

Edited by Matt Clay & Dan Margalit

Office Hours with a Geometric Group Theorist

Office Hours with a Geometric Group Theorist

Edited by Matt Clay and Dan Margalit

PRINCETON UNIVERSITY PRESS PRINCETON AND OXFORD Copyright © 2017 by Princeton University Press

Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock, Oxfordshire, OX20 1TR

press.princeton.edu

Cover art: Illustration of door courtesy of Shutterstock. Other illustrations by Anne Karetnikov

Photo of Tara Brendle by Jonathan Tickner

All Rights Reserved

ISBN 978-0-691-15866-2 Library of Congress Control Number: 2016960820

British Library Cataloging-in-Publication Data is available

This book has been composed in Times

Printed on acid-free paper. ∞

press.princeton.edu

Typeset by Nova Techset Pvt Ltd, Bangalore, India

Printed in the United States of America

 $10\ 9\ 8\ 7\ 6\ 5\ 4\ 3\ 2\ 1$

To Julianna and Magdalena, Lily and Simon

Contents

Preface		xi	
Acknowledgments			
PART 1.	GROUPS AND SPACES	1	
1. Group	os Matt Clay and Dan Margalit	3	
1.1	Groups	5	
1.2	Infinite groups	9	
1.3	Homomorphisms and normal subgroups	13	
1.4	Group presentations	17	
2an	d Spaces Matt Clay and Dan Margalit	21	
2.1	Graphs	23	
2.2	Metric spaces	34	
2.3	Geometric group theory: groups and their spaces	40	
PART 2.	FREE GROUPS	43	
3. Groups Acting on Trees Dan Margalit		45	
3.1	The Farey tree	46	
3.2	Free actions on trees	50	
3.3	Non-free actions on trees	56	
4. Free	Groups and Folding Matt Clay	66	
4.1	Topological model for the free group	67	
4.2	Subgroups via graphs	70	
4.3	Applications of folding	73	
5. The F	ing-Pong Lemma Johanna Mangahas	85	
5.1	Statement, proof, and first examples using ping-pong	85	
5.2	Ping-pong with Möbius transformations	90	
5.3	Hyperbolic geometry	95	
5.4	Final remarks	103	

CONTENTS

6. Autor	norphisms of Free Groups Matt Clay	106			
6.1	Automorphisms of groups: first examples	106			
6.2	108				
6.3	110				
PART 3.	LARGE SCALE GEOMETRY	123			
7. Quas	i-isometries Dan Margalit and Anne Thomas	125			
7.1	Example: the integers				
7.2	Bi-Lipschitz equivalence of word metrics	127			
7.3	Quasi-isometric equivalence of Cayley graphs	130			
7.4	Quasi-isometries between groups and spaces	133			
7.5	Quasi-isometric rigidity				
8. Dehn	Functions Timothy Riley	146			
8.1	Jigsaw puzzles reimagined	147			
8.2	A complexity measure for the word problem	149			
8.3	Isoperimetry	156			
8.4	A large-scale geometric invariant	162			
8.5	The Dehn function landscape	163			
9. Нуре	rbolic Groups Moon Duchin	176			
9.1	Definition of hyperbolicity	178			
9.2	Examples and nonexamples	182			
9.3	Surface groups	186			
9.4	Geometric properties	193			
9.5	Hyperbolic groups have solvable word problem	197			
10. Ends	of Groups Nic Koban and John Meier	203			
10.1	An example	203			
10.2	The number of ends of a group	206			
10.3	Semidirect products	208			
10.4	Calculating the number of ends of the braid groups	213			
10.5	Moving beyond counting	215			
11. Asym	ptotic Dimension Greg Bell	219			
11.1	Dimension	219			
11.2	Motivating examples	220			
11.3	Large-scale geometry				
11.4	Topology and dimension				
11.5	Large-scale dimension				
11.6	Motivating examples revisited	231			
11.7	Three questions				
11.8	Other examples				

CONTENTS		ix		
12. Growt	h of Groups Eric Freden	237		
12.1	Growth series	238		
12.2	Cone types	244		
12.3	2.3 Formal languages and context-free grammars			
12.4	The DSV method	256		
PART 4.	EXAMPLES	267		
13. Coxet	er Groups Adam Piggott	269		
13.1	Groups generated by reflections	269		
13.2	Discrete groups generated by reflections	275		
13.3	Relations in finite groups generated by reflections	278		
13.4	Coxeter groups	281		
14. Right-	Angled Artin Groups Robert W. Bell and Matt Clay	291		
14.1	Right-angled Artin groups as subgroups	293		
14.2	Connections with other classes of groups	295		
14.3	Subgroups of right-angled Artin groups	298		
14.4	The word problem for right-angled Artin groups	301		
15. Lamp	lighter Groups Jennifer Taback	310		
15.1	Generators and relators	311		
15.2	Computing word length	315		
15.3	Dead end elements	318		
15.4	Geometry of the Cayley graph	321		
15.5	Generalizations	327		
16. Thom	pson's Group Sean Cleary	331		
16.1	Analytic definition and basic properties	332		
16.2	Combinatorial definition	336		
16.3	Presentations	340		
16.4	Algebraic structure	345		
16.5	Geometric properties	352		
17. Map	ping Class Groups Tara Brendle, Leah Childers, and			
	Dan Margalit	358		
17.1	A brief user's guide to surfaces	359		
17.2	Homeomorphisms of surfaces	363		
17.3	Mapping class groups	367		
17.4	Dehn twists in the mapping class group	370		
17.5	Generating the mapping class group by Dehn twists	373		
18. Braids	s Aaron Abrams	384		
18.1	Getting started	384		
18.2	Some group theory			

18.3	Some topology: configuration spaces	395
18.4	More topology: punctured disks	400
18.5	Connection: knot theory	404
18.6	Connection: robotics	407
18.7	Connection: hyperplane arrangements	409
18.8	A stylish and practical finale	411
Bibliogra	aphy	419

I	n	d	ex
•	•••	~	~ ~ ~

Х

437

Preface

This book grew out of an invited paper session, *An Invitation to Geometric Group Theory*, at the 2010 MAA MathFest in Pittsburgh. Geometric group theory is an exciting, relatively new field of mathematics that explores groups using tools and techniques from geometry and topology. A common theme throughout the field is:

Algebraic properties of a group are reflected in the geometry of spaces on which the group acts.

While the notion of a group forms a key component of an undergraduate education in mathematics, the above viewpoint is rarely taken. The typical undergraduate algebra curriculum focuses mainly on finite groups and students can graduate without being exposed to the vast and interesting world of infinite groups. The purpose of this book is to invite undergraduates, beginning graduates students, and curious researchers to explore this world by means of considering examples central to geometric group theory.

Each chapter—rather, office hour—of this book, driven by examples, covers one topic from geometric group theory. The authors were charged with the task of writing their answer to an undergraduate who had posed one of the following questions:

- Will you tell me about your research?
- Will you help me find a topic for a senior project?

As such, the office hours have an intentionally informal tone, each conveying the individual voice of the author. Ideally, the student reading this book should feel as if he or she is visiting a professor during office hours. When a student comes to your office curious about advanced topics in mathematics, you don't start the conversation by: "An amalgamated free product decomposition is." Rather, you draw a picture of $SL(2, \mathbb{Z})$ acting on the hyperbolic plane and explain how this action naturally decomposes the group. We hope that this style conveys the authors' enthusiasm for the subject and helps to inspire the next generation of researchers.

To facilitate the introduction to mathematical research, each office hour ends with suggestions for further reading and a selection of advanced exercises and research projects.

We do not think of this book as a replacement for John Meier's excellent text, *Groups, Graphs and Trees*, but instead we see it more as an undergraduate version of Pierre de la Harpe's book, *Topics in Geometric Group Theory*. In choosing topics for this book, we focused on the concepts that appear in several areas of geometric group theory—such as quasi-isometry—or on the classes of groups that prevail

throughout geometric group theory—such as Coxeter groups and right-angled Artin groups. We also tried to choose topics that we felt were accessible to an undergraduate who had taken a first course in abstract algebra, at least with the help of a patient adviser. Thus, we omitted some key geometric group theory concepts such as finiteness properties. In the rare office hours that deal with a specialized topic, such as asymptotic dimension, we felt that a student could find some interesting concept to latch onto, such as that of finding appropriate arrangements of bricks.

With the exception of the two introductory office hours, the office hours do not depend on one another, although connections within the material are highlighted. The topics covered in introductory office hours are:

- Groups
- Free groups and free abelian groups
- Group actions
- Cayley graphs
- Metric spaces

After reading the introductory office hours, an undergraduate should be able to browse the remaining office hours independently and focus on the topics that grab his or her interest.

We would be remiss not to list some of the other excellent general reference books on geometric group theory:

- A Course in Geometric Group Theory, Brian H. Bowditch
- Metric Spaces of Non-positive Curvature, Martin Bridson & André Haefliger
- Topics in Geometric Group Theory, Pierre de la Harpe
- Groups, Graphs and Trees, John Meier
- Classical Topology and Combinatorial Group Theory, John Stillwell

It has been a great pleasure working with the 15 other authors to compile some of our favorite gems of geometric group theory. We would like to thank them all for their hard work, persistence, patience, and enthusiasm. Most of all we would like to thank them for sharing their mathematical knowledge and their unique voices, and for believing in this project when we started with little more than an idea six years ago. The process has been invigorating and inspiring for us. We hope that you enjoy your journey through geometric group theory as much as we enjoyed laying down the path.

Matt Clay and Dan Margalit

Fayetteville and Atlanta, November 2016

Acknowledgments

Matt Clay and Dan Margalit would like to thank Mladen Bestvina and Benson Farb for first teaching us about geometric group theory and for being constant sources of support and inspiration. We would also like to thank Linda Chen, Yen Duong, Sudipta Kolay, and Caglar Uyanik for sending several comments and corrections. We are grateful to Ryan Dickmann for help with the typesetting of the book. We would also like to thank Justin Lanier for many insightful suggestions. We thank Bhisham Bherwani for his careful copyediting. Finally, this project would not have been possible without Vickie Kearn at Princeton University Press, whose unwavering support of our idea got us started and kept us going.

Tara Brendle, Leah Childers, and Dan Margalit are grateful to Joan Birman and Saul Schleimer for helpful conversations. They would also like to thank Mante Zelvyte for comments on an earlier draft of their office hour.

Moon Duchin would like to thank the participants in the 2012 Undergraduate Faculty Program at the Park City Mathematics Institute, especially Thomas Mattman and David Peifer, for their comments. She is also grateful to Sarah Bray and Brendan Foley for attempting several of the exercises.

Johanna Mangahas would like to express gratitude for comments and corrections offered by the participants of the 2012 PCMI Undergraduate Summer School taught by Jennifer Taback and the Undergraduate Faculty Program led by Moon Duchin, especially Allison Miller and Joe Ochiltree and David Peifer.

Timothy Riley thanks Margarita Amchislavska, Costandino Moraites, and Kristen Pueschel for careful readings of drafts of this chapter; Sasha Ushakov for conversations on the difficulty of describing "puzzle solutions"; and Robert Kleinberg and Jonathan Robbins for conversations on Project 4 of Office Hour 8.

Jennifer Taback thanks Daniel Studenmund, Peter Davids, Michael Ben-Zvi, Kevin Buckles, and all the students in Math 2602 and Math 3602 at Bowdoin College for helpful comments on her office hour.

Matt Clay, Sean Cleary, Johanna Mangahas, Dan Margalit, Timothy Riley and Jennifer Taback gratefully acknowledge support from NSF Grant Nos. DMS-1006898, DMS-1417820, DMS-1204592, DMS-1057874, DMS-1101651, and DMS-1105407, respectively.

The non-editing office hour authors thank Matt Clay and Dan Margalit for conceiving this book, inviting our participation, providing encouragement and vision along the way, and skillfully and patiently seeing the project to completion.

PART 1 Groups and Spaces

Groups

Matt Clay and Dan Margalit



Symmetry, as wide or as narrow as you define its meaning, is one idea by which man through the ages has tried to comprehend and create order, beauty and perfection.

Hermann Weyl

An important problem in mathematics is to determine the symmetries of objects. The objects we are interested in might be concrete things such as planar shapes, higher-dimensional solids, or the universe. Or they might be more abstract, such as groups, spaces of functions, or electric fields.

We'll begin in this office hour with a general discussion of groups; the focus will be on interpreting every group as a group of symmetries of some object. In the second office hour we will make precise what it means for a group to be a group of symmetries of an object. These ideas are at the very heart of geometric group theory, the study of groups, spaces, and the interactions between them.

Symmetry. By understanding the symmetries of an object we come closer to fully understanding the object itself. For instance, consider a cube in Euclidean 3–space:



A symmetry of the cube is a rigid transformation of the cube (for now, no reflections allowed). In other words, a symmetry of the cube is what you get if you pick the cube up, rotate it with your hands, and then put in back in the same spot where it started. One way to obtain a symmetry is to skewer the cube through the centers of two opposite faces and rotate by any multiple of $\pi/2$. We could also skewer the cube through the midpoints of two opposite edges, or through two opposite corners.

Here is a question:

What are all of the symmetries of the cube?

If we don't know the answer to this question, then we don't know the cube!

There is a fairly straightforward analysis which tells us that there are exactly 24 symmetries of the cube: there are eight places to send any given corner, and once a corner is fixed, there are three choices for how to turn the cube at that corner.

This is a start, but leaves a lot to be desired: What is a good way to list the 24 symmetries? If we do one symmetry followed by another symmetry, we get a third symmetry; where is this third symmetry on our list?

Here is a simple idea that completely illuminates the problem. Draw the four long diagonals of the cube:



We can check the following two facts:

- Any symmetry of the cube gives a permutation of the four long diagonals.
- Any permutation of the four long diagonals gives a unique symmetry of the cube.

The validity of the first statement is not too hard to see. The second one takes a little thought. First notice that if we take the permutation that swaps two long diagonals, then we obtain a symmetry of the cube, namely, a symmetry obtained by skewering the cube through the midpoints of the two edges that connect the endpoints of the two long diagonals in question. But an arbitrary permutation of the four long diagonals can be obtained by swapping two long diagonals at a time, so we are done.

It follows that the set of symmetries of the cube is in bijection with the set of permutations of its four long diagonals. There are 4! = 24 permutations of the four long diagonals, agreeing with our calculation that there are 24 symmetries of the cube. Moreover, since a permutation is just a bijective function from a set to itself, we know what happens when we do one permutation, then another—we get the composition of the two functions. We'll say more about permutations in a few minutes, when we talk about symmetric groups.

That was satisfying—we certainly understand the cube better now than we did at the start. In geometric group theory we aim to understand symmetries of much more complicated (and beautiful!) objects.

Here is the plan for this first office hour. We'll start by revisiting the basic examples of groups from any undergraduate course in abstract algebra. But we will look at these groups from the point of view that *every* group should be the collection of symmetries of some geometric object, just as the symmetric group on four letters is the collection of symmetries of a three-dimensional cube. To some extent, a lot of the material in this first office hour is a review of some of the important ideas from a first course in abstract algebra; you may want to skim through it if you are comfortable with those concepts.

In the next office hour, we will make sense of and prove the following theorem (Theorem 2.2 below):

Every group is naturally identified with the collection of symmetries of some geometric object, namely, the Cayley graph of the group.

The goal of this book is to convince you that this theorem—along with the many other theorems in this book—is the beginning of a beautiful and fruitful dictionary relating the algebraic and geometric structures arising in the wild and fascinating world of infinite groups.

1.1 GROUPS

Let's recall the definition of a group. Then we'll explain why you should think of a group in terms of symmetries.

First, a *multiplication* on a set G is a function

$$G \times G \to G.$$

The image of (g, h) is usually written gh. In other words, a multiplication is a way of combining two elements of G in order to get a third one.

Then, a *group* is a set *G* together with a multiplication that satisfies the following three properties:

- *Identity.* There is an element 1 of G such that 1g = g1 = g for all $g \in G$. The element 1 is called the *identity element* of G.
- *Inverses.* For each g ∈ G, there is an element h ∈ G such that gh = hg = 1. The element h is often denoted by g⁻¹ and is called the *inverse* of g.
- Associativity. The multiplication on G is associative, that is:

$$(fg)h = f(gh)$$

for all $f, g, h \in G$.

What does this abstract, formal definition have to do with symmetries? Let's first think about this in the case of the cube. We defined a symmetry of the cube as a rigid transformation of the cube. The way we multiply two symmetries g and h is:

"do h, then do g." This is just like composing functions: in the composition $g \circ h$ we apply h first.

We now see that the set of symmetries of the cube satisfies the definition of a group. The identity symmetry 1 is the rigid transformation of the cube that leaves it alone. Indeed, doing one rigid transformation g and then doing 1 is the same as doing g. The inverse of a symmetry g is the symmetry that undoes g. Associativity follows from the associativity of composition of functions: doing h, then g and f is the same as doing h and g, then f.

As we already said, we'd like to eventually convince you that every group that is, every set with a multiplication as above—is the set of symmetries of some object. First, let's look at a few examples of groups and explain how they each can be thought of as a set of symmetries.

The dihedral group. The dihedral group D_n is the set of symmetries (rigid transformations) of a regular *n*-gon in the plane. The multiplication is the same as in the example of the cube. Two symmetries of the regular pentagon are the reflections *s* and *t* shown here:



For general *n* we can take *s* to be the reflection through any line that passes through the center of the *n*-gon and a vertex of the *n*-gon, and *t* to be the reflection about the line that differs from the first line by a π/n rotation. The elements *s* and *t* generate the group D_n , by which we mean that every symmetry of D_n is obtained by multiplying, *s*, *t*, *s*⁻¹, and *t*⁻¹ (actually, in this example $s = s^{-1}$ and $t = t^{-1}$). For example, the clockwise rotation by *m* "clicks" is $(st)^m$.

Adam Piggott's Office Hour 13 on Coxeter groups discusses examples of groups that are all based on the idea of the dihedral groups.

The symmetric group. The symmetric group S_n is the set of permutations of the set $\{1, ..., n\}$, with multiplication given by composition of functions. This makes sense since a permutation of $\{1, ..., n\}$ is really just a bijective function $\{1, ..., n\} \rightarrow \{1, ..., n\}$. If we want to think of S_n as the set of symmetries of some object, we can think of it as the symmetries of a set of *n* points.

A *cycle* in S_n is a permutation that can be described as $i_i \mapsto i_2 \mapsto \cdots \mapsto i_k \mapsto i_1$ (and all other elements of $\{1, \ldots, n\}$ stay where they are). We write this cycle as $(i_1 i_2 \cdots i_{k-1} i_k)$. For instance, the cycle (1 2 4) in S_6 sends 1 to 2, 2 to 4, and 4 to 1, while sending 3, 5, and 6 to themselves.

Every element of S_n is a product of disjoint cycles (disjoint means that each element of $\{1, ..., n\}$ appears in at most one cycle). For instance, in S_6 the element

(124)(35)

is the product of the disjoint cycles $(1 \ 2 \ 4)$ and $(3 \ 5)$. This is the permutation that sends 1 to 2, 2 to 4, and 4 to 1, and also 3 to 5 and 5 to 3, and finally 6 to itself.

As we said, the multiplication is function composition. Say, for example, we want to multiply $(1\ 2\ 4)(3\ 5)$ by $(2\ 6)$:

$$(2\ 6)\cdot(1\ 2\ 4)(3\ 5).$$

Where does this product send 1? Well, the first permutation (on the right) sends 1 to 2, and the second permutation sends 2 to 6, so the product (= composition) sends 1 to 6. You can use the same procedure to determine where this product sends 2, 3, 4, 5, and 6. You will find that $(2 \ 6) \cdot (1 \ 2 \ 4)(3 \ 5) = (1 \ 6 \ 2 \ 4)(3 \ 5)$.

A *transposition* is a cycle of length 2, for instance, (3 5). It is an important fact that S_n is generated by transpositions of the form $(i \ i + 1)$, where $1 \le i \le n - 1$ (we actually used this fact implicitly in our discussion of the cube at the beginning). Let us check that we can write $(1 \ 2 \ 4)(3 \ 5)$ as a product of such elements in S_6 . We can draw a diagram of this permutation as follows:



Each crossing in the diagram corresponds to an element of S_6 of the form $(i \ i + 1)$. Reading top to bottom (and writing right to left), we find:

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

The point is that multiplication in S_n can be realized by stacking diagrams. And the diagram for $(1\ 2\ 4)(3\ 5)$ can be obtained by stacking the diagrams for the five permutations on the right-hand side of the equation. Since we can always make a diagram where the crossings occur at different heights, this argument can be used to show that every element of S_n is equal to a product of transpositions $(i\ i+1)$.

Exercise 1. Use the idea of stacking diagrams to prove that S_n is generated by $\{(i \ i+1) \mid 1 \le i \le n-1\}$.

The picture we drew of the permutation $(1\ 2\ 4)(3\ 5)$ can be called a braid diagram. See Aaron Abrams' Office Hour 18 on braid groups for more on this idea.

Exercise 2. We showed that S_4 is the collection of symmetries of a threedimensional cube. Can any of the other symmetric groups be thought of as the symmetries of higher-dimensional cubes? Or other shapes?

The integers modulo *n***.** Let *n* be an integer greater than 1. The *integers modulo n* is the set

$$\mathbb{Z}/n\mathbb{Z} = \{0, 1, \dots, n-1\}$$

with the multiplication

$$(a,b) \mapsto \begin{cases} a+b & a+b \le n-1\\ a+b-n & a+b \ge n. \end{cases}$$

The identity for $\mathbb{Z}/n\mathbb{Z}$ is 0. The inverse of 0 is 0, and the inverse of any other *m* is n - m. Associativity is a little trickier: to check that (a + b) + c = a + (b + c), there are a few cases, depending on which sums are greater than *n*. We'll leave this as an exercise.

Exercise 3. Show that if *m* is any element of $\{0, ..., n-1\}$ with gcd(m, n) = 1, then $\{m\}$ is a generating set for $\mathbb{Z}/n\mathbb{Z}$.

You are most familiar with $\mathbb{Z}/n\mathbb{Z}$ in three cases:

- When *n* = 12, the multiplication we defined is clock arithmetic. For example, 3:00 plus 5:00 is 8:00, and 9:00 plus 4:00 is 1:00 (although perhaps we should think of 12:00 as 0:00 instead).
- When *n* = 10, the multiplication we defined is an operation you use when balancing your checkbook. By just keeping track of the last digit, you have a quick first check for whether the sum of many big numbers is equal to what your bank says it is.
- When n = 2, the multiplication is light switch arithmetic. Identify 1 with flipping the switch and 0 with do nothing. Then flip plus flip is the same as doing nothing, flip plus do nothing is the same as flip, etc.

The idea of light switch arithmetic inspires a whole class of groups called lamplighter groups; see Jen Taback's Office Hour 15.

For what object is $\mathbb{Z}/n\mathbb{Z}$ the group of symmetries? Using the idea of the clock we might try to use a regular *n*-gon. We can certainly see $\mathbb{Z}/n\mathbb{Z}$ as a set of symmetries of the *n*-gon, namely, the rotational symmetries. Thus, $\mathbb{Z}/n\mathbb{Z}$ can be regarded as a *subgroup* of the full symmetry group, D_n . This means that we have realized $\mathbb{Z}/n\mathbb{Z}$ as a subset of D_n in such a way that the multiplication in $\mathbb{Z}/n\mathbb{Z}$ agrees with the multiplication in D_n . We can see that $\mathbb{Z}/n\mathbb{Z}$ is a proper subgroup of D_n (i.e., it is not the whole thing), since D_n also has reflectional symmetries.

That's pretty good. But we can modify the *n*-gon so that $\mathbb{Z}/n\mathbb{Z}$ is exactly the group of symmetries. One way to do this is to draw a little arrow in the middle of each edge:



We think of the arrow as defining a preferred direction on each edge. Then we define a symmetry of this modified *n*-gon to be a rigid transformation that preserves the directions of the arrows. This is still a perfectly fine notion of symmetry! The group $\mathbb{Z}/n\mathbb{Z}$ is the full group of symmetries of the *n*-gon in this modified sense.

1.2 INFINITE GROUPS

Many undergraduate courses in abstract algebra give short shrift to infinite groups. As we will see in this book, this is where all of the fun is! Let's start with the simplest infinite group.

The integers. You are probably familiar with the set of *integers*, or whole numbers:

 $\mathbb{Z} = \{\ldots, -3, -2, -1, 0, 1, 2, 3, \ldots\}.$

The multiplication $\mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$ is the usual addition of integers:

$$(m, n) \mapsto m + n.$$

The group \mathbb{Z} is generated by {1}. But it is also generated by {2, 3}, for example.

Of what object is \mathbb{Z} the set of symmetries? A first guess is that \mathbb{Z} is the symmetries of the number line *L*, by which we mean the real line \mathbb{R} with a distinguished point at each integer:



Here, by a symmetry of L, we just mean a rigid transformation of L that preserves the set of distinguished points. Examples of symmetries of L are: translate L to the left or right by n units, where n is a positive integer; reflect it about an integer point; and reflect it about a half-integer point.

The group \mathbb{Z} gives a set of symmetries of *L* as follows: the integer *n* is the symmetry that translates *L* by *n* units to the right if $n \ge 0$ and translates *L* by -n units to the left if n < 0. If we think of *L* as an ∞ -gon, this is analogous to what we did for $\mathbb{Z}/n\mathbb{Z}$. As in the case of $\mathbb{Z}/n\mathbb{Z}$, this identifies \mathbb{Z} as a proper subgroup of the set of symmetries of *L*, since we don't get any of the reflections.

We can use the same trick as in the $\mathbb{Z}/n\mathbb{Z}$ case to modify *L* in such a way that \mathbb{Z} is the full symmetry group: on each edge we draw a little arrow that points to the right:

The integers squared. Our next group is

$$\mathbb{Z}^2 = \{(m, n) : m, n \in \mathbb{Z}\}$$

with the multiplication

$$((m, n), (m', n')) \mapsto (m + m', n + n').$$

If we think of \mathbb{Z}^2 as the set of vectors in \mathbb{R}^2 with integer coordinates, then this is simply vector addition. One generating set for \mathbb{Z}^2 is $\{(1, 0), (0, 1)\}$.

We can use the same idea as above to build an object whose symmetries are exactly described by \mathbb{Z}^2 :



In order for this to work, we need to again refine what we mean by symmetry: this time we insist that colors are preserved as well as the red points and the arrows. The symmetry corresponding to $(m, n) \in \mathbb{Z}^2$ is the translation by *m* units to the right and *n* units up. (If *m* or *n* are negative, then translate left and down accordingly.)

Multiplicative groups of numbers. We would be remiss to not mention two important groups, \mathbb{R}^* and \mathbb{C}^* . These are the groups of nonzero real numbers and nonzero complex numbers. The group multiplication in each group is the usual multiplication in \mathbb{R} and \mathbb{C} . In both cases the identity is 1. An important distinction between these groups and our previous examples is that \mathbb{R}^* and \mathbb{C}^* are not finitely generated, that is, there does not exist a finite generating set.

If you've already seen these groups, you might want to learn about the quaternions, a generalization of the complex numbers which is very important in physics. Instead of just i, the quaternions have i, j, and k! There is a corresponding multiplicative group, consisting again of the nonzero elements.

Matrix groups. There is another set of examples of groups that you learned about in linear algebra (although in such a class you might not have used the word "group"). Denote by $GL(n, \mathbb{R})$ the set of real $n \times n$ matrices with nonzero determinant. Why is this a group? Well, given two matrices, we already know how to multiply them. Note that this multiplication is closed (that is, the product of two matrices with nonzero determinant is another matrix with nonzero determinant). Indeed, this follows immediately from the following basic fact you learned in linear algebra:

$$\det(MN) = \det(M) \det(N).$$

Two other things you learned in linear algebra is that every square matrix with nonzero determinant is invertible and that multiplication is associative. The identity matrix has determinant 1, which is not zero, so we have checked that $GL(n, \mathbb{R})$ is indeed a group. It is called the *general linear group*. There are really lots of general linear groups, since *n* can be any nonnegative integer and \mathbb{R} can be replaced with any ring (see below), but when the context is clear we still call it *the* general linear group.

There are many variations. The group $SL(n, \mathbb{R})$ is the subgroup of $GL(n, \mathbb{R})$ consisting of matrices with determinant 1 (this is the *special linear group*). The group $GL(n, \mathbb{Z})$ is the subgroup of $GL(n, \mathbb{R})$ consisting of all invertible integral matrices (note: these are exactly the integral matrices with determinant ± 1). Then $SL(n, \mathbb{Z})$ is the subgroup of $GL(n, \mathbb{Z})$ consisting of integral matrices with determinant 1. There is also $GL(n, \mathbb{C})$, the group of $GL(n, \mathbb{C})$ consisting of $n \times n$ matrices with nonzero determinant, and $SL(n, \mathbb{C})$, the subgroup of $GL(n, \mathbb{C})$ consisting of $R(n, \mathbb{C})$, and there are many other examples of interesting subgroups of $GL(n, \mathbb{C})$ (subgroups of $GL(n, \mathbb{C})$ are called *linear groups*).

We will study $SL(2, \mathbb{Z})$ in detail in Office Hour 3, giving a finite presentation for this group as well as many of its relatives.

Exercise 4. Construct a matrix group whose entries lie in a finite field.

The free group of rank 2. Our next example is the most complicated (and interesting!) one yet. And it is most probably not a group that you saw in your undergraduate algebra course.

First we define a *word* in the letters *a* and *b* to be an arbitrary finite string made up of the symbols *a*, *b*, a^{-1} , and b^{-1} . Some examples are:

 $abababa, aaaaaaa, aba^{-1}b^{-1}$, and aa^{-1} .

We allow the empty word as well; this is the string containing no letters. We can multiply two words by concatenating them:

$$(ab, b^{-1}a) \mapsto abb^{-1}a.$$

Next, we define a *reduced word* to be a word with the property that we never see an a followed by an a^{-1} (or vice versa) or a b followed by a b^{-1} (or vice versa). If we have a word that is not reduced, we can make it shorter by deleting an offending pair of symbols:

$$aa^{-1}abb^{-1}a \rightsquigarrow abb^{-1}a.$$

Performing this procedure inductively will eventually lead to a reduced word. You should check that this reduction procedure always results in the same reduced word. In the above example, there are three places where we could start reducing. In this case, though, no matter what we do we will end up with the reduced word *aa*.

Exercise 5. Show that every word in a, b, a^{-1} , and b^{-1} can be reduced to a unique reduced word.

We are finally ready to define the *free group of rank 2*. It is the set

 $F_2 = \{$ reduced words in a and $b\}$

with the multiplication:

concatenate, then reduce.

Let's check that this is a group. The identity is the empty word. The inverse of a word is obtained by reversing the word, and then replacing each *a* by a^{-1} , a^{-1} by a, b by b^{-1} , and b^{-1} by *b*. For example, the inverse of

$$aba^{-1}b$$

is

$$b^{-1}ab^{-1}a^{-1}$$

Indeed, if we concatenate these two (in either order!) and reduce, we obtain the empty word. Associativity follows immediately from the fact that the reduction procedure always results in the same reduced word (but this fact needs to be proved!). By the very definition of F_2 , we see that this group is generated by $\{a, b\}$. **Exercise 6.** Check the details of the construction of the free group of rank 2, in particular the assertion that the multiplication is associative. *Hint: Show by induction on the length of (unreduced) words that if there are two choices for a reduction, then either choice leads to the same reduced word. There are two cases, according to whether the two reductions are adjacent or not.*

Free groups are important examples of groups in geometric group theory. Office Hours 3, 4, 5, and 6 explore these groups in more detail.

How can we construct an object whose symmetry group is F_2 ? This is not so easy! In Office Hour 2 we will explain a way to do this for an arbitrary group.

Other free groups. Given a set *S*, we can define a free group F(S) as follows. First, we artificially create inverses for each element of *S*. For $s \in S$, we can write this inverse as s^{-1} (let's not allow *s* and s^{-1} to both be elements of *S*). Then the elements of F(S) are reduced words in the elements of *S*, and the multiplication is defined in the same way as for F_2 . As such, when *S* is a finite set with *n* elements, we can denote F(S) by F_n (different choices of *S* with |S| = n technically give different groups, but they are all isomorphic in the sense defined below). Convince yourself that F(S) still is a well-defined group when *S* is infinite, or even uncountable. In general, the cardinality of *S* is the *rank* of the corresponding free group.

As we will explain in more detail in the next section, free groups are important because every group is the quotient of a free group.

And many more. One aspect of geometric group theory is to explore the zoo of infinite groups that are out there. Some examples of interesting classes of groups in this book are the Coxeter groups already mentioned (Adam Piggott's Office Hour 13), the right-angled Artin groups (Bob Bell and Matt Clay's Office Hour 14), the lamplighter groups (Jen Taback's Office Hour 15), Thompson's groups (Sean Cleary's Office Hour 16), the mapping class groups (Tara Brendle and Leah Childers' Office Hour 17), and the braid groups (Aaron Abrams' Office Hour 18).

1.3 HOMOMORPHISMS AND NORMAL SUBGROUPS

We'll now take a few minutes to give a quick overview of the most basic and important ideas from a first course in group theory. If you already know all about the first isomorphism theorem and about group presentations, then you might want to skip ahead to the next office hour.

Homomorphisms. We have talked about groups and we now have some examples. Next we would like to discuss how groups relate to each other. A *homomorphism* from a group G to a group H is a function

$$f: G \to H$$

so that

$$f(ab) = f(a)f(b)$$

for all $a, b \in G$. We can say this succinctly as: f is a function that preserves the multiplication in G. A homomorphism is an *isomorphism* if it is bijective, and two groups are *isomorphic* if there is an isomorphism from one to the other. Convince yourself that isomorphic groups are essentially the same group.

To illustrate the idea of an isomorphism, let's look at a pair of groups. Consider the group \mathbb{Z} on one hand, and the free group F_1 on one generator a on the other hand. Elements of F_1 are finite, freely reduced strings in a and a^{-1} , and the multiplication is given as usual by concatenation and free reduction. What is an isomorphism $\mathbb{Z} \to F_1$? (There is more than one!)

Besides isomorphisms, we might also be interested in injective homomorphisms and surjective homomorphisms. A homomorphism is injective if and only if the preimage of the identity in H is the identity in G. (Why?)

Note that any homomorphism $f: G \to H$ can be turned into a surjective homomorphism by replacing H with f(G). For that reason, the surjectivity of a homomorphism is not nearly as crucial as the injectivity. As such we will focus on injective and non-injective homomorphisms.

Injective homomorphisms. If we have an injective homomorphism $f: G \rightarrow H$, then we can think of f as realizing G as a subgroup of H. Here are a few examples:

- 1. There is an injective homomorphism $\mathbb{Z}/n\mathbb{Z} \to D_n$ where *m* maps to rotation by $2\pi m/n$.
- 2. There is an injective homomorphism $\mathbb{Z}/2\mathbb{Z} \to D_n$ where the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$ maps to any reflection.
- 3. There is an injective homomorphism $\mathbb{Z}/2\mathbb{Z} \to S_n$ where the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$ maps to any product of disjoint transpositions.
- 4. For any integer $m \neq 0$, there is an injective homomorphism $\mathbb{Z} \to \mathbb{Z}$ given by multiplication by *m*.
- 5. For any choice of $(a, b) \neq (0, 0)$, there is an injective homomorphism $\mathbb{Z} \to \mathbb{Z}^2$, where *n* maps to (na, nb).
- For any choice of nontrivial element w ∈ F₂, there is an injective homomorphism Z → F₂ given by n → wⁿ.

Exercise 7. How many other injective homomorphisms can you find between the groups we have already discussed?

Non-injective homomorphisms. Non-injective homomorphisms can be just as interesting and just as useful. Since we have multiple elements of G mapping to the same element of H, we are definitely losing some information. But the fact that a homomorphism preserves the multiplication of G means we are still remembering some important data. Often we can set up a non-injective homomorphism to

remember exactly the data we care about and throw away everything else. Here are some examples:

- The most basic example is the homomorphism Z → Z/2Z where 1 maps to 1. A light switch is a great example of this: it does not remember how many times it has been flipped, just whether it has been flipped an even or an odd number of times.
- 2. There is a homomorphism $GL(n, \mathbb{R}) \to \mathbb{R}^*$ given by the determinant. So if you want to know the determinant of a product of matrices, you do not need to remember the actual matrices, just their determinants.
- 3. There is a homomorphism ℝ* → ℤ/2ℤ where positive numbers map to 0 and negative numbers map to 1. The fact that this is a homomorphism can be rephrased as: in order to tell if a product of nonzero numbers is positive or negative, you do not need to remember what the numbers are, just whether they are positive or negative.
- 4. The composition $GL(n, \mathbb{R}) \to \mathbb{R}^* \to \mathbb{Z}/2\mathbb{Z}$ is a homomorphism that remembers whether an element of $GL(n, \mathbb{R})$ has positive or negative determinant.
- 5. There is a homomorphism $S_n \to \mathbb{Z}/2\mathbb{Z}$ that records whether the diagram for a permutation (as above) has an even or an odd number of crossings. In other words, this homomorphism records the parity of the number of transpositions needed to write a given element of S_n .
- 6. There is a homomorphism $D_n \to \mathbb{Z}/2\mathbb{Z}$ where rotations map to 0 and reflections map to 1. This homomorphism remembers whether the polygon has been flipped over or not.
- 7. There is a homomorphism $F_2 \to \mathbb{Z}^2$ where the generators *a* and *b* for F_2 map to (1, 0) and (0, 1), respectively. This homomorphism remembers the sum of the exponents on *a* and *b* in any element of F_2 .
- 8. Any linear map $\mathbb{R}^n \to \mathbb{R}^m$ is a homomorphism. For example, the map $\mathbb{R}^2 \to \mathbb{R}$ that projects to the *x*-axis is a homomorphism that remembers just the first coordinate. This is a fancy way to state the obvious fact that in order to know the first coordinate of a sum of vectors, you need to know just the first coordinate of each vector.

Exercise 8. How many other non-injective homomorphisms can you find between the groups we have already discussed?

Normal subgroups. Next we will present a definition that will seem unrelated to homomorphisms; and then we will make a deep connection. A *normal subgroup* of a group G is a group N where N is a subset of G and where gng^{-1} lies in N for all $n \in N$ and all $g \in G$.

Why are normal subgroups important? Well, one way to make a normal subgroup of *G* is to find a group homomorphism *f* from *G* to another group *H*. The *kernel* of *f* is the set of $g \in G$ so that f(g) is the identity (we use the same terminology in linear algebra). We have the following basic fact: the kernel of a homomorphism $f: G \to H$ is a normal subgroup of *G*.

Here is why: if n is in the kernel of f, then

$$f(gng^{-1}) = f(g)f(n)f(g^{-1}) = f(g) 1 f(g^{-1})$$
$$= f(g)f(g^{-1}) = f(gg^{-1}) = f(1) = 1,$$

and so gng^{-1} is also in the kernel of f. Thus the kernel of f is normal.

The above examples of non-injective homomorphisms have kernels as follows:

- 1. The kernel of our $\mathbb{Z} \to \mathbb{Z}/2\mathbb{Z}$ is $2\mathbb{Z}$.
- 2. The kernel of our $GL(n, \mathbb{R}) \to \mathbb{R}^*$ is $SL(n, \mathbb{R})$.
- 3. The kernel of our $\mathbb{R}^* \to \mathbb{Z}/2\mathbb{Z}$ is \mathbb{R}_+ .
- 4. The kernel of our $GL(n, \mathbb{R}) \to \mathbb{R}^* \to \mathbb{Z}/2\mathbb{Z}$ is the subgroup $GL^+(n, \mathbb{R})$ consisting of matrices with positive determinant.
- 5. The kernel of our $S_n \to \mathbb{Z}/2\mathbb{Z}$ is the *alternating group* A_n ; this consists of the elements that can be written as the product of an even number of transpositions.
- 6. The kernel of our $D_n \to \mathbb{Z}/2\mathbb{Z}$ is the subgroup consisting of rotations; this subgroup is isomorphic to $\mathbb{Z}/n\mathbb{Z}$.
- 7. The kernel of our $F_2 \to \mathbb{Z}^2$ is the set of elements whose *a*-exponents and *b*-exponents both add to zero, e.g., $a^7 b^{-5} a^{-10} b^5 a^3$. This group is also known as the commutator subgroup of F_2 and is denoted by $[F_2, F_2]$.
- The kernel of the projection R² → R is the same as the kernel from linear algebra; in this case it is isomorphic to R.

So non-injective homomorphisms give us nontrivial normal subgroups. Can we go the other way? That is, if N is a normal subgroup of G, is there a homomorphism $G \rightarrow H$ so that N is the kernel? The answer is yes! Let's explain this. The first step is to define the *quotient group* G/N. Informally, this is the group obtained from G by declaring every element of N to be trivial.

Let us make this more precise. Declare two elements g_1, g_2 of G to be equivalent if $g_1g_2^{-1} \in N$, and write [g] for the equivalence class of g. Then the elements of G/N are the equivalence classes of the elements of G (notice that the equivalence class of the identity is precisely N). Now we need to say how to multiply two equivalence classes. We declare that the product of the equivalence class $[g_1]$ with the equivalence class $[g_2]$ is the equivalence class $[g_1g_2]$. This is precisely the place where the definition of a normal subgroup comes from—it is exactly what is needed so that this multiplication is well defined.

Exercise 9. Check that the multiplication in G/N is well defined if and only if N is normal.

There is an obvious homomorphism $G \to G/N$, where $g \in G$ maps to $[g] \in G/N$. It is not hard to see that the kernel of this homomorphism is precisely N. So we have completed the loop: normal subgroups and surjective homomorphisms are the same thing!

The first isomorphism theorem. The so-called first isomorphism theorem succinctly summarizes the relationship between homomorphisms and normal subgroups that we have just described.

THEOREM 1.1 (First isomorphism theorem). If $G \rightarrow H$ is a surjective homomorphism with kernel K, then H is isomorphic to G/K.

Again using our above examples, we have the following isomorphisms:

- 1. $(\mathbb{Z})/(2\mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}$ (this explains the notation $\mathbb{Z}/2\mathbb{Z}$)
- 2. $\operatorname{GL}(n, \mathbb{R}) / \operatorname{SL}(n, \mathbb{R}) \cong \mathbb{R}^{\star}$
- 3. $\mathbb{R}^*/\mathbb{R}_+ \cong \mathbb{Z}/2\mathbb{Z}$
- 4. $\operatorname{GL}(n, \mathbb{R}) / \operatorname{GL}^+(n, \mathbb{R}) \cong \mathbb{Z}/2\mathbb{Z}$
- 5. $S_n/A_n \cong \mathbb{Z}/2\mathbb{Z}$
- 6. $D_n/(\mathbb{Z}/n\mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}$
- 7. $F_2/[F_2, F_2] \cong \mathbb{Z}^2$
- 8. $\mathbb{R}^2/\mathbb{R}\cong\mathbb{R}$

Exercise 10. Use the first isomorphism theorem to come up with your own isomorphisms.

1.4 GROUP PRESENTATIONS

You may not have covered group presentations in your first course in abstract algebra, but they are closely tied to the other topics in this section and will loom large in this book.

One naive way to completely describe a group is to list all of the elements and write down the multiplication table (that is, explicitly give the product of any two elements). For a very small finite group, this is a reasonable way of understanding the group. But once you start considering larger—or even infinite—groups this approach becomes unwieldy.

We already have a more economical way of listing the elements of a group: we can list a set of generators. Often the infinite groups we care about have finite sets of generators. This is good, but it certainly doesn't tell us enough about a group: the free group F_2 and the free abelian group \mathbb{Z}^2 both have generating sets with two elements and we know that these are different groups since one is abelian and the other is not. We need a way of describing the multiplication table. But obviously we don't want to just list all of the entries.

The key observation is that some entries in the multiplication table imply other entries. In other words, equations between group elements imply other equations. For instance, if gh = k and h = pq, then (substituting pq for h in the first equation) we know that gpq = k. So what we can hope for is that we don't need to write down the entire multiplication table, but rather a small set of equalities that imply all the others. So now we have exactly hit upon the idea of what a group presentation is: it is firstly a list of generators and secondly a list of defining equalities so that all equalities follow from the listed ones.

Here is a first example, the group $\mathbb{Z}/n\mathbb{Z}$. Let's denote the elements of $\mathbb{Z}/n\mathbb{Z}$ by a^0, \ldots, a^{n-1} instead of $0, \ldots, n-1$. Naturally, the multiplication is $a^j a^k = a^{j+k}$, where the exponents are all taken modulo *n*. With this notation, $\mathbb{Z}/n\mathbb{Z}$ has the

presentation

$$\mathbb{Z}/n\mathbb{Z}\cong\langle a\mid a^n=1\rangle.$$

So the generating set is *a* and the only defining equality is $a^n = 1$. Every entry in the multiplication table for $\mathbb{Z}/n\mathbb{Z}$ follows from this one equality. For instance, one entry in the multiplication table is $a^{n-3}a^4 = a$. But we can easily derive this using the defining equality $a^n = 1$ as follows:

$$a^{n-3}a^4 = a^{n+1} = a^n a^1 = 1 \cdot a = a.$$

A second important example of a group presentation is the following:

$$\mathbb{Z}^2 = \langle x, y \mid xy = yx \rangle.$$

Can you see why this is a presentation of \mathbb{Z}^2 ? We'll come back to this in a little bit.

Exercise 11. Convince yourself that \mathbb{Z}^2 indeed has the purported presentation.

Hopefully these examples give you an idea of what a group presentation is. The way we make things formal—and the way we really think about group presentations—is through the use of free groups. At first our definition of a presentation won't look exactly like the examples we just gave, but we'll eventually explain why it is really the same thing.

A group presentation is a pair (S, R), where S is a set and R is a set of words in S. If we freely reduce the elements of R, then we can regard them as elements of the free group F(S), and indeed it will be important to take this point of view. The presentation is finite when both sets S and R are finite.

We say that a group G has the presentation $\langle S | R \rangle$ if G is isomorphic to the quotient of F(S) by the normal closure of R, the smallest normal subgroup of F(S) containing R (equivalently, the subgroup of F(S) generated by all conjugates of all elements of R). We write:

$$G \cong \langle S \mid R \rangle.$$

Here is how we think of this: the group *G* consists of reduced words in $S \cup S^{-1}$ but with the added caveat that words in *R* are the same as the empty word (this is where the normal closure comes in: if some element of *R* is supposed to represent the identity, then any conjugate of that element in *F*(*S*) should also represent the identity, and any products of conjugates of such elements ...). The elements of *S* are called *generators* for *G* and the elements of the normal closure of *R* are called *relators* for *G*, and the elements of *R* are called defining relators. If we have a presentation (as in our above discussion) with a relation like ab = ba or aba = bab, then we turn this into a relator by moving everything to one side: $aba^{-1}b^{-1}$ or $abab^{-1}a^{-1}a^{-1}$.

How does this formal definition match up with what we said before? Let's look more closely at \mathbb{Z}^2 . This group has the presentation

$$\mathbb{Z}^2 \cong \langle a, b \mid aba^{-1}b^{-1} \rangle.$$

(Notice that this is different from the presentation we gave a few minutes ago, but bear with us!) What are the elements of the group $\langle a, b | aba^{-1}b^{-1} \rangle$? Well, we

start with the free group on *a* and *b*, which consists of all freely reduced words in *a* and *b*. Then we need to take the quotient by the normal closure of $aba^{-1}b^{-1}$. What is the quotient as a set? Well, in the quotient we can change a word *w* into an equivalent word by inserting $aba^{-1}b^{-1}$ anywhere in *w*, or removing it from anywhere in *w*.

Why is that the same as our first presentation of \mathbb{Z}^2 as $\langle x, y | xy = yx \rangle$? Well, we claim that two elements of the free group on $\{a, b\}$ are equivalent in the quotient (by the normal quotient by $aba^{-1}b^{-1}$) if and only if they differ by a sequence of swaps, where we swap ab for ba or vice versa. Let's convince ourselves of this. Start with a word wbaw', where w and w' are both words in a and b. We said we can insert $aba^{-1}b^{-1}$ anywhere we want; so in the quotient,

$$wbaw' = w(aba^{-1}b^{-1})baw' = wab(a^{-1}b^{-1}ba)w' = wabw'.$$

By applying swaps like this we see that every element of the quotient is equivalent to $a^m b^n$ for some *m* and *n*. But this gives us an obvious isomorphism between our quotient and \mathbb{Z}^2 ; the isomorphism is given by $a^m b^n \leftrightarrow (m, n)$.

We can also see from the above discussion how our new presentation of \mathbb{Z}^2 (with generators *a* and *b*) matches up with our old one (with generators *x* and *y*). Both presentations are really saying the same thing. So which one you use just depends on your preference. When we have an equality between generators like xy = yx, then the equality is called a *relation*. And when we have an equality like $aba^{-1}b^{-1} = 1$ we say that $aba^{-1}b^{-1}$ is a *relator*, as in our formal definition. We can always change a relation to a relator by moving all of the generators to one side (as in our \mathbb{Z}^2 example), and the net effect is the same. In this book we will tend to use relators more than relations, but both will come up.

Every group can be given by a group presentation. In particular:

Every group is a quotient of a free group.

This is one reason why free groups are so important. In many ways, geometric group theory grew out of the study of combinatorial group theory, a subject largely concerned with studying groups via their presentations.

Exercise 12. Show that every group has a presentation. *Hint: Your presentation does not have to be efficient.*

Exercise 13. Find (nice) group presentations for F_n , D_n , and S_n .

Armed with the notion of a group presentation, you can start to list lots of different groups. Which presentations represent isomorphic groups? It turns out that this is an unsolvable problem in general; there is no algorithm that will take two presentations and decide in finite time if they present the same group. It is good to keep this in mind: group presentations are very helpful, but they do have limitations.

That's it for the first office hour—we introduced a number of examples of groups that will play key roles in the rest of the book. We want to understand these groups from as many points of view as possible. At first glance you might wonder what there is to learn about a free abelian group or a free group. The answer: more than you might imagine! Hopefully we have convinced you that the concept of a group is deeply intertwined with the idea of symmetries: groups lead to symmetries and symmetries lead to groups. In order to study infinite groups, we will want to take advantage of this and think of groups geometrically. A key step on this path is to show that every group really is the group of symmetries of some geometric object. That's the goal of the next office hour—see you there!

... and Spaces

Matt Clay and Dan Margalit



As we said at the outset of Office Hour 1, our goal is to exhibit every group as a collection of symmetries of some geometric object. We have given some examples of finite and infinite groups and in some cases we have given an appropriate geometric object with that group as its symmetries.

Our first task in this office hour is to say precisely what it means for a group to be a group of symmetries of a geometric object (answer: group actions). We will then define, for any group G, the Cayley graph of G and show that the group of symmetries of this graph is precisely the group G.

We will then introduce metric spaces, which formalize the notion of a geometric object. We'll see lots of metric spaces that groups can act on. Furthermore, we will see that groups themselves *are* metric spaces; in other words, groups themselves can be thought of as geometric objects. At the end we will use these ideas to frame the motivating questions of geometric group theory.

Group actions. In Office Hour 1 we got the idea that groups can *do* things to geometric objects (or sets) in the sense that elements of the group give symmetries of the object. For instance, each element of the dihedral group D_n does something to the *n*-gon, each element of the symmetric group S_n does something to the set $\{1, \ldots, n\}$, and each element of \mathbb{Z} does something to the real line. In other words, we can say that these groups *act* on those objects. This is precisely the idea of a group action.

Formally, an *action* of a group G on a set X is a function

 $G\times X\to X$

where the image of (g, x) is written $g \cdot x$ and where

- $1 \cdot x = x$ for all $x \in X$ and
- $g \cdot (h \cdot x) = (gh) \cdot x$ for all $g, h \in G$ and $x \in X$.

If A is a subset of X then $g \cdot A$ is the set $\{g \cdot a \mid a \in A\}$. When there is no confusion we sometimes write gx and gA instead of $g \cdot x$ and $g \cdot A$.

Notice that for any group G and any set X we always have the trivial action, defined by $g \cdot x = x$ for all $g \in G$ and $x \in X$.

Let S_X denote the symmetric group of X, that is, the group of permutations of X (this makes sense even if X is infinite). So not only does each element of G give rise to a permutation of X, but it does so in such a way as to respect the multiplications in the two groups.

The group S_X is the group of symmetries of X, thought of as a set (hence the name). An action of G on X is the same thing as a homomorphism $G \to S_X$. So an action of G on X is the formal way to realize G as a group of symmetries of the set X.

Exercise 1. Check that actions of *G* on *X* are the same as homomorphisms $G \rightarrow S_X$.

Here are some examples of group actions on sets:

- 1. S_n acts on $\{1, ..., n\}$
- 2. D_n acts on the set of vertices (or edges, or diagonals ...) of an *n*-gon
- 3. GL(n, \mathbb{R}) acts on the set of vectors (or planes, or bases ...) in \mathbb{R}^n
- 4. $\mathbb{Z}/2\mathbb{Z}$ acts on \mathbb{Z} via multiplication by ± 1
- 5. \mathbb{Z} acts on \mathbb{R}^n via $m \cdot v = A^m v$, where A is some fixed $n \times n$ matrix
- 6. \mathbb{Z} acts on the unit circle in \mathbb{C} via the formula $m \cdot z = e^{im\theta} z$, where θ is any real number (so *m* rotates the circle by $m\theta$)

Exercise 2. List some other group actions.

If we have an action of a group G on a set X, then the *orbit* of $x \in X$ under G is the set

$$G \cdot x = \{g \cdot x \mid g \in G\}.$$

So, for example, the orbit of a vertex of an *n*-gon under the action of D_n on the set of vertices is the full set of all vertices; in such a situation we say that the action is *transitive*. At the other extreme, the orbit of *x* under the trivial action is simply $\{x\}$. If no nontrivial element of *G* fixes any $x \in X$, we say that the action of *G* on *X* is *free*.

Exercise 3. List all of the orbits for your favorite group actions (for instance, the ones listed above).

Groups acting on themselves. If G is any group, then G acts on itself in two ways:

- 1. *G* acts on itself by left multiplication: $g \cdot h = gh$
- 2. *G* acts on itself by conjugation: $g \cdot h = ghg^{-1}$

The first action is transitive and free, but the second usually is not (the orbits in the second case are called *conjugacy classes*).

Exercise 4. Check that the purported actions of G on itself are really actions.

Exercise 5. Describe the conjugacy classes in S_n , D_n , F_n , and \mathbb{Z}^n .

The action of G on itself by left multiplication gives a homomorphism $G \rightarrow S_G$. This homomorphism is injective (why?). Therefore, we have the following theorem.

THEOREM 2.1 (Cayley's theorem). Every group is isomorphic to a subgroup of some symmetric group.

Cayley's theorem already tells us that every group can be thought of as a set of symmetries (of the group itself). However, there is no geometry here: we are just thinking of the group as a discrete set of points. The real goal is to find group actions on geometric objects so that we can use geometry to study the group. In this book the geometric objects we will think about are graphs and metric spaces, the subjects of the next two sections.

2.1 GRAPHS

We'll now describe our first class of geometric objects—graphs. We'll say what it means for a group to act on a graph, and finally we will show that every group is the group of symmetries of some graph, its Cayley graph.

Graphs. A graph Γ is a combinatorial object defined as follows: there is a nonempty set V (the set of vertices) and a set E (the set of edges), and a function from E to the set of unordered pairs of (not necessarily distinct) elements of V; we'll call this function the *endpoint function*. We sometimes say that an edge *connects* its two endpoints.

We can draw a picture of Γ by drawing one dot for each element of V and connecting dots with arcs—one for each edge—according to the endpoint function; see, for instance, Figure 2.1.

Exercise 6. Write down a formal description of the graph in Figure 2.1, using sets *V* and *E* and the endpoint function.

The *valence* of a vertex in a graph is the number of times that vertex appears as the endpoint of an edge. In other words, this is the number of half-edges incident to the vertex. In Figure 2.1 every vertex has valence 4.



Figure 2.1 A graph.

Graphs usually are described not by explicitly listing the sets V and E, but rather by giving some rule for listing the vertices and edges. Here are some examples:

1. The *complete graph* on *n*-vertices K_n has vertex set $\{1, \ldots, n\}$ and an edge for every pair of distinct vertices.

Notice that in our picture of K_4 there is an intersection right in the middle. This is not a vertex! It is simply an artifact of the picture. Can you find a picture of K_4 without any extra crossings like this? What about K_5 ?



2. The *complete bipartite graph* on $K_{m,n}$ has vertex set $V = V_1 \coprod V_2$ where V_1 and V_2 have *m* and *n* elements, respectively. There is an edge for each pair $\{v, w\}$, where $v \in V_1$ and $w \in V_2$. If we think of the elements of V_1 as "red vertices" and the elements of V_2 as "green vertices," then there are edges connecting every red vertex to every green vertex, and no edges connecting vertices of the same color. (More generally, a *bipartite graph* is one where there are "red" and "green" vertices and all edges connect vertices of different colors.)



3. Here's an interesting example: consider the graph where the vertices are the two-element subsets of {1, 2, 3, 4, 5} and where there is an edge between two

subsets if they are disjoint. This graph is known as the Petersen graph; it is a famous and important example in graph theory. Each vertex of the Petersen graph has valence 3. For example, the vertex $\{1, 2\}$ is adjacent to $\{3, 4\}$, $\{3, 5\}$, and $\{4, 5\}$.



4. A *tree* is a connected graph without a cycle. Finite trees are very important when studying searching and sorting algorithms in computer science. They also make an appearance in Sean Cleary's Office Hour 16 on Thompson's group. Infinite trees play an essential role in geometric group theory; see Office Hour 3. Below are two pictures of trees, the one on the right is a portion of T_3 , the tree in which every vertex has valence 3. Sometimes T_3 is called a regular 3–valent tree.



5. Graphs are commonly used to describe networks or to show complex relationships. For instance, we could create a graph whose vertices are the members of a social network (such as Facebook) and whose edges connect two members that are friends. (In some social networks, such as Google Plus, it is possible for friendships to be one-directional; to describe such a social network we need directed graphs, which we'll get to in a little bit.)

Here is a similar example: consider the graph that has as vertex set the set of movie actors and has edges between actors who have appeared in the



Figure 2.2 Two isomorphic graphs.

same movie. Extensive knowledge of this graph would be useful in playing the parlor game Six Degrees of Kevin Bacon.

Variations. In certain contexts, we might be interested in variations on the definition of a graph. For instance, in some cases we will care about *loopless graphs* (no edges connecting a vertex to itself) or *simple graphs* (at most one edge between two given vertices).

We also might be interested in *directed graphs*, where the edges have preferred directions, and in *labeled graphs*, where the vertices and/or edges come with a label (e.g., a letter or a color). When we draw a picture of a directed graph, we usually draw little arrows on the edges to indicate the preferred direction. We already encountered directed graphs in Office Hour 1 with edge labels when we constructed the geometric object that has \mathbb{Z}^2 as its group of symmetries.

Exercise 7. Give the formal definitions of simple graphs, loopless graphs, directed graphs, and labeled graphs.

Graph isomorphisms. You should have some intuitive idea of when two graphs are the same, for example, the two graphs in Figure 2.2. Formally, if Γ and Γ' are graphs with vertex sets V and V' and edge sets E and E', then an *isomorphism* $\Gamma \rightarrow \Gamma'$ is a pair of bijections $V \rightarrow V'$ and $E \rightarrow E'$ that are compatible with the endpoint function in the sense that the endpoints of the image of an edge are the images of the endpoints of the edge. So just as an isomorphism of graphs is a bijection (or rather a pair of bijections) that preserves all of the graph structure. For either groups or graphs, an isomorphism is a certificate that two a priori different objects are really the same object.

Sometimes the isomorphism is not clear, as in Figure 2.2. You might be surprised to learn that both graphs in Figure 2.3 are isomorphic to the Petersen graph!

Exercise 8. Find an explicit isomorphism between the graphs in Figure 2.2.

Exercise 9. Show that $K_{m,n}$ is isomorphic to $K_{n,m}$.

Exercise 10. Show that both graphs in Figure 2.3 are isomorphic to the Petersen graph.



Figure 2.3 Two graphs isomorphic to the Petersen graph.

If we are working with graphs that have extra structure (i.e. directed and/or labeled edges) then an isomorphism must preserve this extra structure. For instance, if we have two graphs with blue and green edges, then an isomorphism should take blue edges to blue edges and green edges to green edges. Similarly, if we have directed graphs, then a graph isomorphism should respect the arrows on the edges.

Graph automorphisms. An isomorphism from a graph to itself is called an *automorphism*. The automorphisms of a graph are exactly what you should think of as its symmetries. The set of automorphisms of a graph Γ forms a group denoted by Aut(Γ).

The graph on the left-hand side of Figure 2.2 has lots of obvious automorphisms, for instance, rotation by $\pi/2$ and reflection about the horizontal. It also has less-obvious automorphisms, for instance, ones that interchange the inner and outer squares.

Since the graph on the right-hand side of Figure 2.2 is isomorphic to the one on the left-hand side, it has the same group of automorphisms. What do the less-obvious automorphisms from the left-hand side look like on the right-hand side?

Exercise 11. Determine the automorphism group of the graph in Figure 2.2. *Hint: the discussion at the beginning of Office Hour 1 should help.*

We already have a number of examples of graphs. Here are some of their automorphisms:

- 1. Since each vertex in K_n is connected to every other vertex, all vertices are interchangeable; thus $Aut(K_n) \cong S_n$.
- 2. Similarly, we find that if $m \neq n$, then $\operatorname{Aut}(K_{m,n}) \cong S_m \times S_n$. What about $\operatorname{Aut}(K_{n,n})$?
- 3. From our original definition of the Petersen graph (using subsets of $\{1, 2, 3, 4, 5\}$) we see that S_5 is a subgroup of the automorphism group.
- 4. The regular 3-valent tree T_3 has lots of automorphisms. For instance, we can rotate by $2\pi/3$ at any vertex or swap two branches at a vertex (of course we need to be careful about how to define the automorphism on the rest of the tree). There are infinite order automorphisms as well. It turns out that Aut(T_3) is uncountable!

5. If two people have the same friends on Facebook, then there is an automorphism of the Facebook graph that interchanges these two people and fixes every other vertex (whether or not the two people are friends with each other!).

Exercise 12. Is *S*₅ the full group of automorphisms of the Petersen graph?

Exercise 13. Describe an infinite order element in $Aut(T_3)$.

Exercise 14. Besides the one mentioned above, can you think of other scenarios that lead to graph automorphisms of the Facebook graph?

If a graph Γ has extra structure (i.e., directed and/or labeled edges), then automorphisms of the graph are required to preserve this extra structure (after all, a graph automorphism is a special case of a graph isomorphism). The result is still a group and is still denoted by Aut(Γ). For instance, in Office Hour 1 we considered automorphisms of a regular *n*-gon, thought of as a directed graph where the edges are all directed in the counterclockwise direction. As we saw, the group of automorphisms of this directed graph is $\mathbb{Z}/n\mathbb{Z}$.

Group actions on graphs. An *action* of a group G on a (possibly directed, possibly labeled) graph Γ is a homomorphism $G \to \operatorname{Aut}(\Gamma)$. So it is based on the same idea as a group action, but an action on a graph needs to preserve the structure of the graph.

If we unravel the definition of a group action of *G* on an (unlabeled, undirected) graph Γ with vertex set *V* and edge set *E*, we have:

- for each $g \in G$ and $v \in V$ a choice of element of V, denoted by $g \cdot v$,
- for each $g \in G$ and $e \in E$ a choice of element of E, denoted by $g \cdot e$,
- for x in either V or E, we have $1 \cdot x = x$,
- for $g, h \in G$ and x in V or E, we have $g \cdot (h \cdot x) = (gh) \cdot x$, and
- if the endpoints of $e \in E$ are v and w, then the endpoints of $g \cdot e$ are $g \cdot v$ and $g \cdot w$.

In (yet) other words, a group action of G on Γ is a pair of group actions of G on V and E that are compatible with the endpoint function.

If we take Γ to be a labeled graph, then the second condition above would further stipulate that the labels on *e* and $g \cdot e$ need to be the same.

If we take Γ to be a directed graph, then the last condition above would say: if the initial and terminal endpoints of $e \in E$ are v and w, respectively, then the initial and terminal endpoints of $g \cdot e$ are $g \cdot v$ and $g \cdot w$, respectively.

We can rephrase our list of automorphisms of graphs in terms of group actions:

- 1. S_n acts on K_n by permuting the vertices
- 2. $S_m \times S_n$ acts on $K_{m,n}$ by permuting the two sets of vertices
- 3. S_5 acts on the Petersen graph via its action on $\{1, 2, 3, 4, 5\}$
- 4. $\mathbb{Z}/3\mathbb{Z}$ acts on the regular 3-valent tree T_3 (as do $\mathbb{Z}/2\mathbb{Z}$ and \mathbb{Z})
- 5. $\mathbb{Z}/2\mathbb{Z}$ sometimes acts on the Facebook graph (for instance, when there are two people with the same set of friends, as above)

Cayley graphs. Finally, we come to the whole reason we have been studying graphs. Let *G* be an arbitrary group and let *S* be a generating set for *G*. For simplicity we assume that *S* does not contain the identity. The *Cayley graph* for *G* with respect to *S* is a directed, labeled graph $\Gamma(G, S)$ given as follows: the vertex set is *G*, there is a directed edge from *g* to *gs* for every $g \in G$ and $s \in S$, and we label the edge from *g* to *gs* by the element *s* of *G*. (If *S* is small we can sometimes use colors for the edges instead of labels.)

You might wonder what happens if you take two different generating sets for the same group. How are the two Cayley graphs related? They might not be isomorphic (for instance, if the cardinalities of the generating sets are different). However, there is a sense in which they are the same; see Office Hour 7 on quasi-isometries for an explanation.

The Cayley graphs for \mathbb{Z} with respect to {1} and for \mathbb{Z}^2 with respect to {(1, 0), (0, 1)} are familiar. Indeed, we already discovered them in Office Hour 1! Each blue edge in the graph for \mathbb{Z} corresponds to the generator {1}. The green and blue edges in the graph for \mathbb{Z}^2 correspond to the generators (1, 0) and (0, 1), respectively. (Of course, 1 here is the number 1, not the identity.)



Here is the Cayley graph for S_3 with respect to the generating set {(1 2), (2 3)}. The blue edges correspond to the generator (1 2) and the green edges correspond to the generator (2 3).





Figure 2.4 Part of a Cayley graph.

Exercise 15. Draw the Cayley graph for \mathbb{Z} with the generating set {2, 3}, for D_4 with the generating set (s, t), for S_3 with the generating set {(1 2), (1 2 3)}, and for S_4 with the generating set {(1 2), (2 3), (3 4)}.

Here's a useful observation about the Cayley graph made evident by Figure 2.4:

If s_1, s_2, \ldots, s_n are the labels (in order) on an edge path in $\Gamma(G, S)$ from 1 to g, then $g = s_1 s_2 \cdots s_n$. And vice versa, if $g = s_1 s_2 \cdots s_n$, then there is an edge path from 1 to g in $\Gamma(G, S)$ whose labels are s_1, s_2, \ldots, s_n (in that order).

(We allow s_i to be the inverse of an element *s* of *S*, meaning that our path crosses an *s*-edge against the direction of the edge.) So there is a correspondence between paths in the Cayley graph from 1 to *g* on one hand and words in the generators representing *g* on the other hand.

Exercise 16. What do the labels on an edge path from *g* to *h* in $\Gamma(G, S)$ tell you about *g* and *h*?

Automorphisms of the Cayley graph. Each element of a group G gives rise to an automorphism of the Cayley graph $\Gamma(G, S)$, as follows. The automorphism Φ_g associated to $g \in G$ is given on the vertices by:

$$\Phi_g(v) = gv.$$

This is just the action of G on itself by left multiplication. Since there is at most one directed edge connecting any two vertices with a specified direction, and since no edges connect a vertex to itself, there is only one way to extend this to an automorphism of the entire directed, labeled graph: Φ_g sends the edge $v \rightarrow vs$ to the edge $gv \rightarrow gvs$. This works since $\Phi_g(v) = gv$ and $\Phi_g(vs) = gvs$.

THEOREM 2.2. Let S be a generating set for G. Then the map $G \to \text{Aut}(\Gamma(G, S))$ is an isomorphism.

The proof requires one new idea. Let S^{-1} denote the set of inverses of elements of *S*. The *word length* of $g \in G$ with respect to *S* is the length of the shortest word in $S \cup S^{-1}$ that is equal to g in G. The identity $1 \in G$ has word length 0. By the definition of a generating set, every element of G has finite word length with respect to *S*.

Our earlier observation about edge paths in Cayley graphs shows that we can also think of word length geometrically: the word length of g is the fewest number of edges in an edge path from 1 to g.

Let's prove Theorem 2.2. To simplify notation, let $\Gamma = \Gamma(G, S)$. It is not hard to see that the map $G \to \operatorname{Aut}(\Gamma)$ we have already defined is a homomorphism. It is injective since different elements of G send the identity vertex of Γ to different vertices. The main thing is to show that it is surjective, that is, every element of $\operatorname{Aut}(\Gamma)$ is equal to some Φ_{g} .

Let $\Phi \in \operatorname{Aut}(\Gamma)$, that is, Φ is an isomorphism $\Gamma \to \Gamma$, where Γ is regarded as a directed graph with edge labels. Suppose that $\Phi(1) = g$ (we are using the same symbol for a vertex and the corresponding element of Γ). We will show that $\Phi = \Phi_g$. As above, since directed edges are determined by their endpoints, it is enough to show that Φ agrees with Φ_g on the vertices of Γ .

We use induction on word length in *G* with respect to *S*. The base case is word length 0, which is just the statement $\Phi(1) = g = \Phi_g(1)$. Assume Φ agrees with Φ_g on all elements of *G* of word length *n* with respect to *S*. Suppose $v \in G$ has word length n + 1. This means that v = ws, where the word length of *w* is *n* and $s \in S \cup S^{-1}$. To simplify matters, assume $s \in S$; the other case is similar.

By assumption, $\Phi(w) = \Phi_g(w)$. There is a unique edge labeled *s* emanating from *w*, namely, the edge from *w* to ws = v. Similarly, there is a unique edge labeled *s* emanating from $\Phi(w)$, with ending point $\Phi(w)s$. Since Φ and Φ_g both respect edge labels, it must be that $\Phi(v) = \Phi_g(v) = \Phi(w)s$. In particular, $\Phi(v) = \Phi_g(v)$, which is what we wanted to show!

The Cayley graph for the free group. By Theorem 2.2, we know there is a geometric object with F_2 as its group of symmetries, namely, the Cayley graph for F_2 (with respect to any generating set). But what does such a Cayley graph look like?

Let *a* and *b* be the standard generators for F_2 . We can draw a picture of the Cayley graph $\Gamma = \Gamma(F_2, \{a, b\})$ as follows. We start by drawing the vertex for the identity element 1 of F_2 (the empty word):

The four edges containing 1 connect 1 to the vertices a, b, a^{-1} , and b^{-1} , the four elements of F_2 with word length 1:

1

