

Group Theory at Mathcamp

July 6–10, 2010

Nic Ford

1 Symmetries and the Definition of a Group

1.1 Symmetries

The point of these notes is to introduce you to the definition of a group and to help you understand how groups work and why one would want to study them. The primary goal here is the development of intuition rather than the introduction of formalism. Therefore, we spend a lot of time, especially in the beginning, with examples. There are exercises at the end of most of the sections, and if you're reading these notes outside the context of a course, it is a good idea to attempt all or most of them.

It will be instructive to begin with a short discussion of symmetry: historically, this is one of the first situations in which the theory of groups was developed, and it's also probably the easiest to understand. When I was in third grade or so, one of the things they made us do in math class was to take pictures of some capital letters and draw "lines of symmetry" on them. The lines we were drawing, they told us, were the ones with the property that the letter "looks the same" on either side of the line. For example, the letter A gets just one line, going vertically through the middle, whereas the letter H gets two, one vertically down the middle, and another one horizontally across the middle. The lines are like mirrors: the part of the letter on one side is the reflection of the part on the other side.

This made for a fun exercise for third-graders, but unfortunately it fails to capture the essence of what we mean when we say that an object is symmetrical. To see why, consider the letters N and R. If you try to draw lines of symmetry for either one, you'll find that you can't. There's no way to place a "mirror" in the middle of either of these letters. (If you don't believe me, try for yourself.) But when one looks at the shapes of both letters, one can't help but feel that N is somehow "more symmetrical" than R.

This feeling, it turns out, is entirely justified. The letter N is symmetrical, just in a different way: if you rotate it 180 degrees, it looks the same, whereas this is definitely not true of R. This suggests a new definition, which is the one we'll use: a *symmetry* of some object is some way of moving the object around¹ which leaves it looking the same as it started.

This new definition of symmetry includes our old one, and we find that even some shapes which have "mirror" symmetries also have rotational symmetries, like the letter N. For example, consider the letter H (Figure 1).

You can flip it vertically or horizontally², or rotate it by 180 degrees. In addition to these three symmetries, there is a fourth, of a type we haven't considered yet: you can simply do nothing. This transformation, called the *identity*, is always a symmetry of any shape you might happen to be looking at. (While this may seem like a strange point to be emphasizing, the fact that doing nothing counts as a symmetry will be important later.) So the letter H has four symmetries. The letter A, by contrast, only has two, as you can see in Figure 2.

It can either be flipped horizontally or left alone.

¹When we talk about "moving an object around," we need to be a bit more precise: we restrict our attention to rigid transformations which keep distances the same, like rotations and reflections. These are called *isometries*.

²There is some ambiguity in the phrase "flip vertically": does it mean a reflection around a vertical line, or a reflection which moves points vertically? In these notes, we'll always mean the latter, as in the leftmost H in Figure 1.




Figure 1: The symmetries of the letter H. In order: a vertical flip, v ; a horizontal flip, h ; a 180-degree rotation, r^{180} ; and the identity, i .




Figure 2: The symmetries of the letter A: the identity, i , and a horizontal flip, h .

We can use this to compare the shapes of the letters we've been looking at. We already observed that some letters have non-identity symmetries (like N) and others don't (like R), but we can also make finer distinctions than that. For example, the symmetries of the letters M, T, and Y are like those of the letter A; the number 8 is like H; the letters S and Z are like N. Indeed, N and A have something in common as well: they both have exactly two symmetries, whereas H has four and R has one.




Figure 3: These two shapes each have four symmetries, but they have different symmetry groups.

But we can do more with symmetries than just count them. Consider the two shapes in Figure 3. It is not difficult to see that the one on the left has exactly the same four symmetries as the letter H. (In fact, it's just an H with three more lines drawn on it, and all the extra lines look the same after performing any of the symmetries of H.) The one on the right also has four symmetries: it can be rotated by 90, 180, or 270 degrees, or it can be left alone. None of its symmetries can be reflections, since a reflection would switch which side of each small arc has the dot on it.

But there is an important difference between the symmetries of these two shapes, which comes out when we look at how the symmetries interact with each other. Imagine performing one of the symmetries of H twice. Whether you reflect twice or rotate twice by 180 degrees, the end result is that everything is back where it started. That is, doing any of the symmetries of the shape on the left twice is the same as just doing the identity. But this isn't true of the shape on the right: for example, rotating twice by 90 degrees doesn't give you the identity, it gives you a rotation by 180 degrees. So even though both shapes have the same number of symmetries, we see that there is a crucial difference between them. For the shape on the left, doing any symmetry twice gives you the identity, but this isn't true of the shape on the right.

1.2 Generalizing Symmetry

The collection of symmetries of some shape, like the four symmetries of H and the two symmetries of A depicted above, are an example of a *group*, a type of mathematical object which we'll define shortly. When we study groups of symmetries, we will be interested only in how the symmetries interact with each other, rather than in how they interact with whatever object we were looking at. For example, if you rotate by 180 degrees twice, or if you flip horizontally twice, the result is the same as if you had done nothing at all. In this sense, the groups of symmetries of A and N look the same: they both have two symmetries, one of which is the identity, and the other of which gives you the identity if you do it twice. The fact that the non-identity element of A's group is a reflection and not a rotation is not the sort of thing we'll be examining; the fact that all the symmetries in both groups interact with each other in the same way is the sort of thing we'll be examining.

So a group of symmetries can be completely described by listing all the symmetries and what happens when you perform one followed by another. The process of performing one symmetry followed by another is called *composing* the symmetries. For example, when you compose a horizontal flip with a vertical flip, you get a 180-degree rotation.

We might express all this information in a table. The table on the left describes the symmetries of the first shape in the picture above, and the table on the right describes the symmetries of the second shape. Here i is the identity, h is a horizontal flip, v is a vertical flip, and r^n is a rotation by n degrees. To get the entry in each spot in the table, compose the entry on the left side with the entry on the top.

	i	h	v	r^{180}
i	i	h	v	r^{180}
h	h	i	r^{180}	v
v	v	r^{180}	i	h
r^{180}	r^{180}	v	h	i

	i	r^{90}	r^{180}	r^{270}
i	i	r^{90}	r^{180}	r^{270}
r^{90}	r^{90}	r^{180}	r^{270}	i
r^{180}	r^{180}	r^{270}	i	r^{90}
r^{270}	r^{270}	i	r^{90}	r^{180}

It will be helpful later to have a notation for the composition of two symmetries, so we'll introduce one now: if a and b are two symmetries, then their composition will be written $b \circ a$. That is, if you perform a and then perform b , the result is the same as performing $b \circ a$. (This notation seems backwards, but there is a good reason for it, which is related to the corresponding backwardness for function applications: $f(g(x))$ means take x , do g , then do f . In Section 3, we will be happy to have made this convention.) Each table above can be thought of as a sort of "multiplication table" for the corresponding group of symmetries. For example, the table on the left tells us that $h \circ r^{180} = v$.

The tables, which we'll call *composition tables*, also tell us that composing with the identity always leaves you with the same symmetry you started with. (For example, $i \circ r^{270} = r^{270}$ and $h \circ i = h$.) This, of course, is because the identity is the transformation that does nothing. The composition tables should also make it clear why we wanted to consider the identity as a symmetry in the first place. If we didn't, we would have spots in the table that we couldn't fill in, like $h \circ h$.

Here are the composition tables for the symmetries of A and the symmetries of N respectively:

	i	h
i	i	h
h	h	i

	i	r^{180}
i	i	r^{180}
r^{180}	r^{180}	i

Notice that the only difference between these two composition tables is the labeling: all we need to do to turn one table into the other is switch the names of " h " and " r^{180} ." But such a relabeling isn't possible for the first pair of tables. No matter what names you give to each of the symmetries, there's no way to stop each symmetry of the left shape from giving you the

identity when you compose it with itself. This property has a name: we say that the symmetry group of A is *isomorphic* to the symmetry group of N , whereas the symmetry groups of the two shapes represented by the first pair of tables are not. (We will give a more precise definition of isomorphism in a later section.) In group theory, we will be studying the structure of objects like these groups of symmetries, so two isomorphic groups will be regarded as being essentially the same object.

Once you have the composition table in front of you, there isn't actually any need to refer to the geometric object that the table came from except for intuition. For the types of questions we are interested in — how these transformations interact with each other — it's enough just to have the table. The geometric information is redundant, and in fact, it can make two groups look different (like the symmetries of A and N) when, for our purposes, we ought to treat them the same.

When modern mathematicians confront a situation like this, their solution is usually to abstract away from the things that are providing the redundant information, that is, to come up with a way of describing the objects they're looking at that doesn't require them to start from, say, a drawing of a shape if they're just going to forget about it later. In this case, the thing that will come out of this abstraction process will be the definition of a group. So what we need to do is find a way of describing what it is that groups of symmetries all have in common other than the fact that they happen to have come from transformations of shapes.

We have already mentioned two of those things. The first is that symmetries can be composed. That is, given two symmetries, a and b , you can produce a new symmetry $b \circ a$. Since we are attempting to describe symmetry groups without referring to the shapes they came from, every property we discuss will have to be phrased as a property of this composing process rather than as a property of shapes.

The second fact that we've already observed is that, no matter which symmetry group you're looking at, there is a special symmetry that is always there: the identity. When you compose with the identity, you always get the element you started with, as mentioned above, and this property of the identity is the sort of thing that can go into our abstract definition of a group. Notice that we can no longer describe the identity as "the symmetry that does nothing to the shape," because that requires us to refer to the shape. The thing that makes the identity special is its relationship to composition.

The last two properties are things we haven't mentioned yet, but are still true of any conceivable symmetry group. The first of these is that any symmetry is reversible. That is, given any symmetry a , there's another symmetry b , called the *inverse* of a , which does a "backwards." In terms of composition, this means that $b \circ a = i$: if you do a and then b , you're back where you started. For example, in the symmetry groups discussed above, h and v are their own inverses, and r^{270} is the inverse of r^{90} , as you can see in Figure 4.




Figure 4: The operations of rotating by 90 degrees and rotating by 270 degrees are inverses.

The last property, called *associativity*, has to do with sequences of three symmetries. In terms of compositions, it says that $c \circ (b \circ a) = (c \circ b) \circ a$. In other words, whether you first do a and then do $b \circ c$, or you first do $a \circ b$ and then do c , you get the same answer: either way, it has the effect of doing a , then b , then c .

With this, we've finally arrived at the definition of a group:

Definition. A *group* is a collection of elements, together with a *composition law*, that is, a way of taking any two elements a and b and producing a third, $a \circ b$, with the following properties:

- There is an element, called the *identity* and written i , with the property that $i \circ a = a \circ i = a$ for every other element a in the group.
- Every element a in the group has an *inverse*, written a^{-1} , with the property that $a \circ a^{-1} = a^{-1} \circ a = i$.
- Composition is *associative*: given any three elements a , b , and c in the group, $a \circ (b \circ c) = (a \circ b) \circ c$.

As promised, this definition doesn't refer to any properties of the elements other than their relationships to each other. The elements don't have to be symmetries of some shape; there doesn't even need to be any shape in sight at all. As we will see in a moment, there are many examples of groups that have nothing to do with symmetries at all, and the theory of groups has many applications that are not directly related to geometry. So by abstracting away from the geometry of symmetries will accomplish two things. The first is the one we've already mentioned: we will be able to speak about the relevant properties of symmetry groups with a clearer sense of which details are important and which are not. But the second, which is perhaps more valuable, is that the things we learn about groups will have applications to more things than just symmetry groups, as we'll see as we go on.

2 Examples and Basic Properties of Groups

Before we go on, we're going to discuss a few examples of groups.

2.1 Groups of Numbers

Our first collection of examples consists of groups whose elements are numbers. In particular, they don't come from symmetry groups of shapes. If there's any fact here that isn't clear to you, you should take some time to convince yourself before moving on.

- The integers form a group, usually written \mathbb{Z} , in which the composition law is addition. That is, we can declare that $a \circ b = a + b$ for two integers a and b . The identity is 0, and the inverse of some number n is $-n$. Addition is associative, so everything works.
- The real numbers also form a group in which the operation is addition, written \mathbb{R} , in exactly the same way.
- Neither the integers nor the real numbers form a group with subtraction as the composition law. There are two problems: 0 looks like an identity if you put it on the right ($x - 0 = x$) but not on the left ($0 - x \neq x$ unless $x = 0$). Also, subtraction isn't associative. For example, $(1 - 1) - 1 = -1$, but $1 - (1 - 1) = 1$.
- The real numbers almost form a group with multiplication as the composition law. Multiplication is associative, and 1 is the identity. The only problem is that 0 doesn't have an inverse. If you throw out 0, though, you do get a group, called \mathbb{R}^\times , in which the inverse of any number x is $1/x$.
- If you try to make multiplication work by doing the same thing to \mathbb{Z} , you run into more problems than just 0: the only integers with multiplicative inverses are 1 and -1 . Even though, say, 3 has a multiplicative inverse in \mathbb{R} , it doesn't have one in \mathbb{Z} , and so the nonzero integers still don't form a group. You can form a group out of just 1 and -1 , and this is isomorphic to the symmetry group of A that we discussed before: it has just one non-identity element whose composition with itself (which in this case means its square) is the identity.


2.2 Cyclic Groups

There is a class of groups related to the first example above which we'll talk about a few more times throughout these notes. Write \mathbb{Z}_n for the collection of numbers $0, 1, 2, \dots, n - 1$. Imagine the numbers arranged in a circle, starting with 0, then 1, and so on, with $n - 1$ at the end next to 0. (For example, if $n = 12$, you'll end up with the standard arrangement of the numbers on a clock face, except with a 0 instead of the 12.) To compose two numbers, say a and b , start at the spot marked a on the circle and move forward b spots. For example, to compose 7 and 3 in \mathbb{Z}_9 , start with the circle of numbers from 0 to 8, and move three spots forward from 7. The number you land on is 1, as seen in Figure 5.


We'll write \oplus instead of \circ for the operation in this group, so, for example, $7 \oplus 3 = 1$ in \mathbb{Z}_9 .

Notice that as long as you don't pass 0 in your path around the circle, the number you end up with is just $a + b$; for example, in \mathbb{Z}_{11} , $4 \oplus 2 = 6$. (In particular, 0 is the identity of \mathbb{Z}_n , though you could also figure this out directly from the definition.)

This leads us to a slightly different description of the \oplus operation which will help us to prove that \mathbb{Z}_n is a group. Given two numbers a and b , there are two things that can happen when

Figure 5: In \mathbb{Z}_9 , $7 \oplus 3 = 1$.

computing $a \oplus b$: either you don't pass 0 on the way from a to $a \oplus b$, or you do. In the first situation, you get $a + b$, and in the second, you get $a + b - n$. This is because, as we move forward, we increase the number by 1 each time, but when we go from $n - 1$ to 0, our count drops by n from what it should be. You can imagine a second circle sitting underneath the first one in which all the numbers are n bigger. When we pass 0, $a + b$ shows up as our position on the other circle. For example, Figure 6 demonstrates that in \mathbb{Z}_6 , $4 \oplus 4 = 4 + 4 - 6 = 2$.

Figure 6: When computing $4 \oplus 4$ in \mathbb{Z}_6 , you can first compute $4 + 4$, then subtract 6 to get back into the range $0, 1, 2, 3, 4, 5$.

Put another way, computing $a \oplus b$ is as simple as taking $a + b$ and then subtracting n until you end up in the range $0, 1, 2, \dots, n - 1$. This tells us two things. The first is that the inverse of a is always $n - a$, because $a + (n - a) = n$, and after we subtract n from that, we end up with 0. The second is that \oplus is associative: regular addition is associative, so $a \oplus (b \oplus c)$ and $(a \oplus b) \oplus c$ are both just $a + b + c$ minus enough copies of n to get the number into the right range. This was all we needed to make \mathbb{Z}_n into a group.

The groups \mathbb{Z}_n are called *cyclic groups*, and they will be an important source of examples later on. The description of \oplus that we just gave can be used to draw up a composition table for the cyclic groups like the ones in Section 1. One for \mathbb{Z}_5 is shown in Table 1.

2.3 Dihedral Groups

The *dihedral groups* are the symmetry groups of regular polygons. The symmetry group of a regular polygon with n sides will be written D_n . In all the groups we've considered up to this point, the order in which we composed elements hasn't mattered, that is, $a \circ b$ is always equal to $b \circ a$. Notice, though, that this isn't part of the definition of a group, and in fact, there are groups in which it doesn't always happen. When $a \circ b = b \circ a$, we say that a and b *commute*, and a group in which elements always commute, like all the ones we've looked at so far, is called *abelian*.

Consider D_3 , the symmetry group of an equilateral triangle. Just as in Section 1, there are two types of symmetries: rotations and reflections. For example, you can rotate the triangle by 120 degrees, or reflect it around a line going through one of the vertices and the opposite edge. We can place a small mark on one of the corners of the triangle to keep track of what's happening to it when we perform these symmetries. By doing this, we can see very easily that the two symmetries we just mentioned don't commute, as shown in Figure 7. Therefore, D_3 is not abelian.




Figure 7: The elements h and r^{120} in D_3 don't commute.

There are exactly two rotations that act as symmetries of the triangle: rotating by 120 degrees or by 240 degrees. There are also three reflections: for each vertex of the triangle, you can reflect around the line from that vertex to the opposite edge. Together with the identity, this makes for a total of six symmetries. If you're not convinced that this is all the symmetries there are, you should take a moment to verify it by considering what each symmetry can do to each of the vertices of the triangle.

Just like we did for the symmetry groups in Section 1, we can give names to each of the symmetries in D_3 in order to write up a composition table. Let's call the rotations r^{120} and r^{240} and the flips h , d_1 , and d_2 according to Figure 8.




Figure 8: Names for the elements of D_3 .

Then the composition table is given in Table 2.

In fact, a regular polygon with n sides always has exactly $2n$ symmetries. To see that this is true, we pick one of the vertices of the polygon. Any symmetry is going to have to send that vertex to another vertex, so there are n possible choices for what can happen to the one we chose. Since

	0	1	2	3	4
0	0	1	2	3	4
1	1	2	3	4	0
2	2	3	4	0	1
3	3	4	0	1	2
4	4	0	1	2	3

Table 1: A composition table for \mathbb{Z}_5 .

the symmetry can't change the shape of the polygon, the neighbors of our chosen vertex have to be taken to the neighbors of the new vertex. The only thing that can change is which one is on which side. There are two choices for this: either the neighbor on the left stays on the left, or it switches to the right, and vice versa (Figure 9.)




Figure 9: How to pick a symmetry of a polygon. First, decide where the chosen vertex goes, then decide whether its neighbors should switch places.

But once these choices have been made — where our chosen vertex goes and where to put its neighbors — everything else about the symmetry is completely determined. The neighbors of the first vertex have been placed, so *their* neighbors only have one possible place to land, so the same is true of *their* neighbors, and so on. Since there were n choices for where to put the first one, and each of those choices gave 2 further choices, there are a total of $2n$ ways to pick our symmetry.

Just as there were two rotations in D_3 , there will be $n - 1$ rotations among the symmetries in D_n : after picking a vertex, you can rotate the polygon until it lands on any other vertex of your choice, and there are $n - 1$ other vertices to choose from. The other n symmetries are all flips, which you should verify for yourself.

Look at the part of the composition table for D_3 that just comes from the rotations and the identity:

	i	r^{120}	r^{240}
i	i	r^{120}	r^{240}
r^{120}	r^{120}	r^{240}	i
r^{240}	r^{240}	i	r^{120}

Notice that the only entries in the table are the rotations and the identity. This makes sense: you can't get a flip by composing two rotations. In fact, even more is true: these elements form a group all by themselves. They have the identity, they can be composed, and the inverse of every element is already present. We say that the rotations and the identity form a *subgroup* of D_3 , meaning a subset which is still a group with the same composition law. (We'll discuss subgroups a bit more later on.) This subgroup is, in fact, isomorphic to \mathbb{Z}_3 : r^{120} can serve the role of 1 and r^{240} can be 2.

	i	r^{120}	r^{240}	h	d_1	d_2
i	i	r^{120}	r^{240}	h	d_1	d_2
r^{120}	r^{120}	r^{240}	i	d_1	d_2	h
r^{240}	r^{240}	i	r^{120}	d_2	h	d_1
h	h	d_2	d_1	i	r^{120}	r^{240}
d_1	d_1	h	d_2	r^{240}	i	r^{120}
d_2	d_2	d_1	h	r^{120}	r^{240}	i

Table 2: A composition table for D_3 .

2.4 Basic Properties of Groups

Working directly from the definition of a group, it's possible to prove facts about all groups at the same time, whether or not they're the groups we've talked about so far. For example, the *cancellation law* says that if you have three elements a , b , and c in any group and $a \circ b = a \circ c$, then $b = c$. You may compose with a^{-1} on the left to get that $a^{-1} \circ (a \circ b) = a^{-1} \circ (a \circ c)$, then use associativity to get $(a^{-1} \circ a) \circ b = (a^{-1} \circ a) \circ c$. By the definition of inverses, we get $i \circ b = i \circ c$, which means that $b = c$.

The cancellation law is itself useful for other things, like the uniqueness of inverses: if $a \circ b = i$, then $b = a^{-1}$; that is, there can only be one element which acts as an inverse for a . This is simply because $i = a \circ a^{-1}$, and we may then apply the cancellation law to the equation $a \circ b = a \circ a^{-1}$.

As we go forward in our study of groups, it will be convenient to have access to a few more pieces of notation and terminology, which we introduce now. For any element a in a group, we write a^n for the element you get by composing a with itself n times. (That is, $a^4 = a \circ a \circ a \circ a$.) Because composition in a group is associative, it doesn't matter how you group the copies of a or in which order you perform the composition.

We then define the *order* of a to be the smallest positive number n so that a^n is the identity. For example, in D_3 , the rotations have order 3, each of the flips has order 2, and the identity has order 1. In the group \mathbb{Z} of integers under addition, no amount of composing 1 with itself will give you the identity, 0. In this case, we say that 1 has "order infinity."

Exercises

1. Write up composition tables for \mathbb{Z}_2 , \mathbb{Z}_3 , and \mathbb{Z}_4 .
2. Recall that two groups are called *isomorphic* if the only difference between their composition tables is the way the elements are labeled, as was the case for the symmetry groups of A and N. Among the four groups whose composition tables appeared in Section 1, three are isomorphic to cyclic groups and one isn't. Which is which?
3. Give an argument like the one for associativity to show that all cyclic groups are abelian.
4. Prove that every group with two elements is isomorphic to \mathbb{Z}_2 and every group with three elements is isomorphic to \mathbb{Z}_3 . Is the same true of groups with four elements?
5. Prove that the cyclic group \mathbb{Z}_{2n} has a subgroup isomorphic to \mathbb{Z}_2 and a subgroup isomorphic to \mathbb{Z}_n .
6. Write up a composition table for D_4 .
7. Prove that in any dihedral group D_n , the rotations form a subgroup which is isomorphic to the cyclic group \mathbb{Z}_n .
8. How many two-element subgroups are there in D_n if n is odd? If n is even?
9. Describe the group of rotational symmetries of the cube. How many elements are there? What do they do to the cube? What are their orders?
10. A *permutation* of a set X is a way of "rearranging" the elements of X , that is, a function from X to itself for which no two elements are sent to the same place. Let S_n be the collection of permutations of the set $\{1, 2, \dots, n\}$. (For example, an element of S_5 is the permutation which sends 1 to 4, 4 to 3, and 3 to 1, and switches 2 and 5.) Permutations can be composed the same way as symmetries: first do one, then the other. Prove that S_n is a group using this composition law. How many elements does it have?

11. Prove that if j is an element of a group which “acts like the identity” in the sense that $j \circ a = a$ for some element a , then j actually is the identity.
12. Prove that, in the composition table for any finite group, every element appears exactly once in each row and exactly once in each column.
13. Prove that, in any group, $(a \circ b)^{-1} = b^{-1} \circ a^{-1}$, and that if $(a \circ b)^{-1} = a^{-1} \circ b^{-1}$, then a and b commute.
14. Prove that an element of a group has order 2 if and only if it is its own inverse.
15. Prove that $(a^n)^{-1} = (a^{-1})^n$. (We usually write this as a^{-n} , since there is no ambiguity about which of these two things we could mean.)
16. Given an element a of a group G , we write $\langle a \rangle$ for the subset of G consisting of all powers (positive or negative) of a . Prove that $\langle a \rangle$ is a subgroup, and that the number of elements it has is equal to the order of a .
17. If the order of a is n (in particular, it's finite), prove that $\langle a \rangle$ is isomorphic to the cyclic group \mathbb{Z}_n .
18. Suppose G is a group with n elements and some element of G has order n . Prove that G is isomorphic to \mathbb{Z}_n .

3 Cosets and Lagrange's Theorem

3.1 The Proof of Lagrange's Theorem

This section is devoted to proving an important theorem about finite groups: if G is a finite group and H is a subgroup of G , then the number of elements in H divides the number of elements in G . This fact, called Lagrange's Theorem, is very useful in the study of group theory, and even has applications to other areas of mathematics, as we'll see at the end of this section.

In order to prove Lagrange's Theorem, we first need to discuss the concept of a coset. Say H is a subgroup of some group G . For any element a in G , we write $a \circ H$ for the set of all elements of the form $a \circ h$ for some element h of H . That is, $a \circ H$ is the set of all elements of G which are expressible as the composition of a and something in H . These are sometimes called *left cosets*; we could equally well be considering *right cosets*, sets of the form $H \circ a$, but the proof of Lagrange's Theorem won't require us to. Before we go any further, we discuss a few examples:

- H itself is always a coset: it's the coset of the identity. In fact, if a is any element H , then the coset $a \circ H$ is just H . Clearly anything in $a \circ H$ is in H , since H is a subgroup and therefore contains all compositions of all of its elements. And if b is some element of H , we can write $b = a \circ (a^{-1} \circ b)$, and $a^{-1} \circ b$ is an element of H , and so b is in $a \circ H$.
- Take $G = \mathbb{Z}$, the group of integers under addition. For any integer m , there is a subgroup of \mathbb{Z} consisting of all multiples of m . (This subgroup is usually denoted $m\mathbb{Z}$.) To see that it's a subgroup, you just need to notice that (1) 0, the identity, is always a multiple of m , no matter what m is, and (2) if ma and mb are two different multiples of m (any multiple of m looks like this by definition) then $ma + mb = m(a + b)$ and $-ma = m(-a)$ are both multiples of m , and so $m\mathbb{Z}$ contains all compositions and inverses of its elements.

If a is some other integer, then the coset $a + m\mathbb{Z}$ consists of every integer of the form $a + mb$ for some b . For example, if $m = 5$ and $a = 2$, then the coset $2 + 5\mathbb{Z}$ contains, among others, the numbers 2, 7, 12, 17, $-3 = -5 + 2$, and $-8 = -10 + 2$.

- Similarly, if $G = \mathbb{Z}_n$ and m is some integer that divides n , then there is a subgroup of \mathbb{Z}_n consisting of the multiples of m , called $m\mathbb{Z}_n$. We check first that this is a subgroup. First, it contains 0. If ma and mb are two elements of $m\mathbb{Z}_n$, then $ma \oplus mb$ is still a multiple of m : we might have to subtract n from $m(a + b)$, but n is a multiple of m , so this is okay. Finally, the inverse of ma is $n - ma$, and since n is a multiple of m , this is as well.

If we take $n = 12$ and $m = 3$, then $3\mathbb{Z}_{12} = \{0, 3, 6, 9\}$, and there are exactly three cosets: there is $0 \oplus 3\mathbb{Z}_{12} = \{0, 3, 6, 9\}$, $1 \oplus 3\mathbb{Z}_{12} = \{1, 4, 7, 10\}$, and $2 \oplus 3\mathbb{Z}_{12} = \{2, 5, 8, 11\}$. Any other coset ends up being the same as one of these three, according to which of these three cosets it belongs to.

- If $G = D_n$ and H is the subgroup consisting of the rotations and the identity, then there are exactly two cosets of H : H itself and the flips. As mentioned above, if a is in H , then $a \circ H = H$. And if b isn't in H , then b is a flip, and the composition of a flip with anything in H is still a flip (which you should check) and so $b \circ H$ consists of the flips.

The last two examples suggest an important fact about cosets, which turns out to always be true: two cosets $a \circ H$ and $b \circ H$ are the same exactly when b is in $a \circ H$ (or the other way around, of course). We can check this directly from the definition of a coset. One direction is easy: if $a \circ H = b \circ H$ then, since $a = a \circ i$ is in $a \circ H$, we get that it's also in $b \circ H$. For the other direction, if b is in $a \circ H$, then $b = a \circ k$ for some k in H . But then anything of the form $b \circ h$ where h is in H can be written as $a \circ k \circ h$, and $k \circ h$ is in H , so this element is actually in the coset $a \circ H$.

This argument shows that $b \circ H$ is contained in $a \circ H$. Since there was nothing special about the elements a and b , we can run the exact same argument but with a and b switched to get that $a \circ H$ is contained in $b \circ H$, and these two facts together let us claim that the sets are equal.

This lets us show that, when G is finite, the cosets of H partition G into equally sized subsets, that is, G is cut up into equally-sized, non-overlapping pieces, each of which is a coset of H . First, every element of G is in a coset: its own. Two different cosets can't overlap: if $a \circ H$ and $b \circ H$ overlap, then they contain some common element, say c . But by what we just showed, this means that $c \circ H = a \circ H$ and $c \circ H = b \circ H$, so our two original cosets were actually the same.

So it just remains to show that all cosets have the same size. In fact, all cosets have the same number of elements as H . Given some coset $a \circ H$, we can pair off the elements of H with the elements of $a \circ H$ by putting an element h in H with the element $a \circ h$ in the coset. This is a one-to-one correspondence: everything in $a \circ H$ is hit by definition, and if $a \circ h = a \circ h'$, then $h = h'$ by cancellation.

So, since we have cut G up into equally-sized pieces and all those pieces have the same size as H , we see that the size of H has to divide the size of G . By the same reasoning, we also get that the number of cosets is $|G|/|H|$, where $|S|$ is the number of elements in some set S . We can see that this is true in both of the finite examples above: in the cyclic group example, there are $\frac{n}{m}$ elements in $m\mathbb{Z}_n$, and there are n elements in the subgroup of D_n consisting of rotations and the identity.

Using Exercise 14 of Section 2, Lagrange's Theorem has an immediate and very useful consequence: since the order of an element a in G is the same as the size of the subgroup $\langle a \rangle$, we also get that the order of a divides the size of G . This will be useful in the number-theoretic application that follows.

3.2 Fermat's Little Theorem

As an application of these ideas, we prove a statement in number theory called *Fermat's Little Theorem*. For this section, we assume some familiarity with a few basic notions from number theory, which we'll go over very briefly here without proving anything.

Pick some positive integer m which we'll keep fixed for this whole discussion. Given two integers a and b , we say that they are *congruent modulo m* , and write $a \equiv b \pmod{m}$, if $a - b$ is a multiple of m . The key fact about this relationship is that it's nice to addition and multiplication. That is, if $a \equiv a' \pmod{m}$ and $b \equiv b' \pmod{m}$, then $a + b \equiv a' + b' \pmod{m}$ and $ab \equiv a'b' \pmod{m}$. (If you're never done this before, it's worth working out why these facts are true on your own. For the second, one method is to replace a with a' , then b with b' , in two different steps.)

The other fact that we need to invoke is sometimes called the Division Algorithm: given any integer a , there is a unique way to write $a = qm + r$ where r is between 0 and $m - 1$. We call q the *quotient* and r the *remainder*. From the division algorithm, it's possible to prove the following useful fact: two integers a and b are relatively prime (that is, they have no common factors) if and only if it is possible to write $1 = as + bt$ for some integers s and t .

This idea of congruence gives us a more sophisticated way to think about the cyclic group \mathbb{Z}_m . We can divide the integers into m different sets, called *equivalence classes* so that any numbers in the same set are congruent modulo m . For example, if m is 3, then one set will contain the multiples of 3, one will contain the numbers which are 1 more than a multiple of 3 (and therefore congruent to 1 modulo 3), and the last the numbers which are 2 more than a multiple of 3. If m is 2, the classes are just the even numbers and the odd numbers.

Every integer will belong to a unique equivalence class. We can then think of \mathbb{Z}_m as consisting not of numbers but of these equivalence classes: instead of 0, we take the equivalence class that 0 belongs to (that is, the multiples of m); instead of 1, we take the class that 1 belongs to, and so on. Then in order to add equivalence classes, we just take one member of each class and add them, and see what class we end up in. The fact that congruence modulo m is nice to addition means that it doesn't matter which elements of each class we pick.

We can use this idea to try to define a multiplicative version of \mathbb{Z}_m , which we'll call \mathbb{Z}_m^\times . Since congruence modulo m is also nice to multiplication, we won't run into any problems defining the composition law in \mathbb{Z}_m^\times , which we'll call \odot . The only problem we might run into is finding inverses. For any integer a , an inverse for the class of a in \mathbb{Z}_m^\times will be some integer b so that $ab \equiv 1 \pmod{m}$. If a and m have a common factor, say d , then this can never happen: d will also be a factor of ab , but since it's also a factor of m , d will divide any number of the form $ab + km$. In particular, such a number can't be 1.

But if a and m have no common factors, we can use the statement we mentioned above after the Division Algorithm: we can write $1 = as + tm$, and s will be our inverse for a , since $as = 1 - tm \equiv 1 \pmod{m}$. So we can make \mathbb{Z}_m^\times into a group so long as we only include numbers which are relatively prime to m .

Fermat's Little Theorem states that, if p is a prime number and a is any number that isn't a multiple of p , then $a^{p-1} \equiv 1 \pmod{p}$. In the language of group theory, this is a statement about the order of a in the group \mathbb{Z}_p^\times . Now, since p is prime, any number that isn't a multiple of p has no factors in common with p . So of the equivalence classes, the only one we need to exclude is the class of 0, the one which contains the multiples of p . This means that \mathbb{Z}_p^\times has $p - 1$ elements. But then Lagrange's Theorem tells us that the order of a divides $p - 1$, say n is the order of a and $p - 1 = nk$. Then $a^{p-1} \equiv (a^n)^k \equiv 1^k \equiv 1 \pmod{p}$, which proves the theorem.

Exercises

1. If f is a flip in a dihedral group D_n , what is the size of the subgroup $H = \langle f \rangle$? How many cosets does H have? What do they look like?
2. Prove that $a \circ H$ is a subgroup of G if and only if a is in H .
3. Prove that $a \circ H = b \circ H$ if and only if $a \circ b^{-1}$ is in H .
4. If G is a group with n elements, prove that for any element a in G , a^n is the identity.
5. Prove that if a group G has an element of order n , and m divides n , then G has an element of order m .
6. Prove that any group with 12 elements has an element of order 2.
7. Use Exercise 15 from Section 2 and Lagrange's Theorem to show that every group with a prime number of elements is isomorphic to a cyclic group.
8. Use the fact that the order of an element divides the order of a group to show that \mathbb{Z}_m^\times has an even number of elements whenever $n > 2$.
9. For any positive integer m , we write $\phi(m)$ for the number of integers in the range $0, 1, \dots, m - 1$ which are relatively prime to m . Generalize the proof of Fermat's Little Theorem to show that, if a is relatively prime to m , then $a^{\phi(m)} \equiv 1 \pmod{m}$.
10. Use Lagrange's Theorem to show that $n | \phi(a^n - 1)$ for any integers $a > 1, n \geq 1$.

4

Group Actions

4.1 Basic Definitions

In spite of the fact that our aim was to abstract the notion of a group away from the geometric objects they came from, it's often necessary to keep track of what a group “does” to some other object in a systematic way. After all, one of the reasons to study the symmetry group of a square is to learn things about the square. This brings us to the notion of a group action, which is meant to capture this idea.

Just as we did for groups themselves, we can extract the definition of a group action from the properties of our basic examples. The basic idea is that if G is some group and S is a set, we will say that G acts on S if all the elements of G “move around” the elements of S , just as the elements of D_3 move around the vertices of the triangle. The point of the beginning of this section is to make this idea precise.

Consider the example of D_4 acting on the vertices of a square. If x is one of those vertices and g is some element of D_4 , we write $g.x$ for the point that x goes to after applying g . (See Figure 10.)




Figure 10: Some examples depicting $g.x$ for various choices of g , where x is the vertex in the upper-left corner.

When we were deciding on the definition of a group in the last section, we needed to figure out which properties of composition which could be described without referring to the geometry that inspired it. Now we need to do the same for the action of a group: we need to see how to describe a group action using properties that depend only on its interaction with group composition.

The first of these has to do with the identity. Remember how we defined the identity in the first place: it was the symmetry which left everything in the same place. This property doesn't require us to know any geometric facts about the set G is acting on, so it can go on our list: if G is a group acting on a set X , i is the identity element of G , and x is some element of X , then $i.x = x$.

The second relates to composition. Remember that we defined $g \circ h$ as the symmetry which you get by doing h , then g . What this means in terms of the group action is that $g \circ h$ should act on an element x by first applying h — which takes it to $h.x$ — and then applying g to that. In other words, $(g \circ h).x = g.(h.x)$. (This is the reason we defined composition the way we did in Section 1.)

These two properties turn out to be enough:

Definition. We say that a group G acts on a set X if, for every pair of elements g in G and x in X , there is a corresponding element $g.x$ in X with the following properties:

- If i is the identity, then $i.x = x$ for any x .
- If g and h are elements of G , then $(g \circ h).x = g.(h.x)$ for any x .

There is one more natural property of symmetries that we could have included in this list: the way inverses act on elements of X . Given the way we defined inverses in Section 1, it seems

natural to insist that g^{-1} take $g.x$ back to x . The reason we didn't include it in our definition is that it already follows from the two properties we gave: $g^{-1}.(g.x) = (g^{-1} \circ g).x = i.x = x$.

We study group actions both to learn more about the groups and about the objects on which they act. In order to help us do both of these things, we introduce a couple of terms.

If x is some element of X , the *orbit of x* , written $G.x$, is the collection of all elements of X which can be written as $g.x$ for some g in G . The *stabilizer of x* , written $\text{Stab}_G(x)$ (or just $\text{Stab}(x)$ if the group G is clear from context), is the collection of all elements g in G for which $g.x = x$. We say that such elements *fix x* . Given some element g in G , the *invariant set of g* , written X^g , is the collection of elements of X for which $g.x = x$. Even though the definitions are similar, it's important to keep the stabilizer and the invariant set apart: the stabilizer lives in the group and the invariant set lives in the set.

Examples.

1. Any dihedral group, say D_n , acts on the vertices of a regular n -gon. Given any two vertices, there is at least one symmetry of the n -gon which takes the first to the second, so all the vertices are in the same orbit. (When this happens, we say that the action is *transitive*.) The identity fixes everything — this is always true by definition — but none of the rotations fix anything.

Every vertex is fixed by exactly one reflection: the one around the axis that goes through that vertex. So the stabilizer of every point has two elements: the identity, and that one reflection. The invariant set of the identity is everything, of a rotation is empty. If n is odd, then every axis of reflection passes through a vertex, so every reflection has that vertex in its invariant set. (See Figure 11.) If n is even, there are two kinds of reflections: the ones through edges and the ones through vertices. (Also see Figure 11.) In the first case, the reflection has no invariants, and in the second, it has two.




Figure 11: The three types of axes of reflection for elements of D_n

2. Any group acts on itself by setting $g.h = g \circ h$. Again, this action is transitive — that is, there's only one orbit — and this time, all the stabilizers contain just the identity and all the invariant sets are empty, except for the identity, which has everything.
3. The multiplication-mod- n group \mathbb{Z}_n^\times acts on \mathbb{Z}_n by setting $a.b = ab$. If n is 6, for example, then \mathbb{Z}_6^\times consists of just 1 and 5, and $\text{Stab}_{\mathbb{Z}_6^\times}(3)$ is all of \mathbb{Z}_6^\times . Since 3 is the only element fixed by 5, all other stabilizers just contain 1.
4. An *involution* on a set X is a function f from X to itself with the property that $f(f(x)) = x$ for any x . An involution on X is essentially the same as an action of \mathbb{Z}_2 on X : just set $1.x = f(x)$. Since $1 \oplus 1 = 0$, and 0 is the identity, saying this is a group action is the same as saying the function is an involution.
5. Any invertible function f from X to itself gives an action of \mathbb{Z} on X : set $1.x = f(x)$ and $(-1).x = f^{-1}(x)$. Suppose that applying f to any element n times is the same as doing nothing. In this case, for the same reason as in the previous example, this also gives an action of \mathbb{Z}_n on X .

6. As an example of this, consider the function from the set $X = \{a, b, c, d, e\}$ defined as follows:

$$a \mapsto b \quad b \mapsto d \quad c \mapsto e \quad d \mapsto a \quad e \mapsto c$$

The elements $a, b,$ and d form an orbit, and the elements c and e form a separate one. Applying the function three times fixes everything in the first orbit, but not the second. Applying it twice fixes the second, but not the first. This means that applying it six times fixes everything. So this function gives us an action of \mathbb{Z}_6 on X . Everything in the a - b - d orbit is fixed by 0 and 3, and everything in the c - e orbit is fixed by 0, 2, and 4.

4.2 The Orbit-Stabilizer Theorem

Studying orbits and stabilizers is useful because of the following fact, called the Orbit-Stabilizer Theorem: if G is a finite group acting on a set X , and x is an element of X , then $|G \cdot x| = |G|/|\text{Stab}(x)|$. (Recall that $|S|$ denotes the number of elements in S .) First, notice that the right-hand side is the number of cosets of $\text{Stab}(x)$. The way we will prove this statement, then, is by finding a one-to-one correspondence between cosets of $\text{Stab}(x)$ and elements of $G \cdot x$.

Take some coset $g \circ \text{Stab}(x)$. To this coset, we assign the element $g \cdot x$. In order for this to make sense, we need to check that we didn't accidentally assign our coset to more than one element of the orbit: the coset can correspond to more than one element of G . If $g \circ \text{Stab}(x) = h \circ \text{Stab}(x)$, though, we know that $g = h \circ s$ for some s in the stabilizer. So in fact, $g \cdot x = h \cdot (s \cdot x) = h \cdot x$, so our definition indeed makes sense. This argument goes the other way as well: if $g \cdot x = h \cdot x$, then $h^{-1} \circ g$ is in the stabilizer, which means that $g \circ \text{Stab}(x) = h \circ \text{Stab}(x)$ by Exercise 3 from Section 3. This shows that different cosets are assigned different elements of the orbit: if they're assigned the same element of the orbit, they were actually the same coset. Everything in the orbit is hit by definition, so we do have a one-to-one correspondence.

Examples.

1. In Example 1 above — the action of D_n on a regular n -gon — the stabilizer of any point has 2 elements, the orbit (which is everything) has n elements, and the group has $2n$ elements, so the theorem holds.
2. The Orbit-Stabilizer Theorem, when combined with Lagrange's Theorem, says that the size of any orbit has to divide the size of the group. In some cases, this can be used to analyze the orbits. For example, if $|G| = 15$ and $|X| = 7$, then there has to be an element fixed by everything in G , that is, an element which is in an orbit all by itself: if not, then every orbit has 3, 5, or 15 elements, and there is no way to get 7 by adding only these numbers.

4.3 Burnside's Lemma

The orbit-stabilizer theorem, in turn, gives us a great way to count the number of orbits of a group action, called Burnside's Lemma. The lemma says that the number of orbits of the action of a group G on a set X is the sum of the sizes of the all the invariant sets divided by the size of G . That is, there are $\frac{1}{|G|} \sum_{g \in G} |X^g|$ orbits.

We can prove this by counting the sum $\sum_{g \in G} |X^g|$ in two different ways. That sum is the same as the number of pairs (g, x) for which $g \cdot x = x$. While we have divided these pairs into invariant sets by looking at different values of g , we can equally well divide them into stabilizers by looking at different values of x . So this sum is the same as the sum of the sizes of the stabilizers, that is, $\sum_{x \in X} |\text{Stab}(x)|$.

By the Orbit-Stabilizer Theorem, this sum is $\sum_{x \in X} (|G|/|G \cdot x|) = |G| \sum_{x \in X} (1/|G \cdot x|)$, so we just need to find out what that sum is. We can group the elements of x into orbits. Each orbit Y will contribute $1/|Y|$ to this sum for each element, and it will do so $|Y|$ times. That is, each orbit contributes 1 to this sum, so the sum is actually just the number of orbits. Dividing by $|G|$ gives the result.

Examples.

- Suppose we are building necklaces of six beads selected from a set of one of two colors. For some examples, see Figure 12.




Figure 12: Some of the necklaces we're talking about in Example 1.

Now, necklaces B and C are, in some sense, the same: if I turn C over, reflecting it along the line running through the beads on the upper-left and lower-right, I get B. Thinking of pairs these as being the same necklaces, how many “essentially different” necklaces can I get?

There is an action of D_6 on the collection X of all possible necklaces which, for any symmetry a in D_6 , simply applies a to the necklace. Any two necklaces in the same orbit of this action are ones we're considering the same, so what we're asking for is the number of orbits of this action. For this, we need to count the sizes of the invariant sets of each element of D_6 . You should check all of the claims in this argument and convince yourself that they're true.

First, consider the reflections. There are two types of reflections: the ones through edges and the ones through vertices. For the ones through edges, the beads on one side have to be the same as the beads on the other, so there are 3 beads which are free for us to pick, giving 8 necklaces in the invariant set. For the ones through vertices, the same is true, but the axis passes through two of the beads, so there are 4 for us to pick, giving 16 in the invariant set. All together then, the reflections contribute $3 \cdot 8 + 3 \cdot 16 = 72$ to the sum of the sizes of the invariant sets.

The identity fixes every necklace, so its invariant set has $2^6 = 64$ necklaces. Rotations by 60 and 300 degrees only fix the all-white and all-black necklaces. Rotations by 120 and 240 degrees fix a necklace as long as the colors “alternate”: a bead in some position needs to be the same color as the one two spots over. So we can divide the bead positions into pairs of adjacent beads; we're free to pick the colors in one of the pairs, but then the others have to match. This gives 4 necklaces. Similarly, the rotation by 180 degrees leaves us free to pick 3 beads, giving 8 necklaces. So our entire sum is $72 + 64 + 2 + 2 + 4 + 4 + 8 = 156$, giving $156/12 = 13$ essentially different necklaces.

- With necklaces of six beads, it is possible to count by hand. To solve this problem for some other number, you need to be able to divide the rotations into groups based on how many choices you're allowed to make for a necklace which is fixed by that rotation. In general, this depends on the order of that rotation, and it's possible to get rotations of lots of different orders, making the calculation very messy.

If p is some prime number, though, all the rotations have order p : the group of rotations has p elements, and if r is a rotation, then $|\langle r \rangle|$, which is the order of r , divides p , so it has to be p unless r is the identity. This makes it very easy to count the necklaces fixed by r :

it's just the all-black and all-white necklaces. Also, unlike in the case with six beads, all the reflections look more or less the same: they pass through one vertex and one edge, giving us $\frac{p+1}{2}$ beads whose colors we can pick. So the rotations contribute $2(p-1)$ to the sum, the identity contributes 2^p , and the reflections $p \cdot 2^{\frac{p-1}{2}}$. All together, after dividing by $2p$, we have $\frac{1}{p}(p-1 + 2^{p-1} + p \cdot 2^{\frac{p-1}{2}})$ essentially different necklaces.

3. We can use Burnside's Lemma to count the number of distinct ways of painting the faces of a cube with n different colors, where two paintings are "distinct" if there is a way to rotate the cube to take one to another. This is the same as the number of orbits of the action of the rotation group of the cube on the collection of colorings. This example is a sketch; you should fill in the details and check all the computations.

Call the group of rotations G . At the risk of ruining a problem from Section 2, we describe this group and the colorings which are fixed by its elements. There are:


- The identity, which fixes all n^6 colorings of the faces.
- Six 90-degree rotations around axes that go through pairs of opposite faces. A coloring fixed by such a rotation can have any color on the top or bottom, and has to have the same color on all the sides, leaving 3 faces to choose colors for, meaning n^3 colorings.
- Three 180-degree rotations around axes that go through pairs of opposite faces. Each of these has n^4 fixed colorings.
- Six 180-degree rotations around axes that go through pairs of opposite edges. Each of these has n^3 fixed colorings.
- Eight 120-degree rotations around axes that go through pairs of opposite vertices. Each of these has n^2 fixed colorings.

All together, then, there are $\frac{1}{24}(n^6 + 3n^4 + 12n^3 + 8n^2)$ colorings. For $n = 2$, this is 10. For $n = 3$, it's 57.

Exercises

1. Recall the definition of the group S_n of permutations of a set of n elements from Exercise 10 of Section 2. There is an action of S_n on the set $\{1, 2, \dots, n\}$ which comes from the definition: a permutation in S_n is a way of rearranging the numbers because of the way we defined it, so that tells us what it does to each number in the set. Describe the stabilizer of an element of the set. What group does it look like?
2. We say that a group G acts *faithfully* on a set X if no elements of G other than the identity fix everything in X . Suppose G acts faithfully on a set with n elements. Show that G is isomorphic to a subgroup of S_n .
3. Consider the action of the group of rotations of the cube on the collection of pairs of opposite vertices on the cube. Use the previous problem to show that G is isomorphic to S_4 .
4. Show that the group of all symmetries (not just rotations) of a regular tetrahedron is isomorphic to S_4 .
5. Say G has n elements, and p is the smallest prime number dividing n . If $k < p$, describe all possible actions of G on a set of k elements.
6. Say p is a prime number, and the number of elements in G is a power of p , say p^n . If G acts on some set X , and there aren't any points which are fixed by everything in G , show that p divides the size of X .

7. Fill in the details in the example about coloring faces of the cube.
8. Suppose we make $n \times n$ tiles by taking a grid and coloring each square black or white. For example, these are possible tiles when $n = 3$:



We regard two tiles as the same if one can be moved around — either rotated or reflected — to get the other. For example, the tiles in each pair above are the same. How many different possible tiles are there?