INTRODUCTION TO ALGEBRAIC STRUCTURES

Tom Denton
Google Research

LibreTexts™

Introduction to Algebraic Structures

Tom Denton

The LibreTexts libraries are Powered by NICE CXOne and are supported by the Department of Education Open Textbook Pilot Project, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

This text was compiled on 02/01/2024

TABLE OF CONTENTS

9: Vector Spaces

Index

Glossary

Detailed Licensing

Licensing

A detailed breakdown of this resource's licensing can be found in ***Back Matter/Detailed Licensing****.*

Preface

This is a set of notes I developed for an e-learning course in Algebraic Structures offered by Maseno, University in Western Kenya. The idea is to introduce the key concepts of algebraic structures without assuming much background in higher mathematics. Math education in Kenya is heavy on calculation (it's relatively easier to teach and evaluate), but often falls short when it comes to teaching students to think creatively about mathematics, and really understand the subject as it relates to the world beyond the test. On the bright side, these same students are usually very ready to take a more creative approach to mathematics: good skills in calculation provides at least a good intuition for working with numbers, and gives a good foundation from which to build. Kenyan students are also generally very enthusiastic when presented with interesting mathematics.

The notes are trying, then, to accomplish the following:

1. Give students a first encounter with algebraic structures: Groups, rings, fields, and vector spaces,
2. Create an intuition for how these objects appear 'in the world,' meaning both in the real world and in the broader scope of mathematics,
3. Encourage students to engage with the material in a creative way, and
4. Teach/Reinforce important points from the foundations of mathematics, such as induction.

It's a lot to ask for a single ten-week term. Let's see where we get.

The notes themselves are divided into eleven 'chapters,' one for each week of Maseno's term, plus this introductory chapter. Taking a cue from computer science, all numbering of chapters and sections starts at 0. As the course becomes fully developed, I will be inserting videos for each section, giving an alternate presentation of the ideas. But the text is primary!

Here are some underlying principles that I believe strongly in, which also guide the formation of these notes.

1. We live in the future. Computers are somewhere between a million and a billion times faster at computation than humans are. Therefore, we should focus our teaching on what humans do better than computers: Understanding, problem solving, and placing things in context. It is often essential to understand how to compute things (indeed, otherwise we would not be able to tell the computer how to do computations for us!), but computation should not be the aim of a course.
2. We live in the future. We can communicate at almost zero-cost at slightly less than the speed of light. Information is governed by post-scarcity economics, and we need to treat it as such. This means we cannot treat information like a scarce resource to be hoarded: we must share our infinite wealth freely. Thus, these notes will remain free, and will be distributed under the Gnu Public License.

Design Principles

This book is also a programming project! As of this writing, I'm learning some modern web-programming tools; this book runs on Django, HTML5, Javascript, JQuery, MathJAX, the Sage Cell Server, and probably more by the time I'm done. HTML5 support is becoming more common in browsers, and should be an available standard for a long time to come.

Here is a list of design principles that I hope to adhere to for the final product:

1. An important principle of the book is to support multiple learning modes: there should be a combination of video and text for every section.
2. Videos, for their part, should be no longer than five or ten minutes long. Likewise, sections of the text should be somewhere south of 1000 words.
3. Each section should have at least one exercise, and these exercises should encourage both basic mechanical understanding and encourage creative approaches to the material.
4. The finished work should be free and freely available.
5. The finished work should meet the standards of a Maseno University e-learning course; in particular, have at least ten 'topics' to be digested at a rate of one-per-week, with clearly marked exercises to act as assignments.
6. Wherever reasonable, interactive elements (probably using Sage) should be included.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

1: Symmetry

Groups arise in nature whenever we can find symmetry. For example, the human body has a *lateral symmetry:* if you imagine reversing left and right, most people would look more-or-less the same.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

1.1: Symmetry

Groups arise in nature whenever we can find symmetry. For example, the human body has a *lateral symmetry:* if you imagine reversing left and right, most people would look more-or-less the same. (In fact, we see an example of this every time we look in a mirror.)

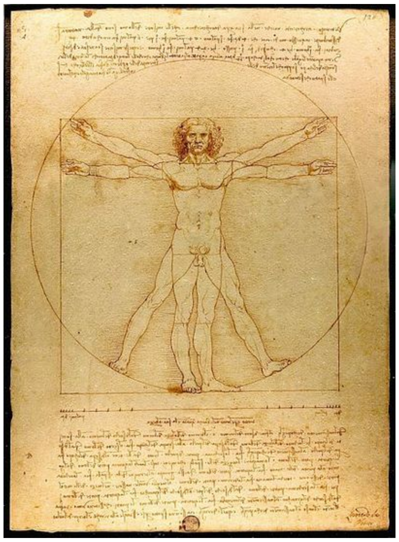

Figure 1: Da Vinci's famous sketch demonstrates lateral symmetry in the human body. (Source)

Another example is in the formation of crystals. In a crystal, the atomic structure arranges itself into a very symmetrical pattern, which you can see even with the unaided eye. The symmetry of the atomic structure means the atoms are packed very regularly, which leads to the nice shapes we see. In the late 1800's, mathematicians used group theory to classify all of the shapes of crystals that could ever exist in the world.

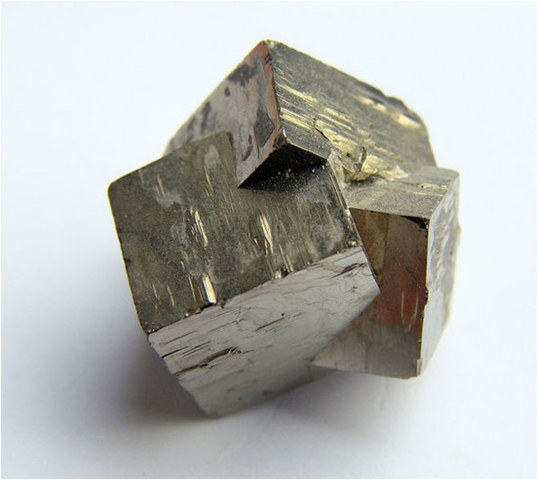

Figure 2: Twinned pyrite crystal. (Source)

Tiling patterns -- two-dimensional pictures that repeat in regular ways -- are also an example of symmetry. Many cultures have explored symmetry through tiling patterns, though the study of tiling was especially refined in the Islamic cultures during the middle ages. Bans on the depiction of human forms led Islamic artists to very deep explorations of abstract designs, with a strong emphasis on tiling. Group theory can also classify all of the possible regular tiling patterns, also known as tessellations, allowing us to verify that all of the possible tiling patterns were actually discovered by Islamic artists! There are many, many interesting tiling patterns out there: here's a place to start reading about them if you're interested.

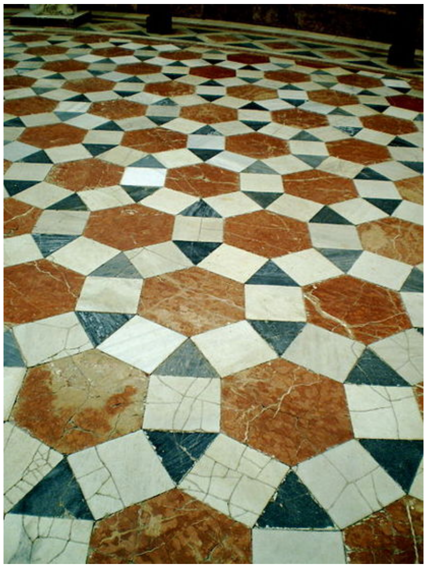

Figure 3: Tiling patterns exhibit interesting symmetries! (Source)

A very deep example of symmetry occurs in our most fundamental assumptions in physics. A basic principle states that what matters in a physical system is the relationship between all of the objects, not their absolute position. This principle allows us to say that the physics we figure out on Earth should work the same on Mars. (Thus far in the history of the world, it's been a very useful assumption.) So if you move the entire system in any direction (or rotate it or reflect it), the system will not 'notice.' This invariance is also a kind of symmetry. Since you can move a physical system by any amount without changing it, the physical system has an infinite number of symmetries!

So what is a group? How can we create mathematics to encompass the study of symmetry? This is what we'll explore in this chapter.

✓ Example 1.1.1:

Find some examples symmetry other than those that we talked about above.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

1.2: Counting Symmetries

So how can we use mathematics to study symmetry? Well, the first thing we learn about in mathematics is counting, so perhaps we should try to count symmetries!

If we think of a perfectly symmetrical face, there are two symmetries: one from flipping left-to-right, and another from leaving the face alone. Some might argue that the face only has one symmetry, and leaving it alone doesn't count. In fact, we could argue all day about this point and make no progress until we became very precise about what is meant by 'a symmetry.'

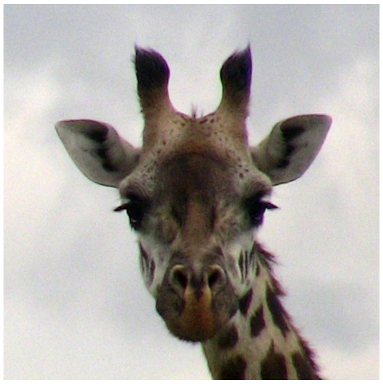

An almost perfectly symmetrical face.

To settle the argument, we require a definition! This first definition of the text is intentionally very loose.

Definition 1.1.0: Symmetry

A *symmetry* of an object is a way of moving the object back onto itself without changing it.

In fact, doing nothing to an object is a way of moving it back onto itself. Thus, we will say that a symmetrical face has two symmetries.

Let's consider some more mathematical objects. A line segment always has two symmetries, just like a face. An equilateral triangle, though, has six symmetries: three rotations (including the rotation by 0°), and three rotations when flipped over. You can keep track of the various symmetries by labelling the corners of the triangle, and seeing where they end up after applying one of the symmetries. (See the illustration.)

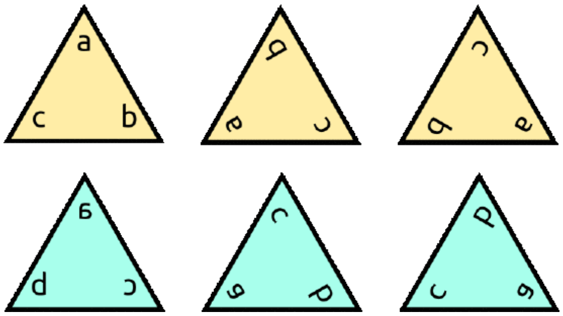

The six symmetries of an equilateral triangle. The top row contains the three rotational symmetries, while the second row has the 'flipped' and rotated symmetries.

? How many symmetries does a square have? How about an n-sided regular polygon?

We can also imagine an object which is symmetrical under some number of rotations, but which can't be flipped over. You can make such an object in many ways; one way is to take a square (or any other regular polygon) and then add an identical 'bump' just to one side of each corner. This object has rotational symmetry, but cannot be flipped.

In three dimensions, we have regular polyhedra. These are three dimensional objects with many symmetries! A tetrahedron has 24 symmetries, for example: twelve of these are rotations, and another 12 can be obtained by reflecting and then rotating.

? There are five regular polyhedra: the tetrahedron, the cube, the octahedron, the dodecahedron, and the icosahedron. How many symmetries does each one have? Try working it out directly for the smaller cases, then see if you can arrive at a formula for the polyhedra with more sides.

Of course, some objects have an infinite number of symmetries. A circle is a good example of this: every rotation is a symmetry, and there are infinitely many angles by which the circle may be rotated, all of which preserves its shape.

Thinking back to our regular tiling patterns, these also have infinitely many symmetries. All of them have *translational* symmetry, since you can translate the whole picture back onto itself. And you can translate in one direction as many times as you like, so there's at least one symmetry for every integer.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

1.3: Symmetric Polynomials

So far, we've considered geometric objects. Let's also have an example of something that isn't geometric. Let f be a polynomial in some number of variables. For now, we'll stick with 3 variables, x, y, and z. We say that f is a *symmetric polynomial* if every way of switching around (ie, *permuting*) the variables leaves f the same.

For example, the polynomial $f(x, y, z) = x + y + z$ is symmetric: switching the x and the z, for example, gives $z + y + x$, which is the same as f. As a more complicated example, you can check that $g(x, y, z) = x^2y + x^2z + y^2x + y^2z + z^2x + z^2y$ is also symmetric.

On the other hand, $h(x, y, z) = x^3 + y^3 + z$ is not symmetric, since switching x and z produces $z^3 + y^3 + x$, which is not equal to h. This polynomial does have *some* symmetry, since switching x and y leaves h the same, but we save the name 'symmetric polynomial' for the *fully* symmetric polynomials.

? Exercise 1.2.0:

Let f be a symmetric polynomial with n variables. how many symmetries does f have?

If you haven't tried a problem like this before - working in n variables - it is extremely important to get some practice. Try writing down some different symmetric polynomials with small numbers of variables. Is there a formula that describes the the number of symmetries in terms of the number of variables?

Symmetric polynomials are really interesting things, and we'll see them again when we talk about rings and vector spaces!

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

1.4: Abstraction. What is a group?

So we now have many examples of symmetry, but what, exactly is a group?

First, let's remember where numbers come from. You can imagine having a set of three pineapples, or a set of three people, or a set of three bottles. All of these share the property of 'three-ness.' So when we say the number three, it describes all of these different sets. The number three is an abstraction of the property of a set containing three things.

What's a group? A group is an abstraction of symmetry! Just like a number describes all sets with a certain number of things in it, a given group describes all objects with a certain kind of symmetry. Here's an example: The 'very symmetrical face' has two symmetries, related by a reflection. Likewise, the symmetric polynomial $f(x,y)=x+y$ has a two symmetries, one from leaving f alone, and one from exchanging the two variables x and y. If we think of switching x and y as a reflection, we see that the face and the polynomial somehow have the same kind of symmetry! The group is that measure of symmetry.

Now, in that last example, the two objects had the same number of symmetries. It turns out that just counting symmetries ins't enough to tell whether the two objects have the same *group of symmetries*. Think of the symmetries of our equilateral triangle: there were *rotational* symmetries, and there was a *reflection* symmetry. And there were six symmetries in all. Now, consider a regular hexagon with some regularly placed bumps on each side, so that there is no reflection symmetry available. Thus, this object has six rotational symmetries, but can't be flipped over like the triangle. Therefore it has the same *number* of symmetries as the equilateral triangle, but the set of symmetries is somehow different.

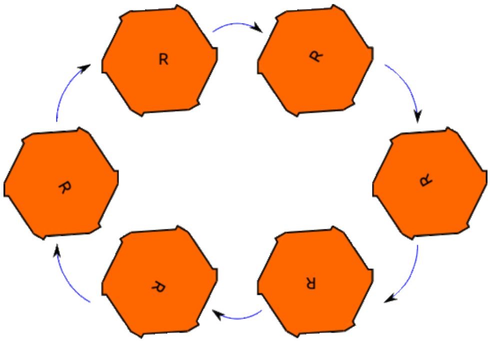

The six rotational symmetries of the bumpy hexagon. The blue arrow represents the rotation r; $r^6=id$.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

2: Groups I

We give a precise definition of a group and explore some different groups in the context of this definition.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

2.1: Symmetry and Functions

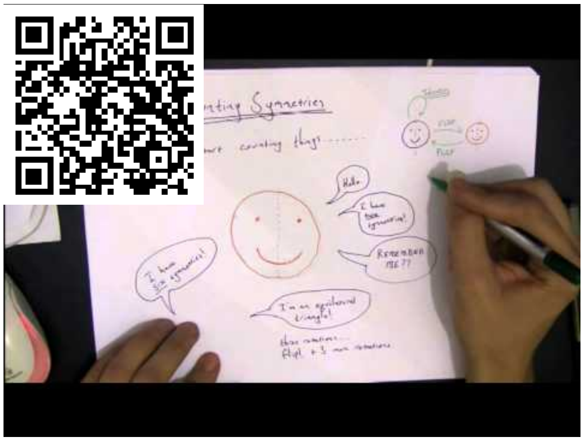

(For best results, view this video in full-screen.)

Now we will work towards a precise definition of a group. In mathematics, we often begin explorations with some general concept (in our case, symmetry) and then work on those concepts until we can arrive at a very precise definition. Having precise definitions allows us better abstractions and generalizations; the more precise our definitions, the better the mathematics we can expect to derive.

Recall our definition of the symmetries of a geometric object: all of the ways of moving that object back onto itself without changing it. Each of these different ways of moving the object back onto itself we can identify as a *function*. Let's call our object X. Then every symmetry of X we can identify as a certain special function $f: X \to X$.

Earlier we noted that leaving X alone should also be a symmetry of X. This is the identity function! $id: X \to X$, with $id(x) = x$ for all points $x \in X$.

Next, we notice that *composition of functions* is a helpful operation: Indeed, if we have two different symmetries f and g of X, then their composition $g \circ f$ will also be a symmetry. The first function applied to X 'moves X onto itself without changing it,' and then the second does as well. Thus, the composition is also a symmetry.

Finally, we note in passing that the composition of functions is associative: for three symmetries $f, g, h,$ we have $(f \circ g) \circ h = f \circ (g \circ h)$.

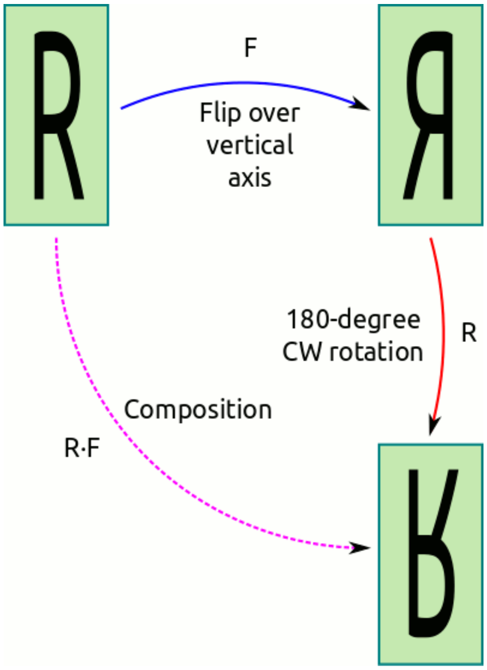

Composition of two symmetries of a rectangle.

We really need some examples here! So let's consider the perfectly symmetrical face. There are only two symmetries: the identity and the left-to-right flip (reflection over the vertical axis). Call the identity e and the flip f. Then we see that, considered as functions from the face to itself, $f \circ f = e$.

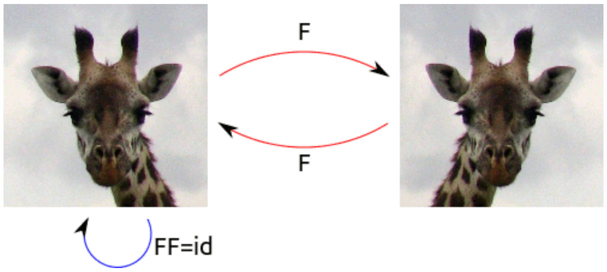

The red arrows describe the flip over the vertical axis. Do it twice, and you end up where you started: $F \circ F = id$.

Now remember our 'bumpy' hexagon, which only had rotational symmetries. Call the identity e and let r be the clockwise rotation by 60°. All of the other rotations we can think of as r composed with itself some number of times; we'll just write this as r^k. So all of the symmetries of the bumpy hexagon are $\{e, r, r^2, r^3, r^4, r^5\}$. We notice that $r^5 \circ r = r^6 = e$. This is quite interesting! In fact, we can make a 'composition table' to keep track of what happens when we compose any two of the symmetries.

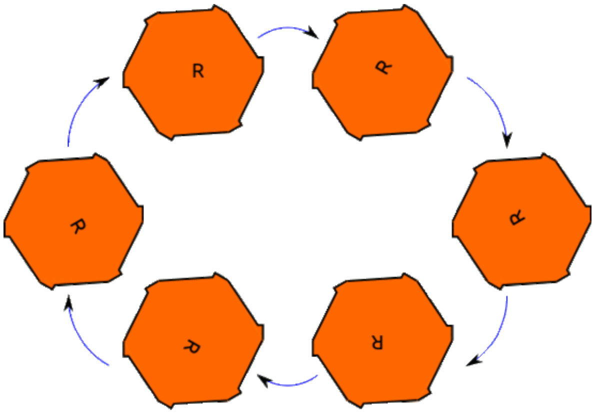

The six rotational symmetries of the bumpy hexagon. The blue arrow represents the rotation r; $r^6 = id$.

And now we can make a slightly less obvious observation: For any of these functions, we can find another symmetry taking the object 'back' to its original orientation. Thus, for any symmetry f, there exists a g such that $f \circ g = e$. The function g is then called the *inverse* of f. We've already seen two examples of this.

? Write down all of the symmetries of an equilateral triangle. Make a 'composition table' of the symmetries, showing what happens when any two of them are composed. Then make a list of each symmetry and its inverse. What do you observe?

Contributors and Attributions

2.2: Definition of a Group

Consider an object X with some symmetries S. We've seen that we can compose any of the symmetries in S and obtain another symmetry of X. We've also seen that these symmetries obey certain rules. We can now, at last, define a group.

Definition 2.1.0: Group

A *group* is a set S with an operation $\circ : S \times S \to S$ satisfying the following properties:

1. Identity: There exists an element $e \in S$ such that for any $f \in S$ we have $e \circ f = f \circ e = f$.
2. Inverses: For any element $f \in S$ there exists $g \in S$ such that $f \circ = e$.
3. Associativity: For any $f, g, h \in S$, we have $(f \circ g) \circ h = f \circ (g \circ h)$.

An essential notion in mathematics is *abstraction*. Note that our definition certainly applies to any collection S of symmetries of an object, but in fact there are other contexts where the definitions apply as well! The operation can be any way of combining two things in S and getting another back; S doesn't need to be a collection of functions, and the operation doesn't need to be composition. A group is defined purely by the rules that it follows! This is our first example of an algebraic structure; all the others that we meet will follow a similar template: A set with some operation(s) that follow some particular rules.

For example, consider the integers $\mathbb{Z}$ with the operation of addition. To check that the integers form a group, we need to check four things:

1. Addition takes two integers and gives another integer back. (Here we're checking the requirement that the operation is one from $S \times S \to S$. Notice that the the output of the operation is always in S! This is called *closure* of the operation.)
2. There's an identity element, 0, where for any integer n, we have $n + 0 = 0 + n = n$.
3. Every integer n has an inverse, $-n$, with $n + (-n) = (-n) + n = 0$.
4. Addition of integers is associative.

Thus, the integers - with the operation of addition - form a group.

On the other hand, the set of integers with the operation of multiplication do not form a group. Multiplication does indeed take two integers and return another integer, and there is an identity 1, and multiplication is associative. But not every element has an inverse that is also an integer. For example, the multiplicative inverse of 2 is $\frac{1}{2}$, but this isn't an integer! Thus, integers with multiplication do not form a group.

? An important note about inverses: An inverse means, roughly, that we can go back to where we started after applying an operation. Algebraically, this means we can cancel elements. When we have something like $gh = gk$, we can multiply both sides on the left by g^{-1} to get $h = k$. We have to be careful to multiply on the the same side on both sides, since groups aren't always commutative! If $gh = kg$, it doesn't necessarily tell us that $h = k$!

? of X.)

? Show that the symmetries of an equilateral triangle are not commutative. In other words, find two symmetries f, g of the equilateral triangle such that $fg \neq gf$.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

2.3: Integers Modulo n

Recall the 'bumpy' hexagon, which had rotational symmetry but no reflection symmetry. The group of symmetries of the bumpy hexagon is called $\mathbb{Z}_6$. In this section, we'll consider the general case, $\mathbb{Z}_n$, which we can initially think of as the group of symmetries of a 'bumpy' n-sided polygon.

There are many different ways in which $\mathbb{Z}_n$ appears in mathematics; it's a very important group! We now describe a number of different ways in which it arises.

1. Well, first we have the group of symmetries of the 'bumpy' n-sided polygon. By the exercise at the end of the last section, we know this is a group.
2. For our second definition, we'll define the 'remainder by n' operation: for any integer a, define $a\%n$ to be the remainder of a when divided by n. For example, $5\%3=2$, because the remainder of 5 when divided by 3 is 2. (You should check that for any integer k, $(kn)\%n=0$.) This operation is usually called 'modulus' or 'mod.' So $12\%5$ is read 'twelve modulo 5' or 'twelve mod 5.' (And is equal, of course, to two!)

   Usually, we don't write $+_n$ for the addition. From now on, whenever you see an expression like $4+3$, you will have to be mindful of the context! If we consider 4 and 3 as plain old integers, the answer is 7. If they are integers mod 5, then the answer is 2!
3. The next definition is really just an easy way to think of the second definition. Imagine a distant planet where the clock has n hours on it instead of 12 (or 24). Then, just as our hours 'wrap around' the circle beyond 12 o'clock, the hours wrap around at n. Now if we imagine the clock is numbered 0 through $n-1$ instead of 1 to n, we have exactly the situation of $\mathbb{Z}_n$.
4. Our last definition will identify $\mathbb{Z}_n$ with the n*-th roots of unity*, which are complex numbers. Recall that any complex number may be written as $re^{i\theta}$, where r is a positive real number and θ is any angle. Now let n and k be some positive integers, and consider the complex number $x_k = e^{\frac{k}{n}2i\pi}$. Then we can see that $x_k^n = (e^{\frac{k}{n}2i\pi})^n = e^{k2i\pi} = 1$. Then we call x_k an nth root of unity, because raising it to the nth power gives us 1 (aka, unity).

All of these are somehow the same; but there's a question of how to *formally* show that two groups are the same. What do we mean by the same? This is an important question to consider, which we will come back to later. For now, an exercise.

? Write out tables for $n=5$ and $n=6$ for:

1. composition of the rotations of the 'bumpy' n-gon,
2. addition in $\mathbb{Z}_n$,
3. addition of hours on an extraterrestrial clock with n hours,
4. and for multiplication of the n-th roots of unity.

In what ways are all of these groups the same? In what ways are they different?

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

2.4: Permutations

A *permutation* of n distinct objects is just a listing of the objects in some order. For example, $[c,b,a]$ is a permutation of the set $\{a,b,c\}$ of three objects. Likewise, [triangle, melon, airplane] is a permutation of three objects as well. From our mathematical point of view, the objects we use don't actually matter; all we care about is the order they are arranged in. So usually we'll just talk about permutations of the numbers 1 through n. You can think of each number as just counting the objects involved: first object, second object, nth object.

Permutations arise in the world in a many, many ways. For example, suppose you are asked to list your preferences amongst a bunch of presidential candidates. The list you make up, from favorite to least favorite, is a permutation of the candidates. In fact, you can use the mathematics of permutations to learn interesting things about different kinds of voting systems.

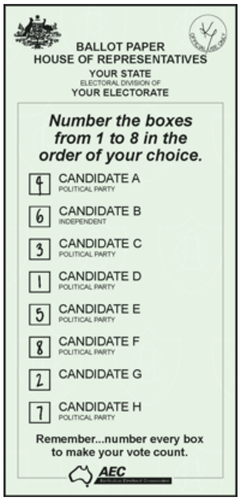

BALLOT PAPER
HOUSE OF REPRESENTATIVES
YOUR STATE
ELECTORAL DIVISION OF
YOUR ELECTORATE

Number the boxes from 1 to 8 in the order of your choice.

4	CANDIDATE A POLITICAL PARTY
6	CANDIDATE B INDEPENDENT
3	CANDIDATE C POLITICAL PARTY
1	CANDIDATE D POLITICAL PARTY
5	CANDIDATE E POLITICAL PARTY
8	CANDIDATE F POLITICAL PARTY
2	CANDIDATE G
7	CANDIDATE H POLITICAL PARTY

Remember...number every box to make your vote count.

AEC

Figure 2.3: Instant run-off voting uses a full permutation of the candidates to find a winner. (Source)

Another example is a deck of playing cards. In a standard deck, each card appears exactly once. When you shuffle the deck, you are just creating a random permutation of the cards. One can use mathematics related to permutations to answer interesting questions about cards. Like: 'How many times do I need to shuffle the deck before it is truly randomized?' The answer, by the way, seems to be 7 for a standard riffle shuffle. But proving that is well beyond the scope of these notes!

Because permutations are so common, problems involving permutations tend to be very applicable! For example, suppose you have two hundred students in a class and they all hand in an exam. The stack of exams they give you is a permutation of the students; most likely, the list of student scores you keep is alphabetical. This suggests a problem: What is the fastest way to sort the exams? (In fact, sorting is a fundamental problem in computer science.)

How many permutations are there of a set of n objects? Suppose we try to build a permutation by successively choosing objects. Then there are n choices for the first object, $n-1$ choices for the second, and so on, until there is only one choice for the last object. Then to get the total number of possible permutations, we multiply these numbers together, and get $n(n-1)(n-2)\cdots 1$. This number, if you haven't seen it before, is called n-factorial, written $n!$.

? Write out all of the permutations of the set $\{1,2,3,4\}$. How many are there in all? Find a sensible way to organize your list!

Suppose we have some initial ordering of our objects. The letters $\{a,b,c\}$, for example, can be organized alphabetically. Then every permutation we can think of as a mixing-up of this initial order. In this sense, the permutation is a special kind of function from the set of objects back to itself. (By special, I mean it's a *bijection*, which is to say a one-to-one and onto function.) (TODO: Wikipedia link) A permutation of these objects is then the list $[\sigma(a),\sigma(b),\sigma(c)]$; this list is called the *one-line notation* for σ.

These permutations-as-functions can be composed: if you think of two permutations σ and τ as different ways to mix up the set, you can mix them up according to σ and then according to τ. Then the composition is specified by the list $[\tau(\sigma(a)),\tau(\sigma(b)),\tau(\sigma(c))]$.

For example, if $\sigma=[2,3,1,4]$ and $\tau=[3,4,1,2]$, then $\tau\circ\sigma=[4,3,1,2]$. (In particular, $\sigma(1)=2$, and $\tau(2)=4$, so the first entry of $\tau\circ\sigma$ is 4. The other three entries are computed similarly.) On the other hand, $\sigma\circ\tau=[3,1,4,2]$. This is different from $\tau\circ\sigma$! So we see that the group of permutations has elements where $f\circ g\neq g\circ f$; we say that S_n is *non-commutative*. (But remember that nothing in our group definition says that a group needs to be commutative, so this is ok.)

A very nice way to keep track of this mixing-up is the braid notation for a permutation. This simply writes the list of objects in two lines, and draws a line connecting an object on the top to the object it is sent to under the permutation.

Braid diagrams for some permutations. At this point, we can ask whether the permutations with the composition operation are in fact a group. In fact, they are! Let's check. Let σ,τ be permutations of the set $X=\{1,2,3,\ldots,n\}$ Then we can specify σ by the list $[\sigma(1),\sigma(2),\ldots,\sigma(n)]$

Composition of two permutations is again a permutation. Since each permutation contains every element of X exactly once, the composition $\tau\circ\sigma$ must also contain each element of X exactly once. Identity: The permutation $[1,2,\ldots,n]$ acts as the identity. Inverses: Roughly speaking, if you can mix things up, you can just as easily sort them back out. The 'sorting permutation' of σ is exactly σ^{-1}. Associativity: Suppose we compose three permutations, σ, τ, and ρ. Int he braid notation, this just means placing the three braids on top of each other top-to-bottom, and then 'forgetting' the two sets of intermediate dots. (TODO: a picture!) Associativity is tantamount to forgetting the two sets of dots in two different orders; the resulting picture is the same either way, so composition of permutations is associative!

? Carefully work through the above and check for yourself that permutations satisfy the definition of a group. For example, where it is stated that the identity permutation has one-line notation $[1,2,\ldots,n]$, you should check that this is actually the identity. Likewise, how can you explicitly compute the inverse of a permutation explicitly?

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

3: Groups II

In this chapter we explore the structure of groups using Cayley graphs and generating sets. We also learn about Lagrange's theorem, which gives an interesting numerical relationship between the size of a group and the size of a subgroup.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

3.1: Generating Sets

We have now seen a few different kinds of groups: groups of symmetries of a geometric object, integers under addition, integers modulo n, and permutations. We can easily visualize the objects related to the group - like the geometric object, numbers, or the braid notation for permutations - but how can we visualize the group itself?

An excellent way to go about this is to identify a set of generators for the group. In a group we can always combine some elements using the group operation to get another group element. Generators are some special elements that we pick out which can be used to get to any other element in the group.

As an example, remember the dihedral group, the symmetries of an n-sided polygon. There are $2n$ symmetries in all, but we can build up any of the symmetries using just a small rotation and a flip. For the symmetries of the equilateral triangle, we let ρ denote the rotation by 120 degrees, and let f be the flip over one of the axes of the triangle. Then the six elements of the dihedral group are given by: $id, \rho, \rho^2, f, f\rho, f\rho^2$. Thus, $\{f, \rho\}$ is a set of generators for the dihedral group.

Here's the formal definition:

Definition 3.0.0:

Let G be a group, and S a subset of G. We say that S *generates* G (and that S is a set of *generators* for G) if every element of G can be expressed as a product of elements of S and their inverses.

We include the inverses of the generators in the definition because we know that every element has an inverse. If we think of the integers under addition, we can write every positive number as a many-times sum of the number 1: for example, 5 is just $1+1+1+1+1$. If we allow inverses as well, we can then get every element of the group from a single generator: the inverse of 1 is -1, so we can write (for example) $-4=(-1)+(-1)+(-1)+(-1)$. (Including the inverses also means we don't need to include the identity, since for any g, $gg^{-1}=e$.)

On the other hand, for any group G, we can certainly take G itself as a generating set! Then every element is considered a 'generator,' so every element can be written as a (trivial) product of generators. This tells us that for any group we can find a generating set. Usually, we try to find a generating set as small as possible. Sometimes, though, a larger generating set might be interesting if it helps us to better understand the group in question.

Once we have a generating set for a graph G, we can produce a very nice visualization of the group called the Cayley graph. By graph, we mean a number of points (called vertices) connected by some arrows (called edges). Graphs are good for keeping track of relationships between things, and appear in many, many places in mathematics and in applications.

The Cayley graph of a group has one vertex for each element x in the group. Each vertex has one arrow coming out of it for each generator g, pointing to the element gx. (This creates the left Cayley graph. The right Cayley graph has arrows pointing from x to xg.) Usually we make the arrows different colors to correspond to the different generators; this is very useful for being able to visualize the structure of the group!

For the dihedral group, we found a set of generators with two elements: the rotation and the flip over one of the axes. In fact, the dihedral group has many different sets of generators of size two! We could have chosen the clockwise rotation instead of the counter-clockwise rotation, for example. Or we could have chosen any of the other flips. But the resulting Cayley graph would have been more-or-less the same.

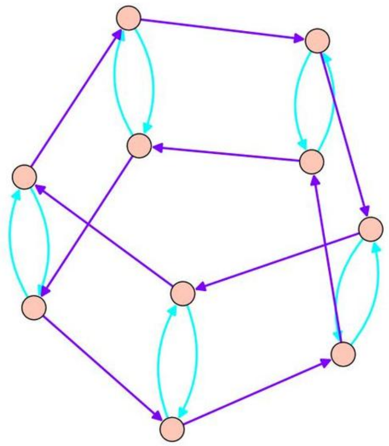

Figure 3.1: Dihedral group Cayley graph, generated by a flip and rotation.

A quite different set of generators for the dihedral group is to take two different flips, across axes that are adjacent to one another. Let's call them f_1 and f_2. You can actually still write any element of the dihedral group as a product of these two flips. And the resulting Cayley graph looks quite different.

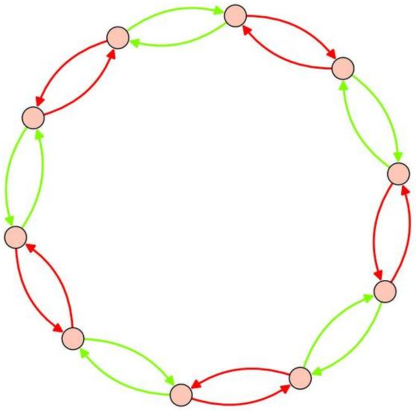

Figure. 3.2: The Cayley graph for the dihedral group with generators given by two different flips.

Suppose we have a generator where $g^2 = 1$. It's tedious to draw arrows in both directions from every element, so we sometimes omit the arrow heads in this case.

? Exercise 3.0.1

Identify generators for the permutation group S_3. Make a Cayley graph.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

3.2: Visualizing Groups- Cayley Graphs

So far, we've seen three different kinds of groups: Groups of symmetries (including the dihedral group of symmetries of a polygon), the integers modulo n, and the permutation group, S_n. We've seen Cayley graphs for the dihedral group; let's see some Cayley graphs for some others.

Integers Modulo n

The integers modulo n only require a single generator to obtain the entire group. 1 is a nice choice of generator, as we know that every number k in $\{0,1,2,...n\}$can be written as $k\cdot 1$, which is to say, the sum of 1 with itself k times. The Cayley graph for this situation is simple: it's just n vertices, arranged in a loop with an arrow pointing from each number to the next. This creates a cycle! When $n=8$, the cycle is this:

$$0\rightarrow 1\rightarrow 2\rightarrow 3\rightarrow 4\rightarrow 5\rightarrow 6\rightarrow 7\rightarrow 0 \tag{3.2.1}$$

. Any group which is generated by a single element (including the usual integers!) is called a cyclic group. (This is yet another interpretation of $\mathbb{Z}_n$!)

We can choose other numbers than 1 as the generator, though! Take $n=8$, and consider the number 3. We can make our Cayley graph by drawing a vertex for each number in $\{0,1,2,\ldots,8\}$and an arrow from each x to $x+3$. Then the cycle draws out as: [0\rightarrow 3\rightarrow 6\rightarrow 1\rightarrow 4\rightarrow 7\rightarrow 2\rightarrow 5\rightarrow 0].

Here's a Cayley graph for $\mathbb{Z}_7$ shown with three generators. Any one of the three generators would work just fine. The red vertex is the identity, 0. The green arrows are for the generator 1, blue for the generator 2, and green for the generator 3. What would happen if we included the generators $4, 5$, or 6?

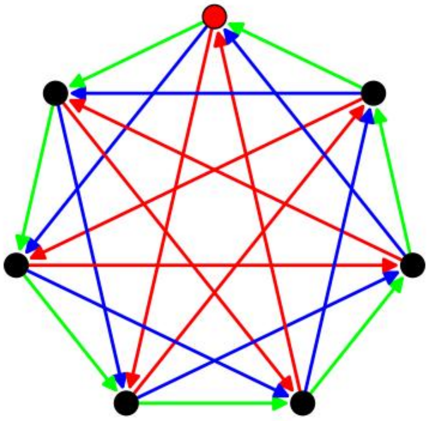

Figure 3.1: Cayley graph for $\mathbb{Z}_7$ with three different generators. The identity is marked as the red dot.

Not every number is a generator of $\mathbb{Z}_n$. For example, in $\mathbb{Z}_8$, if we choose 4, the cycle is just: $0\rightarrow 4\rightarrow 0$. Since the cycle doesn't contain every element of the group, we see that 4 doesn't generate the group on its own.

? Exercise 3.1.0

Suppose $k\in\mathbb{Z}_n$. Show that k generates $\mathbb{Z}_n$ if and only if k is relatively prime to n. (ie, the only common divisor of k and n is 1.)

? Exercise 3.1.1

Suppose $k\in\mathbb{Z}_n$ and k is not relatively prime to n. Is it possible to find another umber m not relatively prime to n such that k and m together generate $\mathbb{Z}_n$? Try some examples! Explain why or why not

Permutation Groups

The permutation group S_n has a number of interesting generating sets. We'll show a few of these generating sets for $n=3$ and $n=4$ for easy comparison.

The first generating set is a minimal set, using just two generators. The first generator is the 'rotation' with list notation $r=[n,1,2,3,4,\ldots,n-1]$. The second is a flip, exchanging only the first two things, $f=[2,1,3,4,\ldots,n]$.

To check that these actually generate S_n, we need to see that we can construct an arbitrary permutation using just these generators. So consider an arbitrary permutation σ, written in list notation. If there are two adjacent entries that are out of order (big to the left of the small), we can apply rotations until the two things sit in the first two entries (suppose we use k rotations to do this). Then we apply the flip. And then we 'unrotate' k times to put the now-sorted numbers back. Then we find two more adjacent numbers and repeat. Once there are no adjacent numbers out of order, then we must be at the identity! Then the reverse of the sequence of moves we just made builds the permutation we wanted. Since the permutation was arbitrary, our two moves must generate the group.

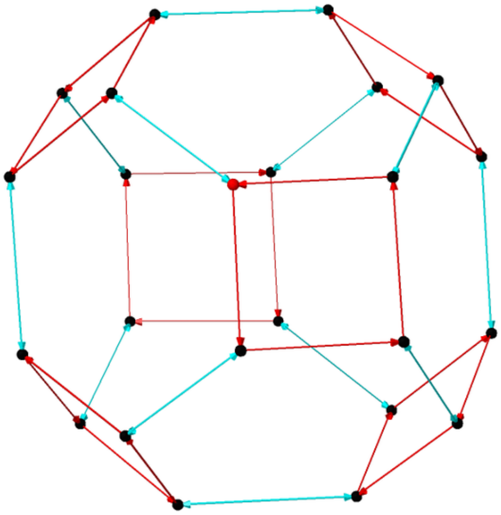

Figure 3.2: A Cayley graph for S_4, generated by the rotation $[4,1,2,3]$(in red) and reflection \([2,1,3,4]) (cyan).

A second set of generators is given by the set of all transpositions. These are all of the permutations that have two things switched and everything else in order. For example, $[4,2,3,1]$ and $[1,2,6,4,5,3]$are transpositions. Modifying the above argument, you can see that the set of all transpositions are a generating set. There are more than 2 transpositions, so this isn't a minimal generating set. But it is an interesting set of generators when studying the permutation group more closely.

Exercise 3.1.2

How many transpositions are there in S_n?

Yet a third set of generators is given by the simple transpositions. This is the set of transpositions $\{s_i\}$ that just exchange i and $i+1$ while leaving everything else alone. There are $n-1$ simple transpositions. This is a very important set of generators in the further study of permutations! But it shows up in one simple context, as well.

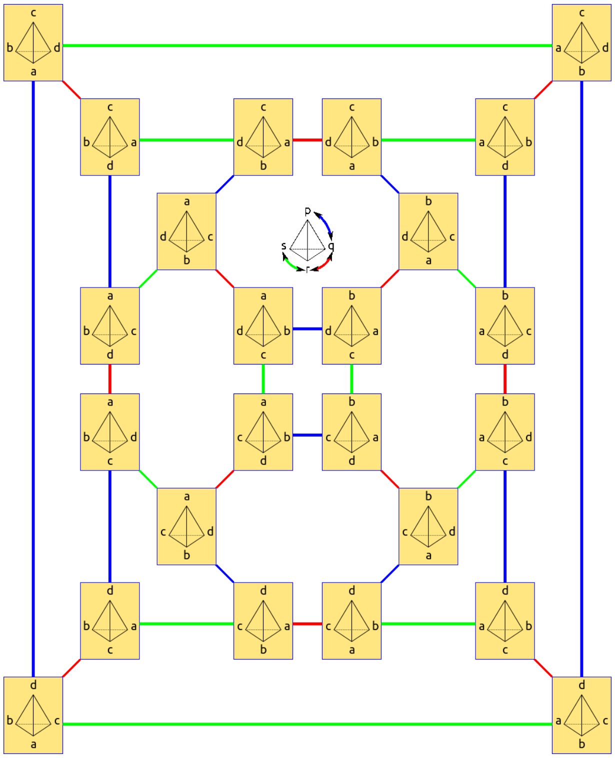

Figure 3.1.3: Cayley graph for S_4; permutations are marked by different symmetries of the tetrahedron. Generators are three 'flips' exchanging two vertices.

A basic problem in computer science is sorting. Given a list of n things, how quickly can they be sorted? What is a good algorithm for sorting an arbitrary list? There are many different sorting algorithms. One of the easiest is called Bubble Sort. For bubble sort, you read through the list, beginning to end, and whenever you see two adjacent entries that are out of order, you switch them. You may have to read through the list performing switches many times, but eventually the list will be sorted. Bubble Sort uses the fact that the simple transpositions are a generating set for the permutations in order to sort an arbitrary list. (This is the first step into the study of complexity theory.)

Exercise 3.1.3

What permutation of n things takes the longest to be sorted by Bubble Sort? How many simple transpositions are necessary to sort that permutation?

Exercise 3.1.4

For the same 'long' permutation from the last exercise, sort the permutation using the first set of generators for S_n, the rotation and the flip. How many steps are needed to sort the permutation this way?

We see that there's a trade-off between having a smaller set of generators and being able to write different group elements as products of fewer generators. (Indeed, if we took the whole group as the generating set, every element could be written as a product of just one generator! But this usually isn't so helpful for understanding the group...)

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

3.3: Subgroups

When we consider the symmetries of, say, a pentagon, we notice that it has rotational symmetries like the 'bumpy' pentagon. From the bumpy pentagon, we see that the rotations themselves form a group; there's a group of rotations inside the group of symmetries of the pentagon! Likewise, if we consider just the flip, we see a group similar to the symmetries of the perfectly symmetrical face. This yields another group inside the symmetries of the pentagon. We can make this precise:

Definition 3.2.0: Subgroup

Let G be a group, and H a subset of G. Then H is a *subgroup* of G if H is itself a group using the same operation as G.

Ostensibly, to check that a subset H is a subgroup, we would need to check all four properties of the group. That is, closure (ie, the operation gives a map $H \times H \to H$; products of things in H are always in H), identity, the existence of inverses, and associativity.

In fact, since H has the same operation as G, we know that the operation in H is associative (since G is a group). Furthermore, if the operation is closed and inverses exist, then we know that for any $h \in H$, $hh^{-1} = e$ must be in H. So really we only need to check two things:

1. Closure: $gh \in H$ for all $g, h \in H$, and
2. Inverses: $h^{-1} \in H$ for all $h \in H$.

Some important things to notice:

1. The group G is always a subgroup of itself! (G is a subset of itself, which is a group with the same operation as G.)
2. The subset containing just the identity element is also a subgroup! This is called the trivial subgroup.
3. The set of all powers of an element h ($\{\ldots, h^{-1}, h^{-2}, e, h, h^2, \ldots\}$) is a subgroup of G. This is called the cyclic subgroup generated by h.

Exercise 3.2.1

Let X be a geometric object. Show that the rotations of X back onto itself forms a subgroup of the group of symmetries of X. (Try this in particular on a regular polygon and a regular polyhedron. What happens with a 'bumpy' polygon?)

Let G be a group, and $g \in G$. Consider a function $f_g : G \to G$ given by $f_g(h) = g \cdot h$. (This is the 'left multiplication by g' function.) What happens if, for some $h, k \in G$, $f_g(h) = f_g(k)$? Then $gh = gk$, so $g^{-1}gh = g^{-1}gk$, and $h = k$. This tells us that f_g is a one-to-one, or injective, function. If G has a finite number of elements, then f_g is also an onto function, and is thus a bijection from G back to itself. Then we can consider f_g as a permutation of G!

If we consider G as a set, we can think of any left multiplication as a permutation of G. But the set of all left multiplications is itself a group. This gives us what is known as Cayley's Theorem!

Theorem 3.2.2: Cayley's Theorem

The ideal gas law is easy to remember and apply in solving problems, as long as you get the **proper values a**

Exercise 3.2.3

Label the six symmetries of the equilateral triangle. Demonstrate that the symmetries of the triangle are a subgroup of S_6, the permutations of 6 objects.

It is worth noticing that for any g in a group G, the powers of g generate a subgroup of G. The set $\{g^i \mid i \in \mathbb{Z}\}$ is closed under the group operation, and includes the identity and inverses. This is called the cyclic subgroup generated by g.

? Exercise 3,2,4

Find all of the subgroups of the permutation group S_3 for three objects. Which subgroups are subgroups of other subgroups? Name each subgroup, and arrange them according to which is contained in which.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

3.4: Cosets and Lagrage's Theorem

In this section, we'll prove Lagrange's Theorem, a very beautiful statement about the size of the subgroups of a finite group. But to do so,we'll need to learn about cosets.

Recall the Cayley graph for the dihedral group D_5 as generated by a flip and a rotation. Notice that the darker blue arrows look like two different 'copies' of $\mathbb{Z}_5$ sitting inside of the dihedral group. Likewise, the light arrow loops look like five copies of $\mathbb{Z}_2$. Both $\mathbb{Z}_5$ and $\mathbb{Z}_2$ are subgroups of D_5, generated by the rotation and flip respectively. These 'copies' of the subgroups that we see in the Cayley graph are examples of cosets.

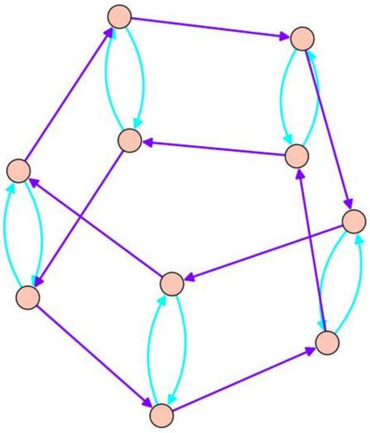

Dihedral group Cayley graph, generated by a flip and rotation. The Cayley graph for the dihedral group with generators given by a flip and a rotation.

Definition 3.3.0: Coset

Let H be a subgroup of G, and $H=\{h_1,h_2,h_3,\dots\}$. Then for any choice of $g\in G$, the *coset* gH is the set $\{gh_1,gh_2,\dots\}$.

These are precisely the 'copies' of the subgroups that we saw in D_5. The elements of D_5 can all be written as f^ir^j with $i\in\{0,1\}$ and $j\in\{0,1,\dots,4\}$. The rotation subgroup consists of the element $R=\{e,r,r^2,r^3,r^4\}$. Then R has two distinct cosets, R and $fR=\{f,fr,fr^2,fr^3,fr^4\}$. For any g we choose, gR is equal to one of these two cosets! Likewise, if we consider the subgroup $F=\{e,f\}$, there are five distinct cosets given by r^iF, where $i\in\{0,1,2,3,4\}$. Notice that the cosets evenly divide up the group; this isn't an accident!

Proposition 3.3.1

Suppose that H is a subgroup of G. Let $x,y\in G$. Then either $xH=yH$ or xH and yH have no elements in common.

Proof 3.3.2

Suppose $z\in xH$ and $z\in yH$. In particular, there exist $h_1,h_2\in H$ such that $z=xh_1=yh_2$, and $xh_1h_2^{-1}=y$. We need to show that $xH=yH$, so take any $h_3\in H$ and consider $yh_3\in yH$. Then $yh_3=xh_1h_2^{-1}h_3\in xH$, since $h_1h_2^{-1}h_3\in H$. Thus, if xH and yH share any elements, then they are equal, and if they share no elements, they are tautologically disjoint!

We'll need one more piece of notation.

Definition 3.3.3: Order of a group

The *order* of a group G, written $|G|$, is the number of elements in G.

So the order of $R\subset D_5$ is $|R|=5$, and the order of the flip subgroup $|F|=2$. We can now prove Lagrange's Theorem!

Theorem 3.3.4: Lagrange's Theorem

Let H be a subgroup of a finite group G. Then $|H|$ divides $|G|$.

Proof 3.3.5: Lagrange's Theorem

Notice that every element of the group G shows up in some coset of H: since $e \in H$, we have $g \in gH$ for every g. Therefore, every element of the group shows up in exactly one coset of H. Also notice that every coset of H has the same number of elements as H. (If the size of gH were less than $|H|$, there would be have to be two different elements $h_1, h_2 \in H$ with $gh_1 = gh_2$. But cancelling the g's gives $h_1 = h_2$, a contradiction.) Then the cosets of H break up G evenly into subsets of size $|H|$. Thus, $|H|$ divides $|G|$, as desired.

? Exercise 3.3.6

Find all of the subgroups of the permutation group S_3 and the dihedral group D_5.

? Exercise 3.3.7

Find a subgroup of the permutation group S_4 with twelve elements. What are it's cosets?

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

4: Groups III

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

4.1: Homomorphisms

When we think of the integers, it is useful not just to think of individual numbers, but relations between numbers; people do this automatically, comparing two numbers to see which is bigger. And if neither number is bigger, we see that the two numbers are equal. In fact, this is a very important way to approach mathematics: We consider not just *objects* but also *how objects are related to one another*. So far, we have thought extensively about group as objects. But are there interesting relationships between groups? If so, can we come up with a notion of when two groups are the same?

We can relate groups using special functions between different groups called homomorphisms. There's a way to think of relationships between numbers as functions, too: if you have k cows and p chickens, we can try to make a one-to-one function from cows to chickens. If it's impossible, we know the number of cows is larger. If there are chickens left over (ie, the function is not onto) then the number of chickens is larger. But if the function is one-to-one and onto, then the two numbers p and k are equal.

Symmetry and Length-Preserving Functions

Think for a moment of the symmetries of the equilateral triangle. These were given by functions from the triangle back to itself. But not just any functions: the symmetries don't distort the triangle in any way: For example, The center of the triangle never gets moved closer to one of the vertices. In particular, the symmetries are *length-preserving* functions, are known as *isometries*. We can be quite precise about what a length-preserving function is.

Definition 4.0.0: Length Preserving Functions

Let $d(x,y)$ denote the distance between two points x and y. Then a function f is *length-preserving* if $d(x,y)=d(f(x),f(y))$ for every pair of points x,y. In other words, distances before we apply the function are the same as distances after we apply the operation.

We should look at lots of examples to build up some intuition! Take the integers with the operation of addition. Define phi:

? Find a function on the interval $[0,1]$ that changes the endpoints but is not a symmetry. In other words, find some $f:[0,1]\to[0,1]$ such that $f(0)=1$, $f(1)=0$, but f is not length-preserving.

So symmetries are a special kind of function which preserve distances. When we try to relate groups to one another, we use special kinds of functions between the groups.

Homomorphisms

A group is a set with an operation which obeys certain rules. So we'll consider functions that *preserve the operation*. That is, functions for which it doesn't matter whether we perform our group operation before or after applying the function. More precisely:

Definition 4.0.2: Homomorphism

Let G and H be groups, and $\phi:G\to H$. Then ϕ is a homomorphism if $\phi(gh)=\phi(g)\phi(h)$. If a homomorphism is also a bijection, then it is called an *isomorphism*.

We should look at lots of examples to build up some intuition! Take the integers with the operation of addition. Define $\phi:\mathbb{Z}\to\mathbb{Z}$ by $\phi(n)=2n$. Note that the definition of homomorphism works regardless of the symbol we're using for the group operation, and for $\mathbb{Z}$ we use addition. Then to show that ϕ is a homomorphism, we need to check that $\phi(n+m)=\phi(n)+\phi(m)$; the operation before applying ϕ is the same as the operation after applying ϕ. So we check! $\phi(n+m)=2(n+m)$, while $\phi(n)+\phi(m)=2m+2n$. Since $2(n+m)=2n+2m$, we see that ϕ is a homomorphism.

We can also have homomorphisms between groups where the operations are written differently! For example, there is a homomorphism between the integers modulo n ($mathbbZ_n$) and the nth roots of unity. Remember that $mathbbZ_n$ is written with an addition operation, while the nth roots of unity are written with multiplication. We define ρ by $\rho(k)=e^{\frac{ik}{2\pi}}$. Then we check that ρ is a homomorphism! Since the operations are written differently (addition and multiplication), we need to check whether $\rho(k+j)=\rho(k)\rho(j)$. This isn't so bad: $\rho(k+j)=e^{\frac{i(k+j)}{2\pi}}$. On the other hand, $\rho(k)\rho(j)=e^{\frac{i(k)}{2\pi}}e^{\frac{i(j)}{2\pi}}=e^{\frac{i(k+j)}{2\pi}}$. So this is a homomorphism; in fact, it is an isomorphism, since the n-th roots of unity and $\mathbb{Z}_n$ have the same number of elements.

Isomorphisms are very special homomorphisms. If two groups are isomorphic, it is impossible to tell them apart using just the tools of group theory. True, the two groups may look very different, but they are *structurally identical*. When we saw the integers modulo n, we saw four different realizations of 'the same' group; they were all isomorphic.

? Define $\rho : \mathbb{Z}_n \to \mathbb{Z}_{2n}$ by $\rho(x) = 2x$ for each x in $\mathbb{Z}_n$. Show that rho is a isomorphism. (Hint: Show that the map is a homomorphism, and argue that the two sets have the same cardinality.)

? Let H be a subgroup of G. Define the inclusion $\iota : H \to G$ by $\iota(x) = x$ for each $x \in H$. Show that ι is a homomorphism.

There are many things we can say about homomorphisms with just a little work. We'll prove two basic statements right away.

Proposition 4.0.5

Let $\phi : G \to H$ be a homomorphism. Then:

1. $\phi(1) = 1$.
2. For any $x \in G$, $\phi(x^{-1}) = \phi(x)^{-1}$.

Proof 4.0.6

1. Choose any element $x \in G$. Then $rho(x) = \rho(1x) = \rho(1)\rho(x)$. So $rho(x) = \rho(1)\rho(x)$. Cancelling the $rho(x)$ on both side leaves us with $1 = \rho(1)$.
2. We have $\phi(1) = 1$, so $1 = \phi(xx^{-1}) = \phi(x)\phi(x^{-1})$, giving us $1 = \phi(x)\phi(x^{-1})$. Then we can multiply both sides on the left by $\phi(x)^{-1}$ to get the result.

? Exercise 4.0.7

1. Show that for any homomorphism ϕ, we have $\phi(x^n) = \phi(x)^n$.
2. Show that if two finite cyclic groups have the same order, then they are isomorphic.

This tells us that group homomorphisms, in addition to preserving the group operation, also preserve inverses and exponents. Thus, group homorphisms also preserve inverses and exponents!

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

4.2: Product Groups

So far, we have a fairly small collection of examples of groups: the dihedral groups, the symmetric group, and $\mathbb{Z}_n$. In this section, we'll look at products of groups and find a way to make new groups from the groups we already know.

A very famous group - though not a very complicated one - is the *Klein Four-Group*. This is the symmetry group of a rectangle. It has a pair of generators, given by the flips over the horizontal and vertical axes.

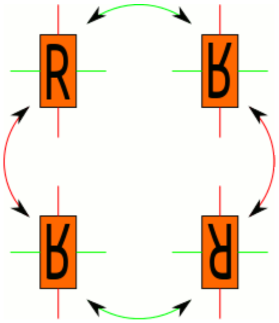

Figure 4.2.1. The symmetries of a rectangle, given by the Klien 4-group.

But the Klein Four-Group can also be thought of as a kind of mash-up of two copies of $\mathbb{Z}_2$. Let H be the additive group with elements $\{(0,0),(1,0),(0,1),(1,1)\}$ and operation given by just adding the elements coordinate-wise as elements of $\mathbb{Z}_2$. (So that $(1,0)+(1,1)=(0,1)$. Then H is a group (check!), and is in fact isomorphic to the Klein Four-Group. It's an example of a product group!

Let's be more precise and set a definition of a product group.

Definition 4.1.0: Direct Product

The *direct product* (or just *product*) of two groups G and H is the group $G\times H$ with elements (g,h) where $g\in G$ and $h\in H$. The group operation is given by $(g_1,h_1)\cdot(g_2,h_2)=(g_1g_2,h_1h_2)$, where the coordinate-wise operations are the operations in G and H.

Here's an example. Take $G=\mathbb{Z}_3$ and $H=\mathbb{Z}_6$, and consider the product $G\times H$. The product group has 18 elements: there are three choices for the first coordinate and 6 choices for the second coordinate. Since we use addition as the operation in both of the coordinate groups, we'll use addition as the operation in the product. So consider elements $(2,4)$ and $(1,3)$. Then $(2,4)+(1,3)=(0,1)$; addition in the first coordinate is according to $\mathbb{Z}_3$, and addition in the second coordinate is according to $\mathbb{Z}_6$.

We should check that the product of any pair of groups G and H is actually a group.

1. The product group has an identity $(1,1)$: $(1,1)\cdot(g,h)=(1g,1h)=(g,h)$.
2. Associativity follows from associativity of G and H.
3. Closure also follows from closure in G and H.
4. The inverse of (g,h) is (g^{-1},h^{-1}).

So $G\times H$ really is a group.

We saw in the example that $\mathbb{Z}_3\times\mathbb{Z}_6$ has 18 elements. This isn't a coincidence! For any finite groups G and H, the product group has $|G||H|$ elements.

An interesting question at this point is suggested by Lagrange's Theorem, which told us that the cardinality of any subgroup divides the cardinality of the original group. We've seen that we can form product group sto 'multiply' groups: Is it also possible to 'divide' groups? Over the next few sections, we'll develop ideas that will let us build *quotient groups*.

? Not all product groups are commutative. How many elements are in $G = S_4 \times \mathbb{Z}_3$? Identify the identity. Write down a few non-identity elements and compute their respective products.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

4.3: Image and Kernel

Definition 4.2.0

The *image* of a homomorphism $\rho : G \to H$ is the set $\{\rho(g) \mid g \in G\} \subset H$, written $\rho(G)$. The *kernel* of ρ is the set $\{g \mid g \in G, \rho(g) = 1\}$, written $\rho^{-1}(1)$, where 1 is the identity of H.

Let's try an example. Recall the homomorphism $\phi : \mathbb{Z} \to \mathbb{Z}$, defined by $\phi(n) = 2n$ for any $n \in \mathbb{Z}$. The image of ϕ is the set of all even integers. Notice that the set of all even integers is a subgroup of $\mathbb{Z}$. The kernel of ϕ is just 0.

Here's another example. Consider the map $\phi : \mathbb{Z}_3 \to \mathbb{Z}_6$ given by $\phi(n) = 2n$. So $\phi(0) = 0$, $\phi(1) = 2$, and $\phi(2) = 4$. This is actually a homomorphism (of additive groups): $\phi(a+b) = 2(a+b) = 2a + 2b = \phi(a) + \phi(b)$. The image is the set $\{0, 2, 4\}$, and, again, the kernel is just 0.

And another example. There's a homomorphism $\rho : \mathbb{Z}_6 \to \mathbb{Z}_3$ given by $\rho(a) = a$ (divide by 3 and keep the remainder). Then $\rho(0) = 0$, $\rho(1) = 1$, $\rho(2) = 2$, $\rho(3) = 0$, $\rho(4) = 1$ and finally $\rho(5) = 2$. You can check that this is actually a homomorphism, whose image is all of $\mathbb{Z}_3$ and whose kernel is $\{0, 3\}$.

So the image is the set of everything in H which has something in G which maps to it. The kernel is the set of elements of G which map to the identity of H. The kernel is a subset of G, while the kernel is a subset of H. In fact, both are subgroups!

Proposition 4.2.1

The image $\rho(G)$ is a subgroup of H. The kernel $\rho^{-1}(1)$ is a subgroup of G.

To see that the kernel is a subgroup, we need to show that for any g and h in the kernel, gh is also in the kernel; in other words, we need to show that $\rho(gh) = 1$. But that follows from the definition of a homomorphism: $\rho(gh) = \rho(g)\rho(h) = 1 \cdot 1 = 1$. We leave it to the reader to find the proof that the image is a subgroup of H.

? Show that for any homomorphism $\rho : G \to H$, $\rho(G)$ is a subgroup of H.

We can use the kernel and image to discern important properties of ρ as a function.

Proposition 4.2.3

Let $\rho : G \to H$ be a homomorphism. Then ρ is injective (one-to-one) if and only if the kernel $\rho^{-1}(1) = \{1\}$.

Proof 4.2.4

If we assume ρ is injective, then we know (from the exercise in the last section) that $\rho^{-1}(1) = \{1\}$. For the reverse direction, suppose $\rho^{-1}(1) = \{1\}$, and assume (for contradiction) that ρ is not injective. Then there exist $x \neq y$ with $\rho(x) = \rho(y)$. But then $\rho(x)\rho(y)^{-1} = \rho(xy^{-1}) = 1$. Since $x \neq y$, $xy^{-1} \neq 1$, giving a contradiction.

The kernel is actually a very special kind of subgroup.

Proposition 4.2.5

Let $\rho : G \to H$ be a homomorphism, and let K be the kernel of ρ. Then for any $k \in K$ and $x \in G$, we have $xkx^{-1} \in K$.

Proof 4.2.6

The proof is a simple computation: $\rho(xkx^{-1}) = \rho(x)\rho(k)\rho(x^{-1}) = \rho(x)1\rho(x^{-1}) = 1$. Therefore, xkx^{-1} is in the kernel of ρ.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

5: Groups IV

We look at quotient groups and group actions.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

5.1: Quotient Groups

Previously we saw product groups; now we'll learn about quotient groups. The construction is a bit more involved than the construction of product groups, just as division of natural numbers is a bit more complicated than multiplication...

Let's think for a moment about how quotients of natural numbers work, for the sake of building an imperfect analogy. When we write $\frac{n}{d} = q$, we have a numerator, denominator, and a quotient. The quotient q can be thought of as the number of times we can divide n into groups of d objects.

In making a quotient group, then, we would like to start with a group G, identify a subgroup H (the divisor) and do something to get a group $G/H = Q$. Using our analogy of dividing natural numbers, we would like to divide the group G into collections according to H. The notion of coset does this quite nicely, and in fact previously allowed us to see that the order of any subgroup H divides the order of G.

Definition 5.0.0

The set of cosets of a subgroup H of G is denoted G/H.

Then we can try to take the cosets of H as the underlying set of our would-be quotient group Q. The question is whether we can now identify a reasonable group operation on the set of cosets of H. The answer is 'sometimes!'

A Bad Choice of Product on Cosets

Suppose we have two cosets of H, aH and bH. We would like to define an operation on the two, so we naively write $aH \cdot bH = abH$, using the group operation in G to multiply a and b. And indeed, sometimes this works, but often it doesn't. What might cause it to fail? A problem arises because the set on which we're defining our new quotient group is the set of cosets, and it isn't generally obvious which element to take as the *representative* of the coset; ie, there is more than one way to write a coset as gH, and different choices might lead to different answers when we multiply our cosets.

Here's an example.

Take the group $G = D_5$, the symmetries of a pentagon generated by a flip f and a rotation r. Let H be the subgroup consisting of just the identity and the flip f. Then $H = \{1, f\}$. This subgroup has five different cosets; suppose we want to multiply the cosets $C = \{r, rf\}$ and $D = \{r^3, r^3 f\}$. Notice that there are two different ways to write C in the from gH: $C = rH$ and $C = rfH$. Each arises from a different *choice of representative* from H. The same is true for D: $D = r^3 H$ and $D = r^3 fH$. Depending on the choice of representatives, our rule for multiplying cosets then yields different answers. For example, $rH \cdot r^3 H = H$, but $rfH \cdot r^3 H = rfr^3 H = r^2 fH \neq H$.

Then we see that a more nuanced approach is necessary: in particular, our notion of a product shouldn't depend on a choice of coset representative!

? depend on the choice of coset representatives!

Products of Cosets

The initial idea for a product on cosets fell down because we were multiplying coset representatives, instead of thinking about how to multiply the actual cosets. So let's try to define an actual product of cosets!

Earlier, we saw what we might call *left cosets*, of the form $aH = \{ah_1, ah_2, \ldots\}$ where h_i are all the elements of H. But we can easily imagine *right cosets* as well, $Ha = \{h_1 a, h_2 a, \ldots\}$, and even *double cosets* $aHb = \{ah_1 b, ah_2 b, \ldots\}$. More generally, we can define *product of sets*: if $A = \{a_1, a_2, \ldots\}$, then AH is the set obtained by taking products of elements of A and H in every possible way: $AH = \{ah | a \in A, h \in H\}$. We can use the product of sets to compute explicit products of cosets.

Proposition 5.0.2

If H is a subgroup of G, then $HH = H$.

Proof 5.0.3

Since H is closed under the group operation, every element of HH is in H. Furthermore, since $1 \in H$, every element h in H appears in HH (for example, as $1h$). Then $HH = H$.

We build up these definitions so we can talk about products of cosets: $(aH)(bH)$. One fear of this approach is that taking a set-product like this may not give back a real left coset of H. In fact, sometimes it does and sometimes it doesn't!

If G is a commutative group, then right cosets and left cosets are the same thing: $ah_i = h_i a$ for every h_i, so $aH = Ha$. In this case, when examining products like $(aH)(bH)$, we have $aHbH = abHH = abH$. Then defining a product on cosets $(aH)(bH) = abH$ makes sense, and will end up giving a nice group structure. The identity is $1H$, associativity follows from the multiplication rule in G, and inverses are easy: $gHg^{-1}H = H$.

We should also check that this product doesn't depend on choice of coset representative. Suppose $aH = xH$, and consider the product $(aH)(bH) = aHHb = aHb$. Then notice that $(aH)(bH) = aHb = abH$, and $(xH)(bH) = xHb = (aH)b = abH$. Thus, we have $aH = bH$.

Normal subgroups

If we look closely at what we've just done, we didn't actually need G to be commutative: all that we needed was $aH = Ha$ for every $a \in G$. For example, we know this is true for the kernel of any homomorphism from the proposition in Section 4.2.

Definition 5.0.4: Normal Subgroups

A subgroup H of a group G is called a *normal subgroup* if $aH = Ha$ for every $a \in G$.

Then we've already proven the following theorem:

Theorem 5.0.5

Let H be a normal subgroup of G. Then GH is a group.

We'll also make explicit an earlier observation.

Proposition 5.0.6

Let G be a commutative group. Then every subgroup H of G is a normal subgroup, and G/H is a group.

We have already noticed that the kernel of any homomorphism is a normal subgroup. We can also define the *quotient map* $\pi : G \to G/H$, defined by $\pi(a) = aH$ for any $a \in G$. So long as the quotient is actually a group (ie, H is a normal subgroup of G), then π is a homomorphism. In fact, the kernel of π is exactly H. So we observe:

Corollary 5.0.7

A subgroup of G is normal if and only if it is the kernel of a homomorphism.

? Let G be a finite group and H a subgroup with $\frac{|G|}{|H|} = 2$. Show that H is a normal subgroup of G.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

5.2: Examples of Quotient Groups

Now that we've learned a bit about normal subgroups and quotients, we should build more examples.

Integers mod n, Again

Recall the group $\mathbb{Z}_n$. This can also be realized as the quotient group!

Let $n\mathbb{Z}$ denote the set of integers divisible by n: $n\mathbb{Z}=\{\ldots,-3n,-2n,-n,0,n,2n,3n,\ldots\}$. This forms a subgroup: 0 is always divisible by n, and if a and b are divisible by n, then so is $a+b$. Since every subgroup of a commutative group is a normal subgroup, we can from the quotient group $\mathbb{Z}/n\mathbb{Z}$.

To see this concretely, let $n=3$. Then the cosets of $3\mathbb{Z}$ are $3\mathbb{Z}$, $1+3\mathbb{Z}$, and $2+3\mathbb{Z}$. We can then add cosets, like so: $(1+3\mathbb{Z})+(2+3\mathbb{Z})=3+3\mathbb{Z}=3\mathbb{Z}$. The last equality is true because $3\mathbb{Z}=\{\ldots,-6,-3,0,3,6,\ldots\}$ so that $3+3\mathbb{Z}=\{\ldots,-3,0,3,6,9,\ldots\}=3\mathbb{Z}$.

? Write out addition tables for $\mathbb{Z}/5\mathbb{Z}$ as a quotient group, and check that it is isomorphic to $\mathbb{Z}_5$ as previously defined.

The Alternating Group

Another example is a very special subgroup of the symmetric group called the *Alternating group*, A_n. There are a couple different ways to interpret the alternating group, but they mainly come down to the idea of the *sign* of a permutation, which is always ± 1. The set $\{1,-1\}$ forms a group under multiplication, isomorphic to $\mathbb{Z}_2$. The sign of a permutation is actually a homomorphism. There are numerous ways to compute the sign or a permutation:

1. Determinants. A *permutation matrix* is the matrix of the linear transformation of n-dimensional space sending the i-th coordinate vector e_i to $e_{\sigma(i)}$. Such matrices have entries all equal to zero or one, with exactly one 1 in each row and each column. One can easily show that such a matrix has determinant equal to ± 1. Since the determinant is a multiplicative function - $\det(MN)=\det(M)\det(N)$ - we can observe the the determinant is a homomorphism from the group of permutation matrices to the group $\{\pm 1\}$.
2. *Count inversions.* An *inversion* in a permutation σ is a pair $i<j$ with $\sigma(i)>\sigma(j)$. For example, the permutation $[3,1,4,2]$ has $\sigma(1)>\sigma(2),\sigma(1)>\sigma(3)$ and $\sigma(3)>\sigma(4)$, and thus has three inversions. If there are i inversions, then the sign of the permutation is $(-1)^i$.
3. *Count crossings.* Draw a braid notation for the permutation where no more than two lines cross at any point and no line intersects itself. Then count the number of crossings, c. Then $s(\sigma)=(-1)^c$. The alternating group is the subgroup of S_n with $s(\sigma)=1$. (To prove that this method of counting works, one needs a notion of Reidemeister moves, which originally arise in the fascinating study of mathematical knots.)

? Find the inversion number for every permutation in S_4, and then sort the permutations by their inversion number.

? Show that each of the three definitions of the sign of a permutation give a homomorphism from S_n to $\{1,-1\}$. (For the third definition, a sketch of a proof will suffice. Be sure to argue that different braid notations for the same permutation give the same sign, even if the total number of crossings is different.)

We call a permutation with sign $+1$ a *positive* permutation, and a permutation with sign -1 a *negative permutation*.

? Show that there are $\frac{n!}{2}$ positive permutations in S_n.

Now we can define the alternating group.

Definition 5.1.4: Alternating Groups

The *alternating group* A_n is the kernel of the homomorphism $s:S_n\to\mathbb{Z}_2$. Equivalently, A_n is the subgroup of all positive permutations in S_n.

? Write out all elements A_4 as a subgroup of S_4. Find a nice generating set for A_4 and make a Cayley graph.

In fact, the alternating group has exactly two cosets. The quotient group S_n/A_n is then isomorphic to $\mathbb{Z}_2$.

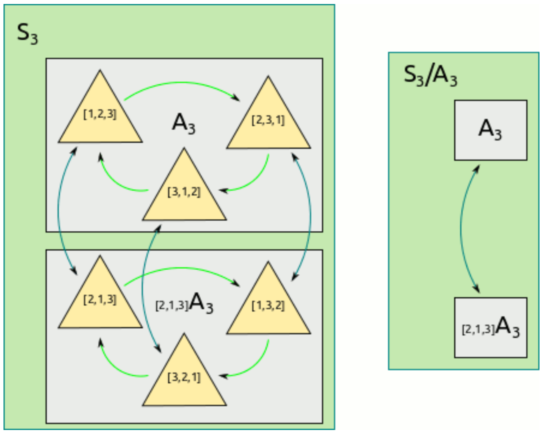

Figure 5.1.2: Quotient of S_3 by A_3.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

5.3: Isomorphism Theorem

We've observed a few cases now where we: 1. Define a homomorphism $\rho : G \to H$, and then 2. Notice that $G/K \sim H$, where K is the kernel of ρ. This isn't an accident!

The proof is just to build a correspondence between the cosets of the kernel gK and elements of the image I. Indeed, in any coset gK all elements map to the same element of the image. $\rho(gk) = \rho(g)\rho(k) = \rho(g)1 = \rho(g)$ for any $k \in K$.

This suggests a homomorphism from the set of cosets to the image: set $\phi(gK) = \rho(g)$. This is a homomorphism, since $\phi(ghK) = \rho(gh) = \rho(g)\rho(h) = \phi(gK)\phi(hK)$.

The map ϕ is also one-to-one: if $\phi(gK) = \phi(hK)$, we have $\rho(g) = \rho(h)$, so that $1 = \rho(g^{-1}h)$, meaning $g^{-1}h \in K$. Then $h = g(g^{-1}h) \in gK$, which tells us that $gK = hK$, since cosets are either equal or disjoint.

The map ϕ is onto, since any element in the image may be written as $\rho(g)$ for some g, which is also the image of gK under ϕ. Therefore, the map ϕ is an isomorphism.

TODO: Pictures!

This theorem is often called the "First Isomorphism Theorem." There are three isomorphism theorems, all of which are about relationships between quotient groups. The third isomorphism theorem has a particularly nice statement: $(G/N)/(H/N) \sim G/H$, which one can relate to the the numerical identity

$$\frac{\frac{n}{m}}{\frac{p}{m}} = \frac{n}{p}. \tag{5.3.1}$$

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

5.4: Classifying Finite Groups

We've seen that group theory can't distinguish between groups that are isomorphic. So a natural question is whether we can make a list of all of the groups!

We can make new groups from old groups using the direct product. So it would be nice to focus on groups that are **not** direct products. In the commutative case, this turns out to be pretty straightforward: a (finite) commutative group is a direct product of subgroups if and only if it has a proper subgroup.

The non-commutative case is much more difficult, though. There are actually a few other ways to build new groups from old groups; the most important of these other ways is the *semi-direct product*; we won't describe how to build semi-direct products here, but you can read about them elsewhere. Importantly, one can 'undo' a semi-direct product using a quotient, the same way one can undo a direct product. To get a sense of how useful the construction is, the symmetric group S_n is the semi-direct product of A_n and $\mathbb{Z}_2$. Also, the dihedral group D_n is a semi-direct product of $\mathbb{Z}_n$ and $\mathbb{Z}_2$.

An interesting question, then, is 'Which groups have no quotients?' We've seen that we can form a quotient group whenever there is a normal subgroup.

✎ Definition 5.3.0: Simple Groups

A group is *simple* if it has no proper normal subgroups. (A proper subgroup is any subgroup of G that is not equal to G or $\{1\}$, which are always normal subgroups.)

We'll now actually classify all of the finite simple groups, and discuss some of the history of the non-commutative case.

The Commutative Case

We can actually classify all of the finite commutative groups pretty easily. First, recall that *every* subgroup of a commutative group is normal.

Proposition 5.3.1

A finite commutative group is simple if and only if it has prime order p. In this case, it is isomorphic to the cyclic group, $\mathbb{Z}_p$.

Proof 5.3.2

If a finite commutative group has prime order then it has no proper subgroups, by Lagrange's theorem. Then it must be simple.
For the other direction, we assume G is a finite commutative simple group. G must be cyclic, or else we could form a proper subgroup by taking powers of a generator. So $G \sim \mathbb{Z}_n$ for some n. But if n is not prime we can find a subgroup using a proper divisor of n. Then $G \sim \mathbb{Z}_p$ for some prime p.

Theorem 5.3.3

Every finite commutative group is a direct product of cyclic groups of prime order.

Proof 5.3.4

Let A be a commutative group with n elements. Take any element x not equal to the identity in A; we know that there is some minimal integer m for which $x^m = 1$. Then A has a subgroup of order m generated by x, isomorphic to $\mathbb{Z}_m$. As a result, we have $A \sim A_1 \otimes \mathbb{Z}_m$, where A_1 is the quotient A/ \mathord \mathbb{Z}_m.
We can repeat that procedure indefinitely (taking an x in A_1 and writing A_1 as a product, and so on), until we obtain a decomposition $A = \mathbb{Z}_{m_1} \otimes \mathbb{Z}_{m_k}$, a product of cyclic groups.
We can then use the same trick to decompose each $\mathbb{Z}_m$ into a direct product of cyclic groups of prime order, completing the proof.

One can extend this trick to some infinite groups: those which have a finite number of generators. (Such groups, unsurprisingly, are called *finitely-generated*.) This gives rise to the *Fundamental theorem of finitely-generated commutative groups*.

? Suppose A is a finitely generated commutative group with infinite cardinality. Show that $A \sim \mathbb{Z} \otimes A'$, where A' is a finitely-generated commutative group.

The Non-Commutative Case

One of the major projects of 20th century mathematics research was to classify all of the finite simple groups; the project took fifty years, and the proof of the classification is estimated to span 10,000 pages written by over 100 authors. There's currently an effort underway to simplify the proof, however.

The classification shows that all finite simple groups are of one of four types:

1. Commutative groups of prime order,
2. Alternating groups A_n with $n \geq 5$,
3. Groups of Lie type,
4. The 26 sporadic groups.

We've already seen the first two types of simple group. It turns out that 'most' finite simple groups are in the third class, groups of Lie type, which are well beyond the scope of these notes to construct. Basically, though, groups of Lie type are certain groups of matrices with entries from a *finite field*, which are we'll see in the next chapter. The 'sporadic' groups are just those groups that don't fit into any of the other three classes!

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

6: Group Actions

In this chapter, we examine group actions and some fun applications of group theory.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

6.1: Group Actions

Group actions bring us back to our original view of groups as measures of symmetry. We begin with a definition.

Definition 6.0.1

Let G be a group. A set S is a (left) *G-set* if there is a function from $G \times S \to S$ (which we will write as $g \cdot s$ for $g \in G, s \in S$) satisfying:

1. $(gh) \cdot s = g \cdot (h \cdot s)$ for all $g, h \in G, s \in S$, and
2. $1 \cdot s = s$ for all $s \in S$.

An analogous definition can be written for a right G-set; a right G-set has a function from $S \times G \to S$.

These conditions say, in plain language, that elements of G move objects in the set to other objects, in a way which respects the group operation. (ie, it doesn't matter whether you perform the operation before or after applying the action.) Furthermore, the identity fixes every element of the set.

It's best to start with some easy combinatorial examples.

1. Let G be a group and S be any set. The *trivial action* of G on S is given by $g \cdot s = s$ for every $g \in G, s \in S$. One can easily check that the conditions for a group action hold!
2. Consider S_n and a set S of n labelled objects. Then the permutations of the objects constitute an action of S_n on S.
3. But S_n can also act on sets with more than n elements! Consider a regular deck of playing cards; each card has one of four *suits* (Clubs, Spades, Hearts, or Diamonds) and are numbered 1 to 13. (Where a Jack is 11, Queen is 12 and King is 13.) We can write each card in a short form: $4D$ is short for 'four of diamonds.' Then S_4 acts on the deck of cards by permuting the suits of the cards. For example, consider the permutation σ which transposes hearts and diamonds while leaving clubs and spades alone. Then $\sigma(4D) = 4H$, and $\sigma(12C) = 12C$.

   On the other hand, S_{13} acts on the *values* of the cards while leaving the suits alone. For convenience, we write the action on values of cards on the right and the action on suits on the left. Let τ be the permutation in S_{13} which sends i to $i+1$, and 13 to 1. Then $4D \cdot \tau = 5D$. Combining the left and right actions, we have $\sigma \cdot 4D \cdot \tau = 5H$, for example.
4. And S_n can act on sets with *fewer* than n elements. Consider a coin, with a 'heads' side (H) and a 'tails' side (T). We can define an action of S_n where σ flips the coin if the sign of σ is negative, and leaves the coin alone if the sign of σ is positive. Here the set we're acting on is actually the set of *states* of the coin: $\{H, T\}$.

Just as we made Cayley graphs of groups, we can also make Cayley graphs of group actions. If our group G has a generating set $\{g_1, g_2, \dots, g_k\}$, then the Cayley graph has one vertex for each element of the set S and colored edges from each s to $g_i \cdot s$.

For the example of S_n acting on a coin, the Cayley graph will just have two vertices, H and T, and arrows according to the action of the generators. If we consider the generating set of simple transpositions (which, we'll recall, exchange i and $i+1$ and leave everything else alone), the Cayley graph will have arrows from H to T and back for each simple transposition. (Because all simple transpositions have sign -1.)

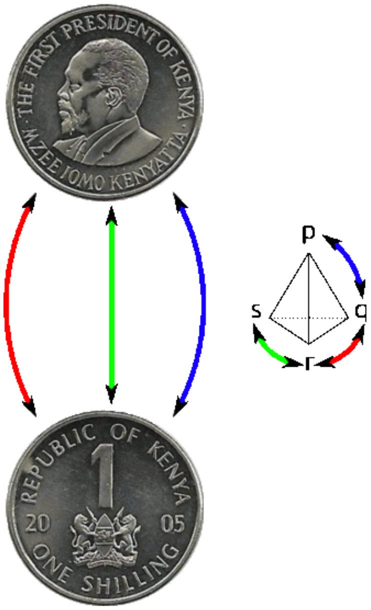

Figure 6.1: Cayley graph for action of S_4 on a coin, with generators given by the simple transpositions.

? S_n has another generating set consisting of the transposition $[2,1,3,4,\ldots,n]$ and the rotation $[n,1,2,\ldots,n-1]$. Draw the Cayley graph of S_n acting on the coin with these generators. (Note that the Cayley graph may be different for different n.)

Let's build a Cayley graph for the cards, but with a smaller set of cards to make things easier to draw. Imagine our deck only has two suits - Spades and Diamonds - and only numbers 1 through 4. Then let S_4 act on the left by permuting numbers. The Cayley graph is below.

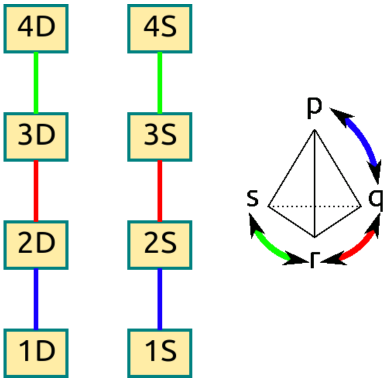

Figure 6.2: Cayley diagram for S_4 acting on some cards numbered 1 through 4, whose suits are all either diamonds (D) or spades (S).

Notice that the action is broken up into two different connected pieces: Exchanging numbers will never allow us to change suits, but we can switch around numbers freely. As a result, the suits each form a 'block' from which the other blocks can't be reached. These blocks are called *orbits*; we'll study them intensively in the next section.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

6.2: Orbits and Stabilizers

In this section, we'll examine orbits and stabilizers, which will allow us to relate group actions to our previous study of cosets and quotients.

Definition 6.1.0: The Orbit

Let S be a G-set, and $s \in S$. The *orbit* of s is the set $G \cdot s = \{g \cdot s \mid g \in G\}$, the full set of objects that s is sent to under the action of G.

There are a few questions that come up when encountering a new group action. The foremost is 'Given two elements s and t from the set S, is there a group element such that $g \cdot s = t$?' In other words, can I use the group to get from any element of the set to any other? In the case of the action of S_n on a coin, the answer is yes. But in the case of S_4 acting on the deck of cards, the answer is no. In fact, this is just a question about orbits. If there is only one orbit, then I can always find a group element to move from any object to any other object. This case has a special name.

Definition 6.1.1: Transitive Group Action

A group action is *transitive* if $G \cdot s = S$. In other words, for any $s, t \in S$, there exists $g \in G$ such that $g \cdot s = t$. Equivalently, S contains a single orbit.

Equally important is the stabilizer of an element, the subset of G which leaves a given element s alone.

Definition 6.1.2: The Stabilizer

The *stabilizer* of s is the set $G_s = \{g \in G \mid g \cdot s = s\}$, the set of elements of G which leave s unchanged under the action.

For example, the stabilizer of the coin with heads (or tails) up is A_n, the set of permutations with positive sign. In our example with S_4 acting on the small deck of eight cards, consider the card $4D$. The stabilizer of $4D$ is the set of permutations σ with $\sigma(4) = 4$; there are six such permutations.

In both of these examples, the stabilizer was a subgroup; this is a general fact!

Proposition 6.1.3

The stabilizer G_s of any element $s \in S$ is a subgroup of G.

Proof 6.1.4

Let $g,h \in G_s$. Then $gh \cdot s = g \cdot (h \cdot s) = g \cdot s = s$. Thus, $gh \in G_s$. If $g \in G_s$, then so is g^{-1}: By definition of a group action, $1 \in G_s$, so:
$s = 1 \cdot s = g^{-1} g \cdot s = g^{-1} s$.
Thus, G_s is a subgroup.

Group action morphisms

And now some algebraic examples!

1. Let G be any group and $S = G$. The *left regular action* of G on itself is given by left multiplication: $g \cdot h = gh$. The first condition for a group action holds by associativity of the group, and the second condition follows from the definition of the identity element. (There is also a *right regular action*, where $g \cdot h = hg$; the action is 'on the right'.) The Cayley graph of the left regular action is the same as the usual Cayley graph of the group!
2. Let H be a subgroup of G, and let S be the set of cosets G/H. The *coset action* is given by $g \cdot (xH) = (gx)H$.

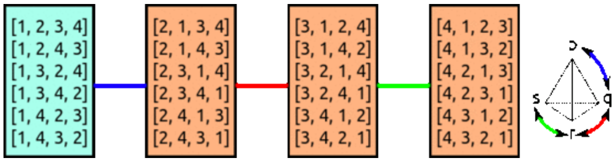

Figure 6.1: H is the subgroup of S_4 with $\sigma(1)=1$ for all σ in H. This illustrates the action of S_4 on cosets of H.

? Consider the permutation group S_n, and fix a number i such that $1 \leq i \leq n$. Let H_i be the set of permutations in S_n with $\sigma(i)=i$.

1. Show H_i is a subgroup of S_n.
2. Now let $n=5$ and Sketch the Cayley graph of the coset action of S_5 on H_1 and H_3.

The coset action is quite special; we can use it to get a general idea of how group actions are put together.

Proposition 6.1.6

Let S be a G-set, with $s \in S$ and G_s. For any $g,h \in G$, $g \cdot s = h \cdot s$ if and only if $gG_s = hG_s$. As a result, there is a bijection between elements of the orbit of s and cosets of the stabilizer G_s.

Proof 6.1.7

We have $gG_s = hG_s$ if and only if $h^{-1}g \in G_s$, if and only if $(h^{-1}g) \cdot s = s$, if and only if $h \cdot s = g \cdot s$, as desired.

In fact, we can generalize this idea considerably. We're actually identifying elements of the G-set with cosets of the stabilizer group, which is also a G-set; in other words, defining a function ϕ between two G-sets. The theorem says that this function preserves the group operation: $\phi(g \cdot s) = g \cdot \phi(s)$.

Definition

Let S, T be G-sets. A *morphism of G-sets* is a function $\phi : S \to T$ such that $\phi(g \cdot s) = g \cdot \phi(s)$ for all $g \in G, s \in S$. We say the G-sets are *isomorphic* if ϕ is a bijection.

We can then restate the proposition:

Theorem 6.1.9

For any s in a G-set S, the orbit of S is isomorphic to the coset action on G_s.

Now we can use LaGrange's theorem in a very interesting way! We know that the cardinality of a subgroup divides the order of the group, and that the number of cosets of a subgroup H is equal to $|G|/|H|$. Then we can use the relationship between cosets and orbits to observe the following:

Theorem 6.1.10

Let S be a G-set, with $s \in S$. Then the size of the orbit of s is $|G|/|G_s|$.

For a somewhat obvious example, considering S_{13} acting on the numerical values of playing cards, we can observe that any given card is fixed by a subgroup of S_{13} isomorphic to S_{12} (switching around the other twelve numbers in any way doesn't change affect the given card). Then the size of the orbit of the card is $|S_{13}|/|S_{12}| = 13$. That's a number we could have figured out directly by reasoning a bit, but it shows us that the theorem is working sensibly!

Now that we have a notion of isomorphism of G-sets, we can say something to classify G-sets. What kinds of actions are possible?

Let G be a finite group, and S a finite G-set. Then S is a collection of orbits. We knw that every orbit is isomorphic to G acting on the cosets of some subgroup of H. So we have the following theorem:

Theorem 6.1.11: Classification of G-Sets

Let G be a finite group, and S a finite G-set. Then S is isomorphic to a union of coset actions of G on subgroups.

For example, S_{13} acting on a full deck of cards decomposes as a union of four orbits, each isomorphic to the coset action of S_{13} on a subgroup isomorphic to S_{12}.

In short, to understand all possible G-sets, we should try to understand all of the subgroups of G. In general, this is a hard problem, though it's easy for some cases.

? Exercise 6.1.12

1. For $n = 15$, draw Cayley graphs of the coset action of $\mathbb{Z}_{15}$ on each of it's cosets.
2. Describe all the subgroups of $\mathbb{Z}_n$ for arbitrary n.

? S_n acts on subsets of $N = \{1, 2, 3, \ldots, n\}$ in a natural way: if $U = \{i_1, \ldots, i_k\} \subset N$, then $\sigma \cdot U = \{\sigma(i_1), \ldots, \sigma(i_k)\}$.

1. Decompose the action of S_4 on the subsets of $\{1, 2, 3, 4\}$into orbits.
2. Draw a Cayley graph of the action.
3. Identify each orbit with the coset action on a subgroup of S_4.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

6.3: Counting

We saw previously that the size of an orbit is equal to $|G|/|G_s|$. We can use this to figure out how many orbits there are in a G-set in all. This is a very useful thing to count: It's useful for counting things 'up to symmetry.' We'll denote the orbits of S by S/G, and thus the number of orbits is $|S/G|$. The notation should be read 'S-mod-G'; it's useful when two things are 'the same' in S if they are related by an application of an element of G.

For example, let's suppose we wanted to count all of the ways of painting the sides of a cube with three colors. We would naturally think of two ways of painting as being the same if we could rotate the cube to line up the colors in the same way. Then the group G might be symmetries of the cube, and the set S would be all ways of painting the cube in fixed position: What we're really trying to count is $|S/G|$.

Let S^g denote the set $\{s \in S \mid g \cdot s = s\}$. This is like the stabilizer in reverse: we're collecting up all of the elements of the set S that are fixed by g.

Theorem 6.2.0:

(Burnside's Lemma) The number of orbits in a G-set S is $|S/G| = \frac{1}{|G|}\sum_{g\in G}|S^g|$.

(Note that there's ample evidence that Burnside didn't actually invent Burnside's lemma; we include the name because it's what everyone knows it by.)

Proof 6.2.1:

Let $G \cdot s$ denote the orbit of s under G. First notice that the sum of the size of the fixed sets S^g is equal to the sum of the size of the stabilizer groups G_s: Both are counting the number of pairs (g, s) such that $g \cdot s = s$. Then:

$$\sum_{g\in G}|S^g| = \sum_{s\in S}|G_s| = \sum_{s\in S}|G|/|G\cdot s| = |G|\sum_{s\in S}\frac{1}{|G\cdot s|} = |G|\sum_{S/G}1 = |G||S/G|$$

Then dividing both sides by $|G|$ gives the desired result.

Let's try an example. We mentioned earlier the question of the number of ways to color a cube with three colors. Let's try it out. There is an initial question of which group of symmetries we're interested in: Do we allow reflections of the cube or only rotations? Since we can't naturally reflect things in three-dimensional space, we'll stick with the rotation group of the cube. (This choice, by the way, has consequences in chemistry. And here's an excellent Radiolab piece on the topic.)

The rotation group has 24 elements: From a base-position of the cube, you can rotate a marked face to any other face (there are six choices), and from there four rotations are available, making 24 symmetries in all. Every rotation in 3-dimensional space has an *axis of rotation*. So each rotational symmetry will have an axis of rotation; we can identify the symmetries by their axis and amount of rotation. We classify these symmetries into five classes, and determine the number of fixed points for each class.

Axes of symmetric rotation for a cube.

The types of rotational axis for a cube which produce symmetries. From left to right, a 'face' axis, a 'vertex' axis, and an 'edge' axis.

There are 3^6 colorings of the base cube to consider in all; for each symmetry, we determine the number of colorings that the symmetry fixes.

1. *The identity permutation.* This permutation fixes all 3^6 colorings.
2. *Face rotations by* $\pm 90^\circ$. These are formed by rotating around the axis through the center of two opposite faces. There are six such rotations. In order to fix a coloring, the coloring must have the four 'moving' sides all colored with the same color. The other two sides may be colored in any way. Thus, each of these symmetries fixes 3^3 colorings.
3. *Face rotations by* $\pm 180^\circ$. Rotation by 180° about the axis through opposite faces. There are three such symmetries. Each requires that opposite 'moving' faces be the same color, while the 'fixed' faces may be colored arbitrarily. Thus, there are 3^4 fixed points for these rotations.
4. *Vertex rotations by* $\pm 120^\circ$. Rotation through an axis between opposite vertices of the cube. There are four such axes, with two non-trivial rotations through each such axis, for a total of eight symmetries in this class. To fix a coloring, the coloring must have same-colored faces touching the vertices of rotation. There are thus 3^2 fixed points for these rotations.
5. *Edge rotations by* $\pm 180^\circ$. Rotation through the axis connecting the centers of opposite edges. There are six such rotations. To fix a coloring, such a rotation needs all opposite sides to have the same color. Thus, each one fixes 3^3 colorings.

We can then use Burnside's Lemma to find the total number of colorings up to rotation: $\frac{1}{G}(1 \cdot 3^6 + 6 \cdot 3^3 + 3 \cdot 3^4 + 8 \cdot 3^2 + 6 \cdot 3^3) = 57$.

? Exercise 6.2.2

1. Find the number of colorings of the cube with n colors up to rotation.
2. Find the number of colorings of the octahedron with n colors up to rotation.

? Up to rotation or flip, how many colorings of the board are possible? How about for a $k \times k$ board?

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

7: Rings I

A brief overview of the study of ring theory.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

7.1: Juggling With Two Operations

We'll now start looking at algebraic structures with more than one operation. Typically, these structures will have rules governing the different operations, and additional rules for how the operations interact. We'll begin by looking at rings, which have two operations, usually written as addition and multiplication, related by the distributive property.

There are many reasons to study ring theory, often having to do with generalizing the properties that we observe in many of the rings we deal with in day-to-day life, like the integers and the rational numbers. By making precise the algebraic structures that (for example) the integers satisfy, we can figure out what makes our favorite facts about the integers true, and easily see where those same facts hold true.

It's also an area where most of the real pay-off comes later. Understanding ring theory is essential for algebraic geometry in particular, which is a major force in modern mathematics. The basic idea of algebraic geometry is to study geometry using zeroes of polynomials: for example, a line in the plane can be thought of as the zeroes of the polynomial $f(x,y)=y-mx-b$ where m and b are constants. In other words, to understand properties of geometry, it is helpful to understand properties of polynomials. And polynomials are an example of a ring, as we'll see.

Definition 7.0.0

A *ring* is a set R with operations $+$ and $\cdot$ such that:

1. R is a commutative group under $+$,
2. (Distributivity) For all $r,s,t \in R$, we have $r\cdot(s+t)=r\cdot s+r\cdot t$, and $(s+t)\cdot r=s\cdot r+t\cdot r$.

Show, using the definition of a ring, that for any ring R with additive identity 0, we have $0\cdot r=0$ for every $r\in R$.

This is the most general type of ring. There are many different types of ring which arise from placing extra conditions, especially on the multiplicative operation. In fact, ring theory is kind of a zoo, divided up into the study of different 'species' of rings. Possibly the most important rings to study are commutative, associative rings with unity, which we define now.

Definition 7.0.2:

Let R be a ring, and $r,s,t\in R$. Then R is:

1. *Associative* if the multiplication operation is associative: $r\cdot(s\cdot t)=(r\cdot s)\cdot t$,
2. A *ring with unity* if there is a multiplicative identity 1, such that $1\cdot r=r=r\cdot 1$,
3. *Commutative* if the operation $\cdot$ is commutative: $r\cdot s=s\cdot r$.

Usually we'll deal with associative rings with unity; in fact, when we write 'ring' we'll mean an associative ring with unity unless otherwise noted. As a result, 'commutative ring' will mean a ring that is commutative, associative and with unity.

There are numerous examples of rings! Here are some familiar examples.

1. *Integers*. The integers are a commutative group under addition, and have the distributive property. Additionally, the integers are associative and commutative under multiplication, and have a multiplicative identity, 1. Thus, the integers are an commutative associative ring with unity.
2. *Rational Numbers, Real Numbers, Complex Numbers*. All of these familiar number systems are examples of commutative associative rings with unity.
3. *Integers modulo* n, $\mathbb{Z}_n$. The multiplication operation works just as the addition operation does: do the normal multiplication, and then divide by n and keep the remainder: $a\cdot b=(ab)$. This is an associative and commutative operation, and there is a multiplicative identity.
4. *Matrices*. Recall that matrix addition is just entry-by-entry, and that the multiplication of matrices adds and multiplies the entries according to a certain rule: if M and N are matrices, then $(MN)_{i,j}=\sum_k M_{i,k}N_{k,j}$. Since this only uses addition and multiplication, we can thus form matrices with entries in any ring R, since R has notions of addition and multiplication. The set of all $m\times n$ matrices with entries in R is denoted $M_{m\times n}(R)$.

5. *Polynomials*. Polynomials can be added and multiplied so long as we know how to add and multiply the coefficients. We let $R[x]$ denote the ring of polynomials with coefficients from the ring R and variable x with exponent ≥ 0. For example, if $R=\mathbb{Z}_2$, we have $(x+1)(x+1)=x^2+1$.
6. *Rings of Functions*. Many spaces of functions have a ring structure. For example, if we consider differentiable functions $\mathbb{R}\to\mathbb{R}$, we can add and multiply functions: $(f+g)(x)=f(x)+g(x)$ and $(f\cdot g)(x)=f(x)g(x)$. Sums and products of differentiable functions are also differentiable, so they are closed. The functions form an additive group, and there's a multiplicative identity: the constant function defined by $1(x)=1$.

? Exercise 7.0.3

1. Generate two 'random' matrices M and N in $M_{3,3}(\mathbb{Z}_6)$. Compute $M+N$, MN, and NM.
2. Consider $f,g\in\mathbb{Z}_6$, defined by $f=x^3+2x^2+3x$ and $g=4x^3+5x+4$. Find $f+g$ and fg.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

7.2: Ring Homomorphisms

As we saw with both groups and group actions, it pays to consider structure preserving functions!

Definition 7.1.0

Let R and S be rings. Then $\phi : R \to S$ is a *homomorphism* if:

1. ϕ is homomorphism of additive groups: $\phi(a+b) = \phi(a) + \phi(b)$, and
2. ϕ preserves multiplication: $\phi(a \cdot b) = \phi(a) \cdot \phi(b)$.

If the homomorphism is a bijection, then it is an *isomorphism*.

Examples:

1. We have the *inclusion* homomorphism $\iota : \mathbb{Z} \to \mathbb{Q}$, which just sets $\iota(n) = n$. This map clearly preserves both addition and multiplication.
2. Consider the map $\phi : \mathbb{Z} \to \mathbb{Z}_n$ sending k to k. We've seen that this is a homomorphism of additive groups, and can easily check that multiplication is preserved. Indeed, $\phi(a) = \phi(1+1+\cdots+1) = \phi(1)+\phi(1)+\cdots+\phi(1) = a\phi(1) = a$.

   Notice that every element in $\mathbb{Z}$ can be written as a sum of many copies of 1. Then we were able to figure out what the homomorphism does simply by knowing $\phi(1)$. As an example, consider the map $\rho : \mathbb{Z} \to \mathbb{Z}_5$ sending k to $(2k)$. (Thus, $\rho(0) = 0, \rho(3) = 1$.) This can be shown, using the same argument as above, to be a ring homomorphism.
3. The *evaluation map* e_k is a function from $R[x]$ to R. For any polynomial $f \in R[x]$ and $k \in R$, we set $e_k(f) = f(k)$. This is a ring homomorphism! Let $f(x) = a_n x^n + \cdots a_0 x^0$, and $g(x) = b_n x^n + \cdots b_0 x^0$, where the $a_i, b_i \in R$. (We'll also allow leading coefficients to be zero in order to make it easy to add f and g formally.) We then check the ring homomorphism conditions:

   b. Since we know that e_k is an additive homomorphism, we only need to check that it is multiplicative on monomials. But that's easy:

$$e_k((ax^n)(bx^m)) = e_k(abx^{n+m}) \tag{7.2.1}$$

$$= abk^{n+m} = e_k(ax^n)e_k(bx^m). \tag{7.2.2}$$

Exercise 7.1.1

1. Show that $\rho : \mathbb{Z} \to \mathbb{Z}_5$ defined by $\rho(k) = (3k)$ is a ring homomorphism. Find the kernel and image of ρ.

Show that $\rho : \mathbb{Z} \to \mathbb{Z}_6$ defined by $\rho(k) = (3k)$ is a ring homomorphism. Find the kernel and image of ρ. As with groups, we also have direct products of rings.

Definition 7.1.2

Let R and S be rings. Define the *direct product* $R \times S$ as the set $\{(r,s) \mid r \in R, s \in S\}$ with coordinate-wise operations: $(r_1, s_1) + (r_2, s_2) = (r_1 + r_2, s_1 + s_2)$, and $(r_1, s_1) \cdot (r_2, s_2) = (r_1 \cdot r_2, s_1 \cdot s_2)$.

Of course, one should verify that this is a ring by checking the ring axioms.

Exercise 7.1.3

1. Show that for any rings R and S that the product $R \times S$ is a ring.
2. Show that the *inclusion map* $\iota : R \to R \times S$ given by $\iota(r) = (r, 0)$ is a ring homomorphism.
3. Show that the *projection* $\pi : R \times S \to R$ given by $\pi((r,s)) = r$ is a ring homomorphism.

A Small Digression on the Relationship Between Good Computer Science and Good Mathematics

Recall that when we worked with groups the kernel of a homomorphism was quite important; the kernel gave rise to normal subgroups, which were important in creating quotient groups.

For ring homomorphisms, the situation is very similar. The kernel of a ring homomorphism is still called the *kernel* and gives rise to *quotient rings*. In fact, we will basically recreate all of the theorems and definitions that we used for groups, but now in the context of rings. Conceptually, we've already done the hard work.

In computer programming, people often speak of the DRY principle: *Don't Repeat Yourself*, meaning that you shouldn't write the same code more than once. The reason, in computer science, is that it's easier to fix mistakes or make modifications if a given piece of code appears in one distinct place.

In mathematics, we have a similar principle: *generalization*. When you find yourself doing the same thing in different contexts, it means that there's something deeper going on, and that there's probably a proof of whatever theorem you're re-proving that doesn't matter as much on the context. It would be nice, for example, to remember just one concept for quotient groups, quotient rings, quotient vector spaces, and whatever else, instead of a hodgepodge of specific cases of the same basic idea.

For the game of homomorphisms, kernels, and quotients, the generalization involves *category theory* and *universal properties*. Category theory is a bit beyond the scope of these notes, but is an essential part of modern mathematics and serves as a bridge between many different fields of mathematical study.

Subring, Kernel, Image, Quotient.

We begin with some definitions.

Definition 7.1.4

Let R be a ring. A subset S of R is a *subring* if S is itself a ring using the same operations as R. (We don't require that S has a multiplicative identity, though.)

For example, take $R[x]$, the polynomial ring over R. The set of degree 0 polynomials is closed under addition and multiplication; indeed, this set is just a copy of R. Thus, R is a subring of $R[x]$.

On the other hand, consider the set of all polynomials of degree greater than or equal to 2 in $\mathbb{Z}[x]$, which we'll denote $P_{\geq 2}$. This is closed under addition (the sum of two polynomials has degree equal to the max of their degrees), and is closed under multiplication (the degree of the product is the sum of the degrees). Thus, it is a subring. However, the multiplicative identity in $R[x]$ is 1, which has degree 0. So there is no unit in $P_{\geq 2}.

Another example: Take $2\mathbb{Z} \subset \mathbb{Z}$, the set of even integers. This set is closed under addition and multiplication, and is thus a subring. (The sum and product of two even integers is still even.) However, the even integers don't have the number 1, and so there is no unit in $2\mathbb{Z}$.

? Let $P^n_{\geq 2}$ denote all polynomials in $\mathbb{Z}_n[x]$ with degree ≥ 2. Is $P^4_{\geq 2}$ a subring of $\mathbb{Z}_n[x]$? Why or why not?

Definition 7.1.6

Let $\phi : R \to S$ be a ring homomorphism. The *kernel* of ϕ is $\{r \in R \mid \phi(r) = 0\}$, which we also write as $\phi^{-1}(0)$. The *image* of ϕ is the set $\{\phi(r) \mid r \in R\}$, which we also write as $\phi(R)$.

We immediately have the following.

Proposition 7.1.17

Let $\phi : R \to S$ be a ring homomorphism. Then the kernel of ϕ is a subring of R and the image of ϕ is a subring of S.

Proof 7.1.8

Since ϕ is a homomorphism of commutative additive groups, we know that the kernel and image are closed under addition. The kernel is closed under multiplication, because if $\phi(a) = \phi(b) = 0$, then $\phi(ab) = \phi(a)\phi(b) = 0$. The image is closed because if $x, y \in \phi(R)$, then there exist $a, b \in R$ such that $\phi(a) = x, \phi(b) = y$. Then $xy = \phi(a)\phi(b) = \phi(ab) \in \phi(R)$.

Just as kernels of group homomorphisms were special kinds of subgroups, kernels of ring homomorphisms are special kinds of subrings.

Definition 7.1.9

A subring I of a ring R is an *ideal* if for any $x \in I, r \in R$, $rx \in I$ and $xr\in I.

Proposition 7.1.10

Let K be the kernel of a ring homomorphism $\phi : R \to S$. Then K is an ideal.

Proof 7.1.11

For any $x \in K$, we have $\phi(x) = 0$. Then $\phi(rx) = \phi(r)\phi(x) = \phi(r)0 = 0$. Similarly, $\phi(xr) = 0$. Thus, K is a two-sided ideal.

Ideals are playing exactly the same role as normal subgroups in the groups context; in fact, an ideal **is** a normal subgroup of the additive group of the ring. In particular, we can form cosets and consider the quotient R/I. Since it's an additive group, cosets of an ideal I are of the form $r + I = \{r + x | x \in I\}$.

Theorem 7.1.12

If I is an ideal, then R/I is a ring.

Proof 7.1.13

We know that under addition R/I is a commutative group. So we just need to show that the multiplication distributes over addition. For this we have:
$((r + I) + (q + I))(s + I) = rs + qs + I = (r + I)(s + I) + (q + I)(s + I)$.
One can also check that the multiplication is associative and commutative if R is associative and commutative. Likewise, if R has a unit, then $1 + I$ acts as a unit in R/I.

Finally, we have the isomorphism theorem.

Theorem 7.1.14: Isomorphism Theorem

Let R and S be rings, and $\phi : R \to S$ a homomorphism. Then the image of ϕ is isomorphic to R/I.

Proof 7.1.15

To prove the isomorphism theorem, build a homomorphism from R/I to the image of ϕ, just as we did for groups, and show that it is a bijection.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

8: Rings II

In working with different rings (and different kinds of rings) questions quickly arise about which familiar properties of one ring might carry over to another ring. To illustrate this kind of question, we'll spend this chapter talking about division.

Topic hierarchy

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

8.1: The Problem of Division

Let's consider the problem of division. To get at a notion of division in general rings, let's recap what we know about division for familiar number systems.

Definition 8.0.0:

For numbers x, y, we say that y *divides* x if there exists a number z such that $\frac{x}{y} = z$ if $x = zy$. We call z the *quotient* of x by y.

We can try to offload the problem of division to a problem of finding multiplicative inverses. If y has a multiplicative inverse, then division by y is easy: we can set $z = xy^{-1}$, so that $zy = xy^{-1}y = x$. If every element other than 0 has a multiplicative inverse, then R is called a field. You should already know three examples of fields: $\mathbb{Z}$, $\mathbb{R}$, and $\mathbb{C}$. Part of the reason for the importance of fields is that most of the basic facts in linear algebra work for any field.

Definition 8.0.1: Field

A *field* F is a commutative ring in which every element other than 0 has a multiplicative inverse.

Since the field must already be a commutative, associative ring with unity, we see that the set $F \setminus \{0\}$ is a group! Then another way to define a field is as a ring that is a commutative group under addition, and where $F \setminus \{0\}$ is a commutative group as well.

For example, we've already seen the group $\mathbb{Q}^\times$, which is just $\mathbb{Q}$ with the 0 removed. Since this is a commutative group, $\mathbb{Q}$ is a field.

✓ For which values of n is $\mathbb{Z}_n$ a field?

All of this works fine in $\mathbb{R}$, $\mathbb{Q}$ and $\mathbb{C}$: in these rings, for every x and y, we can find a unique number z such that $zy = x$. In other rings, though, things can go wrong in a number of different ways.

1. The first problem that could arise is that y has no multiplicative inverse. For example, in $\mathbb{Z}_6$, there is no number z such that $2 \cdot z = 1$. Likewise, almost no element of $\mathbb{Z}$ has a multiplicative inverse.
2. It could be that for a given x and y, there is no quotient $z = \frac{x}{y}$. An example of this occurs in $\mathbb{Z}$, where (for example) there is no number $\frac{2}{3}$.
3. It could happen that the quotient z exists but is not unique. For example, consider the product ring $\mathbb{Z} \times \mathbb{Z}$. Let $x = (4, 0)$ and $y = (2, 0)$. Then for any integer k, $(2, k) \cdot y = x$.
4. There's also a problem if the ring R is not commutative. It could occur that $yz = x$ but $zy \neq x$. Which 'side' of x is our division happening on?

We'll see that the different ways of resolving these questions give rise to definitions of different kinds of rings.

Zero-divisors

We'll first consider the question of multiplicative inverses. For a start, in any non-zero ring, 0 does not have a multiplicative inverse: For any x we have have $x \cdot 0 = 0$, so it can't be the case that $x \cdot 0 = 1$. This situation is familiar from working with the rational and real numbers. But there can be other elements without a multiplicative inverse.

For example, consider $\mathbb{Z}_6$. The elements 1 and 5 have multiplicative inverses: $1 \cdot 1 = 1$ and $5 \cdot 5 = 1$. But none of the other elements have a multiplicative inverse! For example, if we multiply 2 times each element of $\mathbb{Z}_6$, we get the list $[0, 2, 0, 2, 0, 2]$ Since 1 isn't in the list, 2 has no multiplicative inverse. Something interesting is happening in that list of multiples of 2, though: there are many zeroes!

Definition 8.0.3

Let $x \in R$. Then x is a *zero-divisor* if there exists y such that $x \cdot y = 0$.

? Describe all of the zero-divisors in the ring $\mathbb{Z} \times \mathbb{Z}$.

We have the following immediate result.

Proposition 8.0.5

For $x \in R$, x cannot be both invertible and a zero-divisor.

Proof 8.0.6

Suppose that x is invertible and a zero-divisor, and let $y \neq 0$ with $xy = 0$. Then $y = (x^{-1}x)y = x^{-1}(xy) = x^{-1}0 = 0$, a contradiction.

As a result, the presence of zero-divisors means that there are non-invertible elements in the ring, and thus throws our division project into jeopardy. Furthermore, zero divisors also contribute to non-uniqueness of division: if $ry = 0$ and $x = zy$, then we also have $x = (z + r)y$, so that both z and $z + r$ can be considered as solutions to $\frac{x}{y}$.

To give a concrete example of this phenomenon, consider again $\mathbb{Z}_6$. What is the quotient $\frac{4}{2}$? Obviously, 2 is an answer, since $2 \cdot 2 = 4$. But we also have $2 \cdot 5 = 4$, so we could also write $\frac{4}{2} = 5$ just as easily. Notice that $(2 + 3) \cdot 2 = 4 + 0 = 4$; this is exactly the case described above.

Interestingly, for elements which are neither invertible nor zero-divisors, we still have a cancelation law:

Corollary 8.0.7

Suppose that $r \neq 0$ is not a zero-divisor and $rx = ry$. Then $x = y$.

Proof 8.0.8

We have $rx - ry = r(x - y) = 0$. Then since r is not a zero-divisor, we must have $x - y = 0$, so that $x = y$.

One can use this result directly to prove the following:

Corollary 8.0.9

If r is not a zero divisor, then the quotient $\frac{x}{r}$ is unique if it exists.

Then we see that the presence of zero-divisors is a major impediment to doing division in rings. Rings without zero-divisors will then be nice to work with!

Definition 8.0.10: Integral Domain

A commutative ring with no zero-divisors is called an *integral domain*.

Every field is an integral domain, since every non-zero element of a field is invertible. The primary example of an integral domain that is not a field is the integers: There are no non-zero integers where $xy = 0$, but most integers don't have multiplicative inverses, so $\mathbb{Z}$ is not a field.

Then we seem to have an answer to the problem of division for commutative rings:

1. The best-case scenario is when every element has an inverse. Such rings are called division rings, or (if the ring is also commutative) fields.
2. The next-best case is when there are no zero divisors. These are the integral domains.

In the next two sections, we'll look at two different ways to 'solve' a division problem in an integral domain. The first way is to introduce fractions, which allow us to find inverses for any element of the ring. The second - available only in some rings - allows us to do division with a remainder.

? Show that $M_{n,n}(\mathbb{Q})$, the ring of $n \times n$ matrices, is not an integral domain.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

8.2: Field of Fractions

In the history of number systems, there is a clear progression: Faced with a void where there could be more numbers, more numbers are invented. First came the natural numbers (the counting numbers), and when people found that certain subtraction problems had no solution, negative numbers were introduced to fill the void. Relevant to our current discussion, the rational numbers come about when one notices that some division problems amongst integers don't have a solution.

Suppose that R is an integral domain. Then the only impediment to division is a lack of actual quotients: if the quotients were to exist, they would have to be unique.

Consider $\mathbb{Z}$. This is of course an integral domain, but wouldn't it be nice if 2 had a multiplicative inverse? We'll *extend* $\mathbb{Z}$ by including $\frac{1}{2}$. But when we include $\frac{1}{2}$, we also have to include all possible sums and products in order to ensure that we still have a ring; the operations of addition and multiplication need to be closed, of course. So in addition to $\frac{1}{2}$, we also need to include every number $\frac{n}{2^m}$, with $n \in \mathbb{Z}$ and $m \in \mathbb{N}$ in order to ensure that the set is closed under multiplication and addition. Call this set R. Then R is a commutative ring with unity. It's also an integral domain, but still not a field, since, for example, 3 has no multiplicative inverse.

In that case, we can go ahead and include the multiplicative inverse of *every* positive integer, along with all possible sums and products of those inverses. The resulting ring, of course, is the rational numbers, $\mathbb{Q}$.

We would like to extend this construction to an arbitrary integral domain: Starting from an integral domain D, we introduce inverses and the appropriate sums and products until every element has an inverse. In fact, this involves copying the whole notion of fractions.

First we construct a ring D'. For an integral domain D, D' is the set $D \times (D \setminus \{0\})$. Each pair (a, b) in D' can be thought of as fractions $\frac{a}{b}$; note that we disallow $b = 0$. The operation $+$ defined by $(a, b) + (x, y) = (ay + bx, by)$ and multiplication is defined by $(a, b) \cdot (x, y) = (ax, by)$. Then D' is a commutative ring; we leave it as an exercise to show this is true.

? Let D be an integral domain. Show that D' is an integral domain. (In particular, check all of the ring axioms, and then show that there are no zero-divisors in D'.)

With rational numbers, it is important to notice that many different fractions are the same: currently our ring D' is much too large! For example, we haven't introduced any mechanism for cancellation of numerator and denominator: $(a, b) \cdot (b, a) = (ab, ab) \neq 1$ in D'.

We'll construct the actual ring of fractions as a quotient of D'. To construct a quotient, we only need to identify a suitable ideal I; the quotient will then just be D'/I. The ideal should contain everything that is 'equivalent to 0' in the ring of fractions. Thinking by analogy to the rationals, we see that this is the set $I = \{(0, x) \mid x \in D\}$.

This set is easily shown to be an ideal: $(0, a) + (0, b) = (0ab, ab) = (0, ab) \in I$. And for any $(x, y) \in D'$, we have $(x, y) \cdot (0, a) = (0, ya) \in I$. Then I is an ideal.

We can now check that $(a, b) \cdot (b, a) = (1, 1)$ in the quotient D'/I: $(a, b) \cdot (b, a) = (ab, ab)$, and $(ab, ab) - (1, 1) = (ab - ab, ab) = (0, ab) = 0$, as desired.

We then define the field of fractions as $Q = D'/I$. This is in fact a field: For any $a, b \neq 0$, we have $(a, b)^{-1} = (b, a)$, so every non-zero element in Q has a multiplicative inverse.

Definition 8.1.1: Field of Fractions

The *field of fractions* of an integral domain D is D'/I, with D' and I as defined above.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

8.3: Euclidean Domains

Creating a field of fractions is one way to definitively solve the problems of division in an integral domain: Make up fractions to have an inverse for every non-zero element. But there's (sometimes) another way to define division without resorting to introducing new elements to the field, familiar from the integers: define division using a 'quotient' and a 'remainder.'

For example, among the integers we can write $25 = 8 \cdot 3 + 1$; then $25/3$ has a quotient 8 and remainder 1. Generally, to find n/m, we write $n = qm + r$, where $0 < r < |m|$. Then q is the quotient and r is the remainder.

We can do something similar with polynomials: Given two polynomials f and g, we can divide f by g and **uniquely** write $f = Qg + R$, where Q is a polynomial and R is a polynomial of lower degree than g.

For example, take $f = 2x^5 + 3x^2 + x + 3$ and $g = x^2 + 1$, we can apply the polynomial long division algorithm and get $f = (2x^3 - 2x + 3)g - x$. Here $2x^3 - 2x + 3$ is the whole part and $-x$ is the remainder.

In both the integer division and the polynomial division, the key ingredient is a way of *ordering* the elements of the ring: in the integers, we order by the usual ordering of the integers, and with polynomials we order by the degree of the polynomial.

Definition 8.2.0: Norm on a Ring

A *norm* on a ring R is a function $n : R \to \mathbb{Z}_{\geq 0}$ with $n(0) = 0$. A *positive norm* has $n(r) > 0$ for all $r \neq 0$.

Any given ring can have many different norms. The norm on the integers is simply the absolute value of the integer; it is a positive norm. The norm on polynomials is the degree of the polynomial.

Definition 8.2.1: Euclidean Domain

A *Euclidean domain* is an integral domain R with a norm n such that for any $a, b \in R$, there exist q, r such that $a = q \cdot b + r$ with $n(r) < n(b)$. The element q is called the *quotient* and r is the *remainder*.

A Euclidean domain then has the same kind of partial solution to the question of division as we have in the integers.

In fact, Euclidean domains further have a *Euclidean algorithm* for finding a common divisor of two elements. The Euclidean algorithm is performed by starting with two elements f and g for which we wish to find a common divisor. Dividing f by g gives a quotient q_0 and a remainder r_0. We then divide g by r_0 and obtain a new quotinet q_1 and a new remainder, r_1. We then repeat this process to get quotients $q_2, q_3, \ldots q_k$ and remainders $r_2, r_3, \ldots r_k$. Each remainder has smaller norm than the previous, so this process must eventually terminate with some $r_k = 0$.

The final quotient, q_k divides both g and f: You can see this by writing $f = q_0 g + r_0$, and then expanding r_0: $f = q_0(q_1 r_0 + r_1) + r_0$. If we imagine the process ending at this point, so that $r_1 = 0$, we then have r_0 divides both f and g. On the other hand, if the process doesn't terminate, we can expand $r_0 = q_2 r_1 + r_2$. Then $f = q_0(q_1(q_2 r_1 + r_2) + r_1) + q_2 r_1 + r_2$. If the process terminates, then $r_2 = 0$, and r_1 divides every term, and thus divides f and g. If the process doesn't terminate, we repeat the same basic argument.

(TODO: Examples in Z and Z[x])

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

CHAPTER OVERVIEW

9: Vector Spaces

We look at vector spaces as algebraic structures.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

9.1: A Return to Linear Algebra

We've now seen numerous examples of *algebraic structures*, which we can think of as sets with some operations which satisfy some axioms. Here's a partial list:

1. Groups,
2. Commutative groups,
3. Group actions,
4. Rings,
5. Commutative rings,
6. Integral domains,
7. Fields,
8. and others...

In this chapter, we'll examine *vector spaces* as algebraic structures. Vector spaces are massively important because these are algebraic structures where the tools of linear algebra are available. Linear algebra is, in some ways, the branch of mathematics which is best developed: when a problem in science is converted into a linear algebra problem, we have a pretty good chance of being able to solve it. This is why, for example, the technique of *linearization* which comes up in differential equations and modeling is so important.

In fact, viewing vector spaces as algebraic structures does two things for us.

1. This viewpoint helps us identify more situations as linear algebra situations, allowing us to use our linear algebra tools in a broader set of circumstances, and
2. Abstracting allows us to better identify precisely what tools we are using when we prove statements in linear algebra, so we can identify exactly which situations those tools are applicable in. As with rings, there are more than one kind of vector space, and some vector spaces are more 'friendly' than others.

So let's see the definition.

Definition 9.0.0: vector space properties

A *vector space* is a set V and a field k with two operations, *addition* $+ : V \times V \to V$ and *scalar multiplication* $\cdot : k \times V \to V$, satisfying the following axioms.

1. V under addition is a commutative group.
2. (Distributivity I) For any $c \in k$ and $v, w \in V$, we have $c(v+w) = cv+cw$.
3. (Distributivity II) For any $c, d \in k$ and $v \in V$, we have $(c+d)v = cv+dv$.
4. (Associativity) For any $c, d \in k$ and $v \in V$, we have $(cd)v = c(dv)$.

The elements of the set V are called *vectors*.

(As an aside: There's a another way to think of vector spaces as well. For any ring R, there is a concept of an *R-module* which is similar to a group action: a module is a set with a ring action. That is to say, a ring pushing around objects in the set in a way that is compatible with both of the ring operations. From this viewpoint, a vector space is just a k-module, where the underlying set is a commutative group itself. As a result, R-modules is a generalization of vector spaces.)

As is traditional, we list some examples. Note that the vector space is a set *and* a field: usually, the choice of field is derived from context, but we'll be specific if the context is non-obvious. Often, we say that 'V is a vector space over k' to mean that V is the commutative group and k is the field.

1. $\mathbb{k}^n$ is the vector space whose underlying set is lists of n elements of k, with coordinate-wise addition and k acting by scalar multiplication. This gives rise to the familiar spaces $\mathbb{R}^n$ and $\mathbb{C}^n$. But we also know about finite fields now: $\mathbb{Z}_p^n$ where p is prime is also a vector space.
2. The set of polynomials $k[x]$ in a single variable is a vector space over k.
3. Let $M_{n,m}(k)$ denote the set of $n \times m$ matrices with entries in k. Then $M_{n,m}(k)$ is a vector space over k.
4. Let V be a vector space over k. Set $V*$ to be the set of functions from V to k. (This is called the *dual* of V.) Addition of functions is given by $(f+g)(x) = f(x)+g(x)$, and scalar multiplication is given by $(cf)(x) = c \cdot f(x)$.

? Exercise 9.0.1

1. For each of the above examples of vector spaces, write some example elements and give examples of addition and scalar multiplication in that vector space.
2. Prove that each of these examples is a vector space.

Some kinds of vector spaces only make sense with certain fields. Here's an example in the form of an exercise.

? Exercise 9.0.2

Show that the set of continuous functions from $\mathbb{R} \to \mathbb{R}$ is a vector space over $\mathbb{R}$. (Be sure to explicitly identify what the operations of addition and scalar multiplication are.)

What extra condition would we need for a vector space V over k in order for the notion of continuous functions $V \to k$ to make sense?

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

9.2: Linear Independence

One of the core concepts in linear algebra is *linear independence*, and this concept translates to general vector spaces with no difficulty.

Definition 9.1.0

Let S be a set with elements s_i. A *linear combination* of elements $\{s_1, s_2, \dots, s_n\}$ is given by any **finite** sum $\sum_{s\in S} c_s s$ with coefficients $c_s \in k$. (If S is an infinite set, then all but finitely many c_s must be equal to 0.)

Definition 9.1.1

Let S be a set of vectors in a vector space V. Then we say that S is *linearly dependent* if there exists a linear combination of elements of S equal to 0.

Example

Let $\mathbb{R}^\infty$ be the vector space of sequences of elements of $\mathbb{R}$. (ie, the space of sequences $r = (r_1, r_2, r_3, \dots)$, with coordinate-wise addition and the usual scalar multiplication.) Let $r_i \in \mathbb{R}^\infty$ be the sequence with $(e_i)_i = 1$ and $(e_i)_j = 0$ for all $j \neq i$. Let n be the element $(-1, -1, -1, \dots)$. Now, let S be the set of all the e_i and n. This is actually a linearly independent set. You might note that the sum of all of the elements in S (with all coefficients in the sum equal to 1) seems to be the 0-vector. But this is an infinite sum, and is thus not considered a linear combination of elements of S.

Contributors and Attributions

- Tom Denton (Fields Institute/York University in Toronto)

Index

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: Introduction to Algebraic Structures (Denton)

Webpages: 49

All licenses found:

- Undeclared: 100% (49 pages)

By Page

- Introduction to Algebraic Structures (Denton) - *Undeclared*
  - Front Matter - *Undeclared*
    - TitlePage - *Undeclared*
    - InfoPage - *Undeclared*
    - Table of Contents - *Undeclared*
    - Licensing - *Undeclared*
    - Preface - *Undeclared*
  - 1: Symmetry - *Undeclared*
    - 1.1: Symmetry - *Undeclared*
    - 1.2: Counting Symmetries - *Undeclared*
    - 1.3: Symmetric Polynomials - *Undeclared*
    - 1.4: Abstraction. What is a group? - *Undeclared*
  - 2: Groups I - *Undeclared*
    - 2.1: Symmetry and Functions - *Undeclared*
    - 2.2: Definition of a Group - *Undeclared*
    - 2.3: Integers Modulo n - *Undeclared*
    - 2.4: Permutations - *Undeclared*
  - 3: Groups II - *Undeclared*
    - 3.1: Generating Sets - *Undeclared*
    - 3.2: Visualizing Groups- Cayley Graphs - *Undeclared*
    - 3.3: Subgroups - *Undeclared*
    - 3.4: Cosets and Lagrage's Theorem - *Undeclared*
  - 4: Groups III - *Undeclared*
    - 4.1: Homomorphisms - *Undeclared*
    - 4.2: Product Groups - *Undeclared*
    - 4.3: Image and Kernel - *Undeclared*
  - 5: Groups IV - *Undeclared*
    - 5.1: Quotient Groups - *Undeclared*
    - 5.2: Examples of Quotient Groups - *Undeclared*
    - 5.3: Isomorphism Theorem - *Undeclared*
    - 5.4: Classifying Finite Groups - *Undeclared*
  - 6: Group Actions - *Undeclared*
    - 6.1: Group Actions - *Undeclared*
    - 6.2: Orbits and Stabilizers - *Undeclared*
    - 6.3: Counting - *Undeclared*
  - 7: Rings I - *Undeclared*
    - 7.1: Juggling With Two Operations - *Undeclared*
    - 7.2: Ring Homomorphisms - *Undeclared*
  - 8: Rings II - *Undeclared*
    - 8.1: The Problem of Division - *Undeclared*
    - 8.2: Field of Fractions - *Undeclared*
    - 8.3: Euclidean Domains - *Undeclared*
  - 9: Vector Spaces - *Undeclared*
    - 9.1: A Return to Linear Algebra - *Undeclared*
    - 9.2: Linear Independence - *Undeclared*
  - Back Matter - *Undeclared*
    - Index - *Undeclared*
    - Glossary - *Undeclared*
    - Detailed Licensing - *Undeclared*