puzzle and then find a way to return the rectangle to its solved state.
- Our “predefined list” consists of two actions: **horizontal flip** and **vertical flip**.


Loosely speaking, we allow these moves because they preserve the “footprint” of the rectangle. Do any other moves preserve its footprint?

Road map for The Rectangle Puzzle

For convenience, let's say that when we flip the rectangle, the numbers automatically become "right-side-up," as they would if you rotated an iPhone.

It is not hard to see that using only sequences of **horizontal** and **vertical** flips, we can obtain only four configurations.

Unlike the Rubik's cube group, the "road map" of the rectangle puzzle is small enough that we can draw it.



Observations? What sorts of things does the map tell us about the group?

Observations

Let G denote the rectangle group. This is a **set** of four actions. We see:

- G has 4 actions: the “identity” action e , a horizontal flip h , a vertical flip v , and a 180° rotation r .

$$G = \{e, h, v, r\}.$$

- We need two actions to “generate” G . In our diagram, each **generator** is represented by a different type (color) of arrow. We write:

$$G = \langle h, v \rangle.$$

- The map shows us how to get from any one configuration to any other. There is more than one way to follow the arrows! For example

$$r = hv = vh.$$

- For this particular group, the order of the actions is irrelevant! We call such a group **abelian**. Note that the Rubik’s cube group is *not* abelian.
- Every action in G is its own **inverse**: That is,

$$e = e^2 = h^2 = v^2 = r^2.$$

The Rubik’s cube group does **not** have this property. Algebraically, we write:


$$e^{-1} = e, \quad v^{-1} = v, \quad h^{-1} = h, \quad r^{-1} = r.$$

An alternative set of generators for the Retangle Puzzle

The rectangle puzzle can also be generated by a **horizontal flip** and a **180° rotation**:

$$G = \langle h, r \rangle.$$

Let's build a Cayley graph using this alternative set of generators.



Do you see this road map has the “same structure” as our first one? Of course, we need to “untangle it” first.

Perhaps surprisingly, this might *not* always be the case.

That is, there are (more complicated) groups for which different generating sets yield road maps that are structurally different. We'll see examples of this shortly.

Cayley diagrams

As we saw in the previous example, how we choose to layout our map is irrelevant.

What is important is that the connections between the various states are preserved.

However, we will attempt to construct our maps in a pleasing to the eye and symmetrical way.


The official name of the type of group road map that we have just created is a **Cayley diagram**, named after 19th century British mathematician Arthur Cayley.

In general, a Cayley diagram consists of **nodes** that are connected by colored (or labeled) **arrows**, where:



- an **arrow** of a particular color represents a specific **generator**;
- each **action** of the group is represented by a unique **node** (sometimes we will label nodes by the corresponding action).
- Equivalently, each **action** is represented by a (non-unique) **path** starting from the **solved state**.

More on arrows


- An arrow corresponding to the generator g from node x to node y means that node y is the result of applying the action $g \in G$ to node x :



- If an action $h \in G$ is its own inverse (that is, $h^2 = e$), then we have a 2-way arrow. This happens with **horizontal** and **vertical** flips. For clarity, our convention is to drop the tips on all 2-way arrows. Thus, these are exactly the same:




When we want to focus on a group's structure, we frequently omit the labels at the nodes. Thus, the Cayley diagram of the rectangle puzzle can be drawn as follows:



The 2-Light Switch group

Let's map out another group, which we'll call the *2-Light Switch Group*:

- Consider two light switches side by side that both start in the off position (This is our “solved state”).
- We are allowed 2 actions: **flip L switch** and **flip R switch**.





Notice how the Cayley diagrams for the Rectangle Puzzle $G = \{e, v, h, r\}$ and the 2-Light Switch Group $G' = \{e, L, R, B\}$ are essentially the same.

Although these groups are superficially different, the Cayley diagrams help us see that *they have the same structure*. (The fancy phrase for this property is that the “two groups are *isomorphic*”; more on this later.)

The Klein 4-group

Any group with the same Cayley diagram as the Rectangle Puzzle and the 2-Light Switch Group is called the **Klein 4-group**, denoted by V_4 for *vierergruppe*, “four-group” in German. It is named after the mathematician Felix Klein.

It is important to point out that the number of different types (i.e., colors) of arrows matters. For example, the Cayley diagram on the right *does not* represent V_4 .



Questions

- What group has a Cayley graph like the diagram on the right?
- How would you give a proof (=convincing argument) that these two groups have truly different structures? Can you find a property that one group has that the other does not?
- Can you find another group of size 4 that is different from both of these?

The Triangle Puzzle group


Let's play our “rectangle puzzle” game but with an equilateral triangle:



The “triangle puzzle” group, often denoted D_3 , has 6 actions:

- The identity action: e
- A (clockwise) 120° rotation: r
- A (clockwise) 240° rotation: r^2
- A horizontal flip: f
- Rotate + horizontal flip: rf
- Rotate twice + horizontal flip: r^2f .

One set of generators: $D_3 = \langle r, f \rangle$.



Notice that multiple paths can lead us to the same node. These give us **relations** in our group. For example:

$$r^3 = e, \quad r^{-1} = r^2, \quad f^{-1} = f, \quad rf = fr^2, \quad r^2f = fr.$$

This group is **non-abelian**: $rf \neq fr$.

Properties of Cayley graphs

Observe that at every node of a Cayley graph, there is exactly one out-going edge of each color.


Question 1

Can an edge in a Cayley graph ever connect a node to itself?

Question 2

Suppose we have an edge corresponding to generator g that connects a node x to itself. Does that mean that the edge g connects every node to itself? In other words, can an action be the *identity action* when applied to some actions (or configurations) but not to others?

Visually, we're asking if the following scenerio can ever occur in a Cayley diagram:



A Theorem and Proof!

Perhaps surprisingly, the previous situation is *impossible*! Let's properly **formulate** and **prove** this.

Theorem

Suppose an action g has the property that $gx = x$ for some other action x . Then g is the *identity action*, i.e., $gh = h = hg$ for all other actions h .

Proof

The identity action (we'll denote by 1) is simply the action hh^{-1} , for any action h .

If $gx = x$, then multiplying by x^{-1} *on the right* yields:

$$g = gxx^{-1} = xx^{-1} = 1.$$

Thus g is the identity action. □

This was our first mathematical proof! It shows how we can deduce interesting properties about groups *from* the rules, which were not explicitly *built into* the rules.

Applications of groups

Thus far, we have introduced groups and explored a few basic examples.

At this point, we will pause to discuss a few practical (but not complicated) applications of groups.

We will see applications of group theory in 3 areas:

1. Science
2. Art
3. Mathematics

Our choice of examples is influenced by how well they illustrate the material rather than how useful they are.

Groups of symmetries

Intuitively, something is symmetrical when it looks the same from more than one point of view.

Can you think of an object that exhibits symmetry? Have we already seen some?

How does symmetry relate to groups?

The examples of groups that we've seen so far deal with arrangements of similar things.

In the next section, we will uncover the following fact (we'll be more precise later):

Cayley's Theorem

Every group can be viewed as a collection of ways to rearrange some set of things.

How to make a group out of symmetries

Groups relate to symmetry because an object's symmetries can be described using arrangements of the object's parts.

The following algorithm tells us how to construct a group that describes (or measures) a physical object's symmetry.

Algorithm

1. Identify all the parts of the object that are similar (e.g., the corners of an n -gon), and give each such part a different number.
2. Consider the actions that may rearrange the numbered parts, but leave the object in the same physical space. (This collection of actions forms a group.)
3. (Optional) If you want to visualize the group, explore and map it as we did in the previous lecture with the rectangle puzzle, etc.

Comments


- We'll refer to the physical space that an object occupies as its **footprint** (this terminology does not appear in the text).
- Step 1 of our Algorithm numbers the object's parts so that we can track the manipulations permitted in Step 2. Each new state is a rearrangement of the object's similar parts and allows us to distinguish each of these rearrangements; otherwise we could not tell them apart.
- Not every rearrangement is valid. We are only allowed actions that maintain the object's physical integrity *and* preserve its footprint. For example, we can't rip two arms off a starfish and glue them back on in different places.
- Step 2 requires us to find *all* actions that preserve the object's footprint and physical integrity; not just the generators.
- However, if we choose to complete Step 3 (make a Cayley diagram), we must make a choice concerning generators. Different choices in generators may result in different Cayley diagrams.
- When selecting a set of generators, we would ideally like to select as small a set as possible. We can never choose too many generators, but we can choose too few. However, having "extra" generators only clutters our Cayley diagram.

Shapes of molecules

Because the shape of molecules impacts their behavior, chemists use group theory to classify their shapes. Let's look at an example.

The following figure depicts a molecule of Boric acid, $\text{B}(\text{OH})_3$.

Note that a mirror reflection is *not* a symmetry of this molecule.




Exercise

Follow the steps of our Algorithm to find the group that describes the symmetry of the molecule and draw a possible Cayley diagram.

The group of symmetries of Boric acid has 3 actions requiring at least one generator. If we choose “120° clockwise rotation” as our generator, then the actions are:


1. the *identity* (or “do nothing”) action: e
2. 120° clockwise rotation: r
3. 240° clockwise rotation: r^2 .



This is the **cyclic group**, C_3 . (We'll discuss cyclic groups in a later lecture.)

Crystallography

Solids whose atoms arrange themselves in a regular, repeating pattern are called **crystals**. The study of crystals is called **crystallography**.




When chemists study such crystals they treat them as patterns that repeat without end. This allows a new manipulation that preserves the infinite footprint of the crystal and its physical integrity: **translation**.

In this case, the groups describing the symmetry of crystals are infinite. Why?

Frieze patterns

Crystals are patterns that repeat in 3 dimensions. Patterns that only repeat in 1 dimension are called **frieze patterns**. The groups that describe their symmetries are called **frieze groups**.

Frieze patterns (or at least finite sections of them) occur throughout art and architecture. Here is an example:



This frieze admits a new type of manipulation that preserves its footprint and physical integrity: a **glide reflection**. This action consists of a horizontal translation (by the appropriate amount) followed by a vertical flip.

Note that for this pattern, a vertical flip all by itself does not preserve the footprint, and thus is not one of the actions of the group of symmetries.


Exercise

Determine the group of symmetries of this frieze pattern and draw a possible Cayley diagram.

Frieze patterns

The group of symmetries of the frieze pattern on the previous slide turns out to be infinite, but we only needed two generators: **horizontal flip** and **glide reflection**.



Here is a possible Cayley diagram:



Friezes, wallpapers, and crystals

- The symmetry of any frieze pattern can be described by one of 7 different infinite groups. Some **frieze groups** are *isomorphic* (have the same structure) even though the visual appearance of the patterns (and Cayley graphs) may differ.
- The symmetry of 2-dimensional repeating patterns, called “wallpaper patterns,” has also been classified. There are 17 different **wallpaper groups**.
- There are 230 **crystallographic groups**, which describe the symmetries of 3-dimensional repeating patterns.

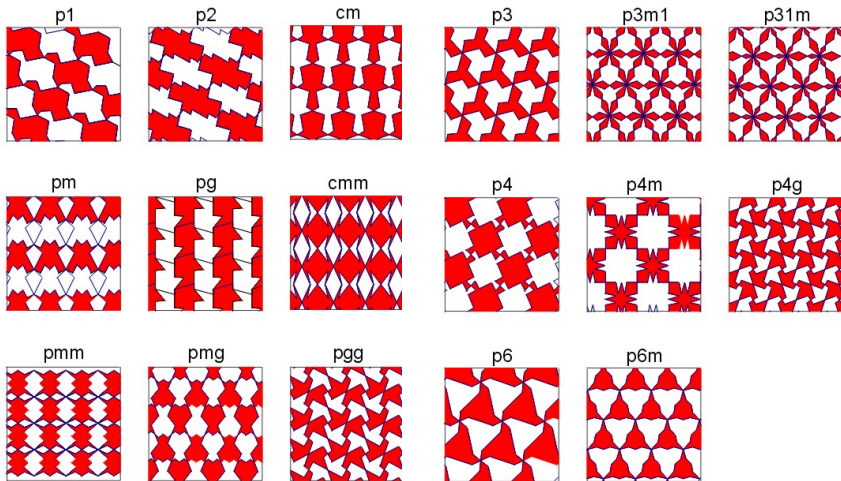
The 7 types of frieze patterns



Questions

- What basic types of symmetries (e.g., translation, reflection, rotation, glide reflection) do these frieze groups have?
- What are the (minimal) generators for the corresponding frieze groups?
- Which of these frieze patterns have isomorphic frieze groups?
- Which of these frieze groups are abelian?


The 17 types of wallpaper patterns



Images courtesy of Patrick Morandi (New Mexico State University).

The 17 types of wallpaper patterns


Here is another picture of all 17 wallpapers, with the official **IUC notation** for the symmetry group, adopted by the International Union of Crystallography in 1952.



Braid groups

Another area where groups arise in both art and mathematics is the study of **braids**.

This is best seen by an example. The following is a picture of an element (action) from the **braid group** $B_4 = \langle \sigma_1, \sigma_2, \sigma_3 \rangle$:



The braid $b = \sigma_1 \sigma_2 \sigma_1 \sigma_3^{-1} \sigma_1^{-1} \sigma_2 = \sigma_1 \sigma_2 \sigma_3^{-1} \sigma_2$.

Do you see why the set of braids on n strings forms a group?


To combine two braids, just concatenate them.

Every braid is reversible – just “undo” each crossing. In the example above,

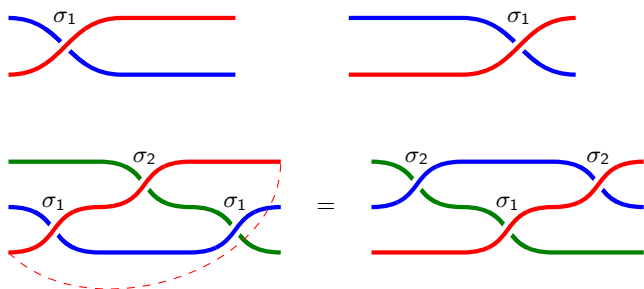
$$e = bb^{-1} = (\sigma_1 \sigma_2 \sigma_1 \sigma_3^{-1} \sigma_1^{-1} \sigma_2)(\sigma_2^{-1} \sigma_1 \sigma_3 \sigma_1^{-1} \sigma_2^{-1} \sigma_1^{-1}).$$

Braid groups

There are two fundamental **relations** in braid groups:


$$\sigma_i \sigma_j = \sigma_j \sigma_i$$

(if $|i - j| \geq 2$)


$$\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$$

We can describe the braid group B_4 by the following **presentation**:

$$B_4 = \langle \sigma_1, \sigma_2, \sigma_3 \mid \sigma_1 \sigma_3 = \sigma_3 \sigma_1, \sigma_1 \sigma_2 \sigma_1 = \sigma_2 \sigma_1 \sigma_2, \sigma_2 \sigma_3 \sigma_2 = \sigma_3 \sigma_2 \sigma_3 \rangle.$$

We will study presentations more in the next lecture; this is just an introduction.

Labeled Cayley diagrams

Recall that **arrows** in a Cayley diagram represent one choice of **generators** of the group. In particular, all arrows of a fixed color correspond to the same generator.

Our choice of generators influenced the resulting Cayley diagram!

When we have been drawing Cayley diagrams, we have been doing one of two things with the nodes:

1. Labeling the nodes with **configurations** of a thing we are acting on.
2. Leaving the nodes unlabeled (this is the “abstract Cayley diagram”).

There is a 3rd thing we can do with the nodes, motivated by the fact that every **path** in the Cayley diagram represents an **action** of the group:



3. Label the nodes with **actions** (this is called a “diagram of actions”).

Motivating idea


If we distinguish one node as the “unscrambled” configuration and label that with the **identity action**, then we can label each remaining node with the action that it takes to reach it from the unscrambled state.

An example: The Klein 4-group

Recall the “rectangle puzzle.” If we use **horizontal flip** (h) and **vertical flip** (v) as generators, then here is the Cayley diagram labeled by configurations (left), and unlabeled Cayley diagram (right):



Let's apply the steps to the abstract Cayley diagram for V_4 , using the upper-left node as the “unscrambled configuration”:



Note that we could also have labeled the node in the lower right corner as hv , as well.

How to label nodes with actions

Let's summarize the process that we just did.

Node labeling algorithm

The following steps transform a Cayley diagram into one that focuses on the group's actions.

- (i) Choose a node as our initial reference point; label it e . (This will correspond to our “identity action.”)
- (ii) Relabel each remaining node in the diagram with a path that leads there from node e . (If there is more than one path, pick any one; shorter is better.)
- (iii) Distinguish arrows of the same type in some way (color them, label them, dashed vs. solid, etc.)

Our convention will be to label the nodes with sequences of generators, so that reading the sequence from **left to right** indicates the appropriate path.

Warning!


Some authors use the opposite convention, motivated by “function composition.”

A “group calculator”

One neat thing about Cayley diagrams with nodes labeled by actions is that they act as a “group calculator”.

For example, if we want to know what a particular sequence is equal to, we can just chase the sequence through the Cayley graph, starting at e .

Let's try one. In V_4 , what is the action $hhhvhvvhv$ equal to?



We see that $hhhvhvvhv = h$. A more condensed way to write this is

$$hhhvhvvhv = h^3vhv^2hv = h.$$


A concise way to describe V_4 is by the following **group presentation** (more on this later):

$$V_4 = \langle v, h \mid v^2 = e, h^2 = e, vh = hv \rangle.$$



Another familiar example: D_3

Recall the “triangle puzzle” group $G = \langle r, f \rangle$, generated by a clockwise 120° rotation r , and a horizontal flip f .

Let’s take the shaded triangle to be the “unscrambled configuration.”



Here are two different ways (of many!) that we can label the nodes with actions:



The following is one (of many!) presentations for this group:

$$D_3 = \langle r, f \mid r^3 = e, f^2 = e, r^2f = fr \rangle.$$

Group presentations

Initially, we wrote $G = \langle h, v \rangle$ to say that “ G is generated by the elements h and v .”

All this tells us is that h and v **generate** G , but not **how** they generate G .

If we want to be more precise, we use a **group presentation** of the following form:

$$G = \langle \text{generators} \mid \text{relations} \rangle$$

The vertical bar can be thought of as meaning “subject to”.

For example, the following is a presentation for V_4 :

$$V_4 = \langle a, b \mid a^2 = e, b^2 = e, ab = ba \rangle.$$

Caveat!


Just because there are elements in a group that “satisfy” the relations above does *not* mean that it is V_4 .

For example, the trivial group $G = \{e\}$ satisfies the above presentation; just take $a = e$ and $b = e$.

Loosely speaking, the above presentation tells us that V_4 is the “**largest group**” that satisfies these relations. (More on this when we study quotients.)

Group presentations

Recall the frieze group from the previous lecture that had the following Cayley diagram:



One presentation of this group is

$$G = \langle t, f \mid f^2 = e, tft = f \rangle.$$

Here is the Cayley diagram of another frieze group:



It has presentation

$$G = \langle a \mid \rangle.$$

That is, “one generator subject to *no relations*.”

Group presentations

Due to the aforementioned caveat, and a few other technicalities, the study of group presentations is a topic usually relegated to graduate-level algebra classes.

However, they are often introduced in an undergraduate algebra class because *they are very useful*, even if the intricate details are harmlessly swept under the rug.

The problem (called the **word problem**) of determining what a mystery group is from a presentation is actually **computationally unsolvable**! In fact, it is equivalent to the famous “halting problem” in computer science!

For (mostly) amusement, what group do you think the following presentation describes?

$$G = \langle a, b \mid ab = b^2a, ba = a^2b \rangle.$$

Surprisingly, this is the trivial group $G = \{e\}$!

Inverses

If g is a generator in a group G , then following the “ g -arrow” backwards is an action that we call its **inverse**, and denoted by g^{-1} .


More generally, if g is represented by a **path** in a Cayley diagram, then g^{-1} is the action achieved by tracing out this path in reverse.

Note that by construction,

$$gg^{-1} = g^{-1}g = e,$$

where e is the **identity** (or “do nothing”) action. Sometimes this is denoted by e , 1 , or 0 .

For example, let's use the following Cayley diagram to compute the inverses of a few actions:



$$r^{-1} = \text{_____} \text{ because } r\text{_____} = e = \text{_____}r$$

$$f^{-1} = \text{_____} \text{ because } f\text{_____} = e = \text{_____}f$$

$$(rf)^{-1} = \text{_____} \text{ because } (rf)\text{_____} = e = \text{_____}(rf)$$


$$(r^2f)^{-1} = \text{_____} \text{ because } (r^2f)\text{_____} = e = \text{_____}(r^2f).$$

Multiplication tables

Since we can use a Cayley diagram with nodes labeled by actions as a “group calculator,” we can create a (group) multiplication table, that shows how every pair of group actions combine.

This is best illustrated by diving in and doing an example. Let's fill out a multiplication table for V_4 .

Since order of multiplication can matter, let's use the convention that the entry in row g and column h is the element gh (rather than hg).



	e	v	h	r
e	e	v	h	r
v	v	e	r	h
h	h	r	e	v
r	r	h	v	e

Some remarks on the structure of multiplication tables

Comments

- The 1st column and 1st row repeat themselves. (Why?) Sometimes these will be omitted (*Group Explorer* does this).
- Multiplication tables can visually reveal patterns that may be difficult to see otherwise. To help make these patterns more obvious, we can color the cells of the multiplication table, assigning a unique color to each action of the group.
- A group is abelian iff its multiplication table is symmetric about the “main diagonal.”
- In each row and each column, each group action occurs exactly once. (This will always happen. . . Why?)

Let's state and prove that last comment as a theorem.

A theorem and proof

Theorem

An element cannot appear twice in the same **row** or **column** of a multiplication table.

Proof

Suppose that in **row** a , the element g appears in columns b and c . Algebraically, this means

$$ab = g = ac.$$

Multiplying everything on the **left** by a^{-1} yields

$$a^{-1}ab = a^{-1}g = a^{-1}ac \implies b = c.$$


Thus, g (or any element) cannot appear twice in the same **row**.

The proof that two elements cannot appear twice in the same **column** is similar, and will be left as a homework exercise. □

Another example: D_3

Let's fill out a multiplication table for the group D_3 ; here are several different presentations:

$$\begin{aligned} D_3 &= \langle r, f \mid r^3 = e, f^2 = e, rf = fr^2 \rangle \\ &= \langle r, f \mid r^3 = e, f^2 = e, rfr = f \rangle. \end{aligned}$$





	e	r	r^2	f	rf	r^2f
e	e	r	r^2	f	rf	r^2f
r	r	r^2	e	rf	r^2f	f
r^2	r^2	e	r	r^2f	f	rf
f	f	r^2f	rf	e	r^2	r
rf	rf	f	r^2f	r	e	r^2
r^2f	r^2f	rf	f	r^2	r	e

Observations? What patterns do you see?

Just for fun, what group do you get if you remove the " $r^3 = e$ " relation from the presentations above? (*Hint*: We've seen it recently!)

Another example: the quaternion group

The following Cayley diagram, laid out two different ways, describes a group of size 8 called the **Quaternion group**, often denoted $Q_8 = \{\pm 1, \pm i, \pm j, \pm k\}$.



The “numbers” j and k individually act like $i = \sqrt{-1}$, because $i^2 = j^2 = k^2 = -1$.

Multiplication of $\{\pm i, \pm j, \pm k\}$ works like the cross product of unit vectors in \mathbb{R}^3 :

$$ij = k, \quad jk = i, \quad ki = j, \quad ji = -k, \quad kj = -i, \quad ik = -j.$$

Here are two possible presentations for this group:

$$\begin{aligned} Q_8 &= \langle i, j, k \mid i^2 = j^2 = k^2 = ijk = -1 \rangle \\ &= \langle i, j \mid i^4 = j^4 = 1, \, iji = j \rangle. \end{aligned}$$

Moving towards the standard definition of a group

We have been calling the members that make up a group “actions” because our definition requires a group to be a collection of actions that satisfy our 4 rules.

Since the standard definition of a group is not phrased in terms of actions, we will need more general terminology.

We will call the members of a group **elements**. In general, a group is a **set of elements** satisfying some set of properties.

We will also use standard set theory notation. For example, we will write things like

$$h \in V_4$$

to mean “ h is an element of the group V_4 .”

Binary operations

Intuitively, an **operation** is a method for combining objects. For example, $+$, $-$, \cdot , and \div are all examples of operations. In fact, these are **binary operations** because they combine two objects into a single object.

Definition

If $*$ is a **binary operation** on a set S , then $s * t \in S$ for all $s, t \in S$. In this case, we say that S is **closed** under the operation $*$.

Combining, or “multiplying” two group elements (i.e., doing one action followed by the other) is a binary operation. We say that it is a binary operation *on* the group.

Recall that Rule 4 says that any sequence of actions is an action. This ensures that the group is closed under the binary operation of multiplication.

Multiplication tables are nice because they depict the group’s binary operation in full.

However, not every table with symbols in it is going to be the multiplication table for a group.

Associativity

Recall that an operation is **associative** if parentheses are **permitted anywhere, but required nowhere**.

For example, ordinary addition and multiplication are associative. However, subtraction of integers is *not* associative:

$$4 - (1 - 2) \neq (4 - 1) - 2.$$

Is the operation of combining actions in a group associative? YES! We will not prove this fact, but rather illustrate it with an example.

Recall D_3 , the group of symmetries for the equilateral triangle, generated by r (=rotate) and f (=horizontal flip).

How do the following compare?

$$rfr, \quad (rf)r, \quad r(fr)$$

Even though we are associating differently, the end result is that *the actions are applied left to right*.

The moral is that we never need parentheses when working with groups, though we may use them to draw our attention to a particular chunk in a sequence.

Classical definition of a group

We are now ready to state the standard definition of a group.

Definition (official)

A set G is a **group** if the following criteria are satisfied:

1. There is a **binary operation** $*$ on G .
2. $*$ is associative.
3. There is an **identity** element $e \in G$. That is, $e * g = g = g * e$ for all $g \in G$.
4. Every element $g \in G$ has an **inverse**, g^{-1} , satisfying $g * g^{-1} = e = g^{-1} * g$.

Remarks

- Depending on context, the binary operation may be denoted by $*$, \cdot , $+$, or \circ .
- As with ordinary multiplication, we frequently omit the symbol altogether and write, e.g., xy for $x * y$.
- We generally only use the $+$ symbol if the group is abelian. Thus, $g + h = h + g$ (always), but in general, $gh \neq hg$.
- Uniqueness of the identity and inverses is *not* built into the definition of a group. However, we can without much trouble, prove these properties.

Definitions of a group: Old vs. New

Do our two competing definitions agree? That is, if our informal definition says something is a group, will our official definition agree? Or vice versa?

Since our first definition of a group was informal, it is impossible to answer this question officially and absolutely. An informal definition potentially allows some technicalities and ambiguities.

This aside, our discussion leading up to our official Definition provides an informal argument for why the answer to the first question should be yes. We will answer the second question in the next chapter.

Regardless of whether the definitions agree, we always have $e^{-1} = e$. That is, the inverse of doing nothing is doing nothing.

Even though we haven't officially shown that the two definitions agree (and in some sense, we can't), we shall begin viewing groups from these two different paradigms:

- a group as a **collection of actions**;
- a group as a **set with a binary operation**.

A few simple properties

One of the first things we can prove about groups is uniqueness of the identity and inverses.

Theorem

Every element of a group has a *unique* inverse.

Proof

Let g be an element of a group G . By definition, it has at least one inverse.

Suppose that h and k are both inverses of g . This means that $gh = hg = e$ and $gk = kg = e$. It suffices to show that $h = k$. Indeed,

$$h = he = h(gk) = (hg)k = ek = k,$$

and the proof is complete. □

The following proof is relegated to the homework; the technique is similar.

Theorem

Every group has a *unique* identity element.