[SQUEAKING] [RUSTLING] [CLICKING]

**ANKUR MOITRA:** All right, let's get started. So today, we're going to continue our discussion of algebra. So last time, we talked about modular arithmetic. And I introduced you to the basic notions. We talked about the Euclidean algorithm, the extended Euclidean algorithm, and multiplicative inverses.

And it turns out that there are some really powerful abstractions that can give us a useful way to think about modular arithmetic. And these come from really the starting point of group theory. So group theory is something that you can take a whole class in, multiple classes in. And it goes quite deep. But we're going to introduce you to the basics of group theory and why this abstraction is powerful.

So let me start off from where we left off last time. So we had this notion of multiplication in modular arithmetic. And one of the questions which was on our mind was whether an element has a multiplicative inverse. So given an element a, is there some other number element we can multiply it by so that mod n we get back 1? And we showed that this was always true whenever n is prime and a is nonzero.

And in fact, what we really showed was that even if n is not prime, the critical thing is whether or not the GCD, the greatest common divisor, between a and n is 1. So we showed in particular that there was this equivalent way to think about the GCD as the smallest integer you can get, non-zero integer, where you take s times a plus t times n. And you want to get that to be equal to 1. So if the GCD is 1, then that gives you the multiplicative inverse. And if you can't get 1, you can show directly that it has no multiplicative inverse.

So this is where we left off last time. But now, I'm going to tell you one very powerful definition. And it'll take us all of today's lecture to really understand why this is such a powerful notion, which takes us to what's called a group. So a group is just a pair, g. It's a set of elements, potentially infinite, along with a binary operation star. So this is some set, potentially infinite. And this is just the binary operation.

And the critical thing about a group is that it has to satisfy a whole bunch of properties. So the first thing we need is closure, that we can apply this binary operation to pairs of elements and get back another element in the group. So this is really just the statement that, for any pair of elements a, b that belong to the group, the operation of a star b is well-defined because it sends us back to some group element.

So last time when I wrote down, for example, this table we got with modular addition, I was really just showing you, for every row that's a choice of a different element, for every column that's a choice of a different element. And I told you in the table what particular group element that pair adds to or multiplies to. That was just the closure property.

We also need associativity, which is also familiar from before. So for all triples a, b, c that belong to your group, doing the operation of a star b, and then taking the result and doing that with star c is the same thing as parenthesizing the expression a different way.

We also need the existence of an identity element, which means that there is some element, which we'll call e, that belongs to the group with the property that, for all elements of the group, doing a star e gets back the same element I started from. And that doesn't matter whether you apply e on the right or on the left. Either way, you get back your original element. And you do nothing to it.

So for example, in the case of the group being-- when we were working with this set of moduli, when we had addition, the identity element is 0. And when we had multiplication, the identity element was 1. And finally, the last property we need from our group is the existence of an inverse. So what this means is natural. We want that, for all a belonging to our group, there is some element, which we'll call a inverse, with the property that a inverse applied on the left versus a inverse applied on the right, both of them send you back to the identity element.

So this is the notion of a group. This is just formalizing algebraic structures that we've already seen in action. We've already seen groups. But we didn't think about them as abstract sets, abstract properties. And I'll also mention that, additionally, if we have commutativity-- so if commutativity holds that a star b is equal to b star a, then what we say is that we call g abelian. So it's an abelian group.

So these are the key definitions for today. And what we're going to do is we're going to understand why this is powerful. So are there any questions about the definition of a group? This makes sense, hopefully. So next, let's go through some examples. We've already seen groups in action. And let's see some examples just to get a feel for what these definitions are.

So one simple example is an infinite group you've been working with all the time. You can look at the set of integers under addition. So we can check, in that case, our set g is just an infinite set. It's the set of all integers, positive and negative. Our binary operation is addition. And you can check that all of these properties hold. I take two integers, I add them up, I definitely get back an integer.

When I add up three integers, it doesn't matter the order that I add them up. I have an identity element that's 0. I have an inverse. That's just the negative of the number I started off with. And in fact, it's also commutative. So it's an abelian group. So this is a very simple example of a group is just the set of integers under addition. Just as a pop quiz, what about the set of integers under multiplication? So what if I take the same set, but my binary operation is now multiplication? Is that a group? Why not?

**STUDENT:** Because 0 doesn't have an inverse.

**ANKUR MOITRA:** 0 doesn't have an inverse. And that's true of all kinds of other elements, too. When I take the element 2, well, I can talk about multiplying by that element. That will still send me back to another integer. But I don't have the existence of an inverse because that would be 1/2. And that's not in my set. So a non-example of a group is integers under multiplication. So it's not just about the set, but it's about how the set behaves under the binary operation I've chosen.

We can do another example with multiplication. We can take the set of rationals minus 0. And we can take our operation to be multiplication. So now, we do have an inverse because we just flip the numerator and denominator. And we, once again, have a group. So these are all examples of infinite groups. And what we'll be particularly interested in today are finite groups. That's the part that's going to connect with modular arithmetic.

So what about finite groups? In fact, I claim we've already seen some examples of finite groups in action. So we can take another example. This will be a very important example for us, which is we can take zn, which is just addition. We can take it with the binary operation of addition, modular addition.

So for example, when n equals 2, this is just two elements. We're looking at evens and odds. And for that finite group, just the set of evens and odds, you can check that when you add up two elements it makes sense. An even plus and even takes you back to even. And odd plus an odd takes you to even. And an even plus an odd takes you to odd.

So in fact, we've already seen examples of finite groups in action. That's what we did last time. But let me give you a more interesting example, which is going to show up. So actually, let's do a simple corollary first. And then we'll get to that example. This will be a corollary of that key lemma I wrote up there.

So let's look at the set zn minus 0. So I'm just looking at the remainders when I take mod n. But I'm removing the 0 element. And if I look at this finite set with multiplication, I claim that this is a group if and only if n is prime. And I'm writing this as a corollary because, really, it's a corollary of this key lemma that we showed last time.

So let's prove this lemma just to get some familiarity with groups. So this is an if and only if statement. Let's prove the reverse direction first. So let's suppose that n is prime. Then why exactly is this a group? Now, some properties, a bunch of properties of the group are tedious to check. And they're very natural.

So if I take multiplication mod n, it's definitely closed. That's what it means to do modular multiplication. We talked about that last time. You just multiply the two numbers and subtract off as many multiples of n as you can. And you'll get back to some remainder. No big deal. It's associative because, really, multiplication is associative. And we're just subtracting off multiples of n.

The main question and the identity element is easy. That's also one. I take any remainder, multiply it by 1, I still get the same remainder. The tricky thing is really the inverse. That's where this corollary becomes interesting. So the key thing to check is that there is an inverse. And really, from our lemma, so if n is prime, we know, by the key lemma from last time, that every element has a multiplicative inverse.

So that proves one direction of our corollary. The converse is equally simple. So let's prove that if n is-- we'll actually prove it in the contrapositive. So let's suppose that n is composite. So what I just proved was that if n is prime, then it's a group. But let's prove that if n is not prime, if it's a composite, that we don't have a group.

And how am I going to prove that this is not a group? What I'm going to do is I'm going to construct an element that doesn't have a multiplicative inverse because that's where all this action is happening in the proof of this corollary. So any guesses for what kind of element might not have a multiplicative inverse? Any good ideas? We already know this lemma tells us a bunch of things do have multiplicative inverses. So where should I look for something that doesn't have a multiplicative inverse? I should look for an element a where the GCD between a and n is not 1.

So let's just take any old element a with the property that the GCD between a and n is not equal to 1. And so the Lemma tells us already-- the work is already done for us-- that that element does not have a multiplicative inverse. So it fails our property for being a group.

So when we were doing modular arithmetic and we were worrying about multiplicative inverses, really, what we were doing was we were characterizing when exactly we had this group structure. So that's the first point, is that we were really doing group theory last time. And what we're going to do now is we're going to understand some of the implications of having this abstraction.

All right. Actually, let me make one comment first. So just as a quick comment, this will be very useful for later. So imagine that we take n, which is composite. We now know that the non-zero elements of zn do not form a group. But it turns out that even though it's not a group, there is a group hiding inside. So this will be very useful for us. So there is a way to fix the failure of this group property.

What we can do is we can let zn star just be the set of elements a belonging to zn with the property that the GCD of a and n is equal to 1. So what I claim-- and I'll just assert this. But the proof is very simple. It just, again, follows from this key lemma last time.

What I claim is that zn star, equipped with the multiplication operation, is a group. So that's our key fact that we're going to be using. So even though there are elements belonging to zn that are not zero that do not have a multiplicative inverse-- so we failed to be a group-- we can just take the set of elements that do have an inverse, the set of all elements a where the GCD is 1. And what I claim is that that's a group sitting inside this other thing that failed to be a group.

So the main point for today's lecture is we're going to develop some understanding of properties of groups. We're going to prove something called Lagrange's theorem, which is going to have a ton of cool, simple to understand applications. So let's go a little further with group theory.

Let me give you a key definition. The order of a group of a finite group G is just the number of elements. And we denote it by size of G, this way. So we're just measuring the number of elements in G. And you shouldn't yet see why this is are a useful notion. But that's where Lagrange's theorem is going to kick in. But let's do some examples just to make sure that we're on the same page.

So we can look at, zn with multiplication for prime n. So we already proved in our corollary that when n is prime, this is a group. So what's the order of this group? Well, the order of this group zn-- sorry, zn minus 0, because the element 0 doesn't work, well, this is just n minus 1 for this group, just because I've removed that element.

So now, the key property we're going to-- so then let's do one more example. And this will be a crucial example that we use for later. So let's imagine that m is composite. We can look at zm star again with multiplication. And let's assume that m is equal to p times q, where both p and q are primes and they're distinct primes.

So this is one simple example. m is composite. So zm minus 0 is not a group. But if we only look at the set of elements that have a multiplicative inverse, we know that it's a group. The question is, what's the order of this group? So how many elements are actually in this set zm star? So this is really just a counting argument because what we want to do is we want to count the number of a's such that the GCD between a and m is equal to 1.

So now, what I can do is I can use my old friend, inclusion exclusion. So what I can do is I can overcount. So I can start off with pq minus 1 because those are all of the elements that are in consideration in this set. And then what I can do is I can subtract off the things which fail to have GCD 1.

So what are the ways that a could fail to have GCD with m 1? Well, it could have GCD equal to p or it could have GCD equal to q. So the question is just, how many elements, how many choices of are there where the GCD between a and m would be p? How many are there where the GCD would be q? Because that's the amount by which I've overcounted. So I should subtract off those things.

And you can subtract off p minus 1. This is just the number of a's where the GCD of a and m is equal to q. In particular, a should be a multiple of q in order for this to be true. And then there are p minus 1 possibilities for what that is. And I can do the same thing the other direction. I can subtract off the things which are multiples of p. And I'm going to get q minus 1 of these elements. And there's nothing that's a multiple of both p and q because the largest my remainder gets to is pq minus 1.

So this, even though it sounds like a pretty tedious example, this turns out to be an extremely important example. So just to write down what the expression is, it's p minus 1 times q minus 1. And what I promise you is that this expression is going to come up in a very magical way later. In fact, I claim that this expression is the key to RSA encryption.

So we'll see a bit of a hint of that. But this expression is very important. It's just the order of this group that has fixed the fact that our original attempt was not a group. So are there any questions about where we are right now? The notion of groups makes sense and the order makes sense hopefully? Any questions? Good? All right.

So I want to get up to Lagrange's theorem. But next, we have one more abstraction before I can tell you Lagrange's theorem. This is our last abstraction for today. We can also talk about the notion of a subgroup. So imagine I gave you a pair, G and star. So G is the set. And we'll be interested in finite sets for today. And then we have this binary operation, star. And let's assume that it's a group. So it satisfies all the axioms 1 through 4 of being a group.

And let's imagine that you give me some subset of the group. You just keep some of the elements. And let's call that H. Well, if this pair H and star is also a group, then what we do is we say that H is a subgroup. So this is our key definition. So not only do we have a group, but sometimes inside a group we have other smaller groups sitting inside them.

And what we're going to do is we're going to show that whenever we have a subgroup, we can say something very interesting about the relationship between H and G. So that's where we're headed. But let's just be explicit about this. So what does it mean for H star to be a group?

So in particular, what this definition is really telling us is that we better satisfy a stronger closure property. So when we have any pair of elements a, b that belong to H, we need that a star b belongs to h as well. See, a and b are both elements of G. So we're automatically guaranteed that a star b belongs to G. But this is the stronger statement that when a and b are actually not just in G, but they're sitting inside H, then the result of multiplying them should be something that's in H, too. So the critical thing here is that this should not just be in G, but in H itself.

And the same thing is true for the inverse, too. We need that, for all elements a belong to H-- well, a is in G. So there is an inverse. But the key thing is that this inverse actually belongs to H and, again, not G. So somehow, when you're doing the binary operations or when you're looking for the inverse, you never have to go outside of the set you started from, even though it lives inside a much bigger group. So that's the idea. So this is a key definition. This is very subtle. Are there any questions about the definition for a subgroup?

All right. So I'm going to state Lagrange's theorem. And then we're going to go through some examples. And we'll prove Lagrange's theorem. And then I'll tell you some awesome applications of it. So what I claim relating all these concepts we've done so far-- so if G is finite, it's a finite group, and H is a subgroup, then what I claim is that the order of H divides the order of G. This is why we care about orders already.

See, when you tell me the group and you tell me the order, how many elements belong to this group, well, that imposes some restrictions on what kinds of subgroups you can find. If I have my original group as order 6, this tells me, no matter where I look, I cannot find a subgroup of order 4 because it wouldn't divide it. So this is a really beautiful fact.

Even though it sounds like we haven't assumed very much, just G satisfies these simple axioms and H is a subgroup, automatically we can say something very striking and powerful about the relationship between H and G that tells us something very important structure on what H could be. So we're going to prove this. But first, let's give some examples just to make this explicit.

All right. So let's let G denote z7 negative 0. 7 is prime. So this is a group. It has six elements. And now, what's going on is that the subgroup H we can take to be the elements 1, 2, 4. And what you can do is you can just check that when I take any of these two elements-- 1 times 2, 1 times 4-- I go back within the subgroup. And 2 times 4, I go back to the subgroup because it gives me 1.

And you can check that every element has a multiplicative inverse because, in fact, the multiplicative inverse of 4 is 2. And similarly, the multiplicative inverse of 2 is 4. In fact, one of the elegant ways to see that this is a subgroup is actually by realizing that there's a way to express this subgroup in a different way.

So an easy and powerful way to see this that will be very useful for what we do later, in terms of applications of Lagrange, is that all of these elements of the subgroup are all powers of 2. So 2 to the 1 is 2 mod 7. 2 to the 3 is 1 mod 7. Sorry 2 to the 2 is 4 mod 7. So all of the elements we got-- 1, 2, and 4-- they were powers of 2.

So why do we have the closure properties? It's just because when I take different powers of 2 and multiply them, I'm just adding their exponents. So really, this is the set of remainders you get from taking powers of 2 and looking at the remainder mod 7. So this is one example of how to build a subgroup, is really with powers of some element. And this will be a crucial way that we use Lagrange's theorem. But I'll have to explain that later.

And of course, you can check the obvious thing, that Lagrange's theorem checks out in this case because the size of G is 6, the size of H is 3. So 3 really does divide 6. All right. So let me pause here and ask you guys if you have any questions again because, if not, we're going to do the proof of Lagrange's theorem. And that's where we're going to assemble all of these definitions from groups and subgroups that we've been talking about so far. So ask me questions to slow me down. They should hopefully be about the course material. But if not, that's fine.

I'm just going to stare awkwardly until someone asks me a question. Everyone understands everything. You promise? Yes, OK. All right. I'll take it. So let's prove Lagrange's theorem. This will be the main proof we do today. But then we'll get a whole bunch of amazing applications out of it and this fact that I promised you is going to come back later on. But we can already see some hint about why this fact is powerful. Counting the number of group elements tells us something very powerful about the types of subgroups we can expect to see inside.

All right. So let's prove it here. I think I should have enough space. Now, this is really just a counting argument using the group structure. After all, we just want to prove some statement about divisibility. So really, what's going on here, the way to prove Lagrange's theorem at a high level is to break G up into different pieces.

We already have a subgroup, which is one piece of G. But the question is, how can I translate that subgroup around to cover all of G in a nice, even way. So that's the intuition behind Lagrange's proof. And making this precise, we have to talk about what are these translates. And this takes us to the notion of a coset.

So as I promised you, the main idea is to partition G. And we're going to build up this partition one by one. And we're going to partition it into something called cosets. These are really just translates of my original subgroup. So if you take any element belonging to the group, what we're going to do is we're going to consider the following set. It's called xH. And it's really the set of all x, y's where y belongs to H.

So what's going on is, for the definition of the subgroup, we needed the property that when I take a pair of elements belonging to H. We go back within H. Here, I'm taking a pair of elements, one of which belongs to H and one of which is general. And these are my different translates. So I'm not just considering an x which belongs to H. But as I range x over the group, I'm going to get different pieces of the group. When x belongs to H, I'm just getting back H. And when x is something else that does not belong to H, it's a translate of the thing I started from.

So these are called cosets. In fact, this is called a left coset because the element that I'm taking from G is multiplied on the left. And right cosets would be the same thing, but on the other side. So now, here's our plan. This is our plan of action, is what we want to do is we want to show that G can be written as x1 H union disjoint union x2 H all the way up to disjoint union xk H.

So what's going on is I have k different cosets, one for each of these xi's. I haven't told you how to find them. We'll talk about that in a minute. But I want the property that I can choose x1 through xk so that their resulting cosets together make up G and there is no overlap between them. So what the symbol means is it's not the union, it's the disjoint union. So I want all of these things to have empty intersection.

So that's my plan. And if this were true, Lagrange's theorem would be done because what is-- each of these elements right here, we're going to prove that each of these cosets has size order of H. And then we've written G as k copies of H. So that'll be our plan. So here are the things we need to show. And I haven't even told you how to build up this set.

What I want to show is that the size of each of these cosets xH is equal to the size of H. And I want to show that the way that I've built it up, the intersection between any pair of elements of any pair of cosets is 0. And this should be true for all i not equal to j. So in fact, we're going to have a very simple plan for how to generate these xi's. All we're going to do is we're just greedily picking uncovered xi.

So I start off before I've built up a single coset. I have nothing. I've covered none of the elements of G. I take any arbitrary element x belonging to G that's not covered, which is all of them. That will be my x1. That defines a coset. I pull that coset out of G. All of those elements are now covered. And then I take any other element that's not yet covered. That's my x2. It defines a coset. I pull it out. And I keep recursing in this way.

But I have to prove that this really does create a disjoint union of G. Everything is covered at the end. And it's covered once and only once. And I have to prove that each time I pull out a piece, a coset, it has exactly the same size. So if we can prove these two properties, we're done. And we've proved Lagrange's theorem. But now, proving these two properties will just be about using the axioms of group theory. So this is where group theory is a powerful definition. We haven't seen it yet, but we will in a minute, is to prove these two properties.

So let's show these two properties. So first, let's show 1. And let's suppose that xh1 equals xh2 for some h1 and h2 belonging to my subgroup, capital H. So what is the way that I wouldn't get the coset has size order H? Really, the way that the coset is defined is once I've fixed x, I look at x star y for all y in H.

And if for different choices of y, I get the same element when I do that multiplication, I get a collision. That's the only way that I don't have an equality here. If that doesn't happen, then I really do have that my coset equals the order of the subgroup. Every distinct choice of y is going to lead to a distinct product.

So let's imagine that we had two elements h1 and h2 that x times h1 and x times h2 were the same thing. And now, what should I do? I claim that the way that my proof strategy wants to work-- this won't be a very complicated proof, but the logic is subtle.

So what I want to do is I want to show that h1 and h2 must be the same thing. I want to show that the only way that you get the same element when you get x times h1 is equal to x times h2 is if h1 equals h2. So how should I do that? How should I conclude that h1 and h2 are the same thing? And you know my hint is that the critical thing in Lagrange's theorem is if G is a group. So what should I use about the group axioms in order to conclude this fact? That's my key question. You promised me you guys understood everything up until that point. So which axiom should I use? I've got you guys on the hook now. Who can make me happy? People are not making eye contact. Yeah?

**STUDENT:** Yeah, maybe you could just use the existence of x inverse.

**ANKUR MOITRA:** Yeah, that's right. That's a great idea. Let's use the existence of x inverse. And that'll work. So now, what we're going to do is we're going to left multiply by x inverse. See, this is where I need the group structure, is that x inverse exists. So once x inverse exists, then what this is going to imply is that x inverse x h1-- and now, I can use associativity-- is equal to x inverse x h2.

And when I use associativity and the property of the inverse, that it equals the identity, and that identity times an element is the element, it's going to imply that h1 equals h2. So in fact, the critical thing was the inverse. But we really used all of the properties, like the fact that we have associativity, the fact that we get back the identity element, and the fact that the identity element doesn't change any of those things.

So in this one step, we've used all the properties of groups, which is great. And it implies the thing we're after. But this finishes this first part in the proof. Every coset has the same size because when I take different elements of H, I cannot get the same thing because of the inverse. So we're halfway home.

So let's show property 2. And then we'll be done with Lagrange's theorem. So now, let's show 2. So what I'm going to do is I'm going to suppose that they're not disjoint, these cosets we get. So let's suppose, for i less than j, we have the property that xi h1 equals xj h2. And this is true for some h1 and h2 belonging to H.

So we just proved that we don't have any collisions within the coset. We don't get back the same element. Now, what I want to do is I want to suppose, for the sake of contradiction, that there's some coset i and there's some coset j that intersect. And that means that there's some element I got out of one, which is xi times h1. And that's the same as the element I got out of the later coset, which is xj times h2. So this is what it means for the two cosets to not be disjoint, is that they have a collision for some h1 and h2.

And now, we're going to try and get a contradiction. So what I want to do-- I'm going to ask you guys to help me out-- is that I want to show that xj actually belonged to an earlier coset. So I want to show that xj belongs to an coset. This would let me be done because what is the way that I built up my cosets? I started off with any x that was uncovered. And I chose it as my x1. And I pulled out the entire coset. And then I took any x2 that was still not covered. I used it to find a coset and pulled it out.

So if some xj later down the line actually belongs to an earlier coset, then that means I shouldn't have done that. It should have been pulled out before. So what I'm really saying logically is that, if this weren't true, if the intersection were non-zero, then I actually haven't implemented this procedure I told you correctly because xj was already covered.

So now, same question-- I'm telling you what I want to show. This is what we're assuming for the purpose of contradiction. How can I show that xj belongs to an earlier coset? And I'll give you a hint. It should be the coset defined by xi. And same thing is going to be true here. We just need to figure out what property of a group or subgroup we should apply.

So what should I do to this expression, xi h1 equals xj h2 to try and derive an expression where xj belongs to a coset, the coset of Xi? What about right-multiplying by h2 inverse? So let's see what happens. So let's right-multiply by h2 inverse. So we get xi h1 h2 inverse as equal to x. I wanted to show that x belongs to an earlier coset. Why is this expression showing that xj belongs to an earlier coset? So what can I say about this element right here, h1 h2 inverse? Yeah?

**STUDENT:** That it's an element in H.

**ANKUR MOITRA:** It's an element in H, right? Why? Because h2 inverse belongs to H. And then I'm multiplying two things belonging to H. So we've used the two key closure properties of the subgroup. And so what this means is that xj belongs to the set xiH. And that means that xj should have been covered. And we shouldn't have considered it to begin with.

So that's our proof. This is not a long proof. But the logic for this is very subtle. We used everything that we could have used about groups and subgroups just to prove these few lines of this logical argument. Once we had the strategy of how to partition it into cosets and we had this definition of a coset, where we needed the group theory and the subgroups was really the structural properties that forced these properties 1 and 2 to be true.

So this is a tricky proof. I will ask again, any questions? Is everyone comfortable with this proof, so I can give you my pop quiz now that will count towards 50% of your grade? I'm just kidding. Boy, would the people who didn't show up today be upset. We should do that at some point.

All right, so now that we've done this hard work of proving Lagrange's theorem, let's see some payoff. So let me tell you some awesome statements that are going to follow as an application of Lagrange's theorem. And it's going to connect all of the tools we talked about-- subgroups, powers of elements, the order, and everything.

So this is a theorem called Fermat's little theorem. We're not going to prove Fermat's last theorem. So what I claim is that if p is prime and a is not divisible by p, then a to the p minus 1 is congruent to 1 mod p. This is a beautiful identity.

In particular, what this means is really that a to the p minus 2 times a is equal to 1. So this gives me an explicit algebraic expression for what the multiplicative inverse is. Last time, we talked about how we could compute the multiplicative inverse by using the extended Euclidean algorithm. It turns out you can also do it by powering.

If you give me a and I know that p is prime, all I have to do is raise a to the p minus 2 power. And that is the multiplicative inverse. You can check that all of our examples satisfy that property. That's why 2 and 4 were multiplicative inverses of each other. It was just that you take one and you raise it to the p minus 2 power. And that will give you the other one.

So this property, it doesn't seem like it's connected to Lagrange, but it is. So it turns out that we're going to prove it through Lagrange's theorem. So what I'm going to do is I'm going to have a key definition, which is I'll let k be the smallest integer such that a to the k is congruent to 1 mod p.

So now, the basic idea is that I have to prove that such a k exists. It's not obvious that there should always be a k. And once I know that there is a k, I'll relate it to the order of the group because this is going to define a subgroup. And that'll tell me something about when I use Lagrange's theorem and I relate that to the order of the actual group, which is p minus 1. That'll tell me that the same thing holds with p minus 1 instead. So really, this is the plan. First, I have to show that such a k exists. So why is there any k?

And now, we get to use a lot of the tools we've talked about in the class so far. So we get to connect it to a lot of earlier tools. I claim we can prove this by pigeonhole principle. So let's assume that there is no k. Suppose there is no such k. Then what I claim is that, by the pigeonhole principle, there must exist some i less than j with the property that a to the i is congruent to a to the j mod p.

So what I'm going to do is I start off with a. I raise it to different powers. And I stop at the first power where what I get is congruent to 1 mod p. Let's suppose I never stop. What if I just kept going? Well, there are only so many other values for what the remainder when I divide by p can be. And eventually, when I consider high enough powers, I must be getting some kind of collision.

So once I get this collision that a to the i equals a to the j, let's figure out why there must exist a k such that a to the k is congruent to 1. So what we can do is we can just divide through by a because there is an inverse to a, because p is prime. And so we can write this as a to the j minus i is congruent to 1 mod p. This is just because a inverse exists and we can get rid of powers of a.

So when we get rid of i powers of from both sides, we get the identity a to the j minus i is congruent to 1 mod p. And that shows that there must exist some k just by pigeonhole principle. So this is the first thing that I owed you, is that k exists. And now, we're going to relate k to p minus 1. And that'll be the proof of Fermat.

So this is the first step. And now, what I claim-- the powers of a form a subgroup. And what is the order of this subgroup? It's k because that's the first time when I got back the identity element. So it was the same way when I did that numerical example up here. I had a subgroup that was really generated by powers of 2. Now, I have a subgroup that's generated by powers of a. And my order is just the first time that I return back to the identity element. That's what k is.

So I have a subgroup, which is the powers of a. And it has order k. So what I can do is now I can appeal to Lagrange's theorem. So by Lagrange's theorem, we know that k, the order of my subgroup, must divide p minus 1. You remember when we did all those computations of what the order of a subgroup is when n is prime, when n is the product of two primes? We know the order is p minus 1. Because of that first example we did. And now, we can appeal to Lagrange's theorem. And we know that k divides p minus 1.

So now, we're golden because a to the p minus 1, that's the thing we want to show is congruent to 1 mod p. And what are we going to do? We're just going to write this as a to the k raised to the c-th power if, for example, p minus 1 equals c times k because k divides p minus 1. So I'm just rewriting it in a different way. And now, what do I know about a to the k? What is it congruent to, mod p? It's congruent to 1. So this is the same thing as raising 1 to the power c, which is the same thing as 1. So that's the proof.

So this is Fermat's little theorem. And even though the statement of it had nothing to do with groups-- I could have stated this for you at the beginning of lecture today because we knew what modular arithmetic was. We know what it means to take powers of a. But somehow, the cleanest proof of Fermat's theorem-- or at least, one of the cleanest proofs-- really goes through the fact that, sitting inside the powers of a, is some understanding of a group structure and how that relates to the larger group that it contains. So are there any questions about Lagrange's theorem? Sorry, the application of Lagrange's theorem to Fermat's little theorem? Who's comfortable with the proof? Give me a thumbs up. OK. All right.

So one of the really cool things you can do-- let me just take a brief digression for a minute, which is that something like Fermat's little theorem, there are actually tons of proofs. So we gave a group theoretic proof right now. Let me give you a totally different proof just because. And this will connect back to some of the earlier topics we've seen.

I claim there's a purely combinatorial way to prove Fermat's little theorem. And we'll prove it in a particular case for a particular setting of a and p. So let me give you a combinatorial proof. And let's do the proof specifically in the case where a equals 2 and p equals 5, just to be concrete.

So the way that I'm going to do this is I'm actually going to think about a counting problem under the hood. And it'll involve something called necklaces. So the idea is that, when I take the statement of Fermat, it's the statement that a to the p minus 1 is congruent to 1. But another way to think about this statement is really the statement that a to the p minus a is divisible by p.

That's really what the statement means because I can take this identity and I can multiply it through by a. So I get a to the p minus 1 is congruent to a mod p. When I move the a on the other side, I get this quantity a to the p minus a. And the statement that that's congruent to 0 mod p is really the statement that that is a multiple of p.

And I claim that there's a combinatorial proof of this fact, which will give us another way to think about what's happening in Fermat's little theorem. So what I'm going to do is I'm going to count certain combinatorial objects. And I'm going to prove as a corollary that statement. So what I can do is I can think about strings with different types of beads. a is 2. So they'll have two kinds of beads, red or blue beads. And their length will be p, which is 5. So I can look at the number of strings of length 5 over an alphabet-- those are just the bead types-- of size 2.

So how many strings are there of length 5 that are composed of red and blue beads? That's very simple. It's 2 to the 5. That's my a to the p, right here, is 2 to the 5, is just counting the number of these strings. But what I'm going to do instead of counting these strings is I'll be interested in the number of necklaces.

So I'll take those strings, which are length 5, and I'm going to tie their ends together. But now, the point is that it doesn't matter for the necklace whether it's-- when you take that necklace and rotate it, it's still the same necklace. So now, I'm losing some information. There are fewer necklaces because some of them are the same thing. So the question is, how many necklaces are there? I don't know. It was very easy to count the number of strings I can get out of two beads of length 5.

And the way to think about it is that, when you take an example, let's say, which is a necklace that looks like red, blue, red, red, red, well, there are actually many strings that could have come from. In fact, there are five strings it could have come from, which all depend on where the B is. Maybe the B was the starting point of the necklace. Maybe it was the second, and so on. And there are five possibilities. And 5 is my p.

So in fact, some of these necklaces really come from equivalence classes of strings that could have generated. Now, there are some necklaces which don't come from five different strings. Can anyone give me an example of a necklace that can only come from one of the 32 possibilities?

**STUDENT:** All reds.

**ANKUR MOITRA:** All reds, perfect. There are all of these other examples. Well, there's two of them. We could take all red. And we could take all blue. This, any way you rotate it, you get back the same thing. So what you can show is that these are the only two possibilities that the proof is using the Euclidean algorithm again, amazingly.

But here, the number of necklaces, really, the way to think about them is in different groups. Either it's a necklace which comes from five strings or it's a necklace that comes from one. So this gives us our proof, is because the statement is that if we look at 5, this should divide 2 to the fifth minus 2. So these are the number of non-monochromatic strings.

And if we just remove all of the monochromatic strings-- the all red and all blue-- the rest of the strings have this property that they come in a 5 to 1 relationship with the necklaces. So in fact, the necklaces gives us a way to group the set of strings into groups of 5. So what does this statement look like more generally? This is just the statement that p divides a to the p minus a because you'll have a monochromatic necklaces. And that is exactly what Fermat's little theorem is telling you.

So once we have enough tools in discrete applied math, there's this amazing nexus between these things. This was a simple statement related to modular arithmetic. I can prove it using group theory. I can prove it using combinatorics. And there are maybe half a dozen different proofs of Fermat's little theorem. There are many routes to this theorem. And they all lead to different kinds of abstractions and different mathematics. All right, any questions about this?

All right. So I will tell you one other statement, which I promised you. This is due to Euler. And this is the thing which is going to have amazing applications in cryptography when we get to it in a bunch of lectures time. We'll do that. We'll do a whole unit on it later. This is due to Euler.

So we're in the same setting of Fermat's little theorem, except for the fact that we have a composite. So if m is equal to p times q where p and q are distinct primes and a has GCD between a and m equal to 1, then what I claim is that a to the p minus 1 times q minus 1 is congruent to 1 mod m.

So this takes a minute to digest. So now, I'm working modulo m that is not prime. So we already know that it's not a group. And not every element has a multiplicative inverse. But some elements do. And in particular, what are the elements that have a multiplicative inverse? Those are the a's where the GCD between a and m is 1.

So what I'm saying is that all of these elements that have a multiplicative inverse, a version of Fermat's little theorem holds, even though we're not working over a prime. The difference is that I'm raising it to a different power. It's not p minus 1 anymore, where p is the modulus. It's now p minus 1 times q minus 1. So we've changed the power we have to raise it to in Euler's theorem.

It turns out that this is really a corollary of the proof I already gave you. There's actually nothing that's really specific about the fact that p is prime. We're just using groups. So what can I do? How does this proof change? Where did I use the fact that it was prime?

Well, let k be the smallest integer so that a to the k is congruent to 1 mod m. I have to prove that such an integer exists. And if not, then there'll be some collision where a to the i is congruent to a to the j. And then I can still rearrange things because I'm working with an a whose GCD between a and m is 1. So it has an inverse. So I can pull out i powers of a. And I'll get a to the j minus i is congruent to 1 mod m. That part still worked. It was literally the same proof. So such a k exists.

And now, that's still a subgroup. When I look at the powers of 2, look at the powers of a, that's still a subgroup. I still know its order because that's the first time I have a collision, is that k. And then by Lagrange's theorem, I know that k divides not p minus 1, but the order of the group. So what did I promise you we were going to use was this simple example, this computation right here.

What is the order of the multiplicative subgroup sitting inside zm, which is where the structure is really happening? It's not m minus 1. It's p minus 1 q minus 1. So what do we know is that k divides p minus 1 times q minus 1. So then p minus 1 times q minus 1 is equal to c times k. Same thing happens here. This is p minus 1 times q minus 1, which in writing as a to the k raised to the c-th power. And that's still 1.

So in fact, what's really cool about this-- I could have told you all of this at the beginning of lecture. We knew what modular arithmetic was. I could have stated Fermat's little theorem. I could have stated Euler's theorem. And they might have seemed a bit mysterious. Why is the power I'm raising things to different depending on whether I have a prime or I have a product of primes/

And it turns out that once you have this group theoretic proof of it, we can see that the reason that we have a different power is because it's no longer a group. And we have to look at the order of the subgroup that's inside. So that's really what's going in the exponent, is just the order of the appropriate subgroup.

So this is one of the amazing things about group theory, is that it can demystify a lot of these statements that otherwise look like they're not quite the same thing. In fact, here, we had literally the same proof. It was 100% the same proof. I just had to cross out a bunch of things and replace them. So are there any questions?

All right. So very last thing I'm going to tell you-- one more abstraction, and then I'm just going to tell you a theorem because we're going to use it later. But we'll just take this for granted. So it turns out a group has a decent amount of structure, but not everything. So there's actually something you can do which imposes even more structure than a group. Actually, there are many things you can do. There are things called rings. We won't be interested in those.

But we'll be interested in something called a field, which has a lot more structure than a group. And when you're dealing with a field, a lot more nice things can happen. So last bit of abstract algebra is a field is a set F, potentially infinite, which now has two operations, addition and multiplication. And it satisfies a bunch of properties.

So first of all, we need that F and the operation plus, the binary operation, is an abelian group. And it'll have identity which is some distinguished 0 element. We're going to need that multiplication is also a group. We need that F minus 0, when we remove out the element that's the identity for addition, we need that the binary operation on whatever remains is a group. We need that it's an abelian group. And it'll have identity 1, some special element.

And the very last property that we need is we need that it satisfies the distributive property. So we need that, for all a, b, c belonging to my field, a times b plus c is the same thing as a times b plus a times c. So this is a field. A lot of times, we're not just dealing with groups. We're dealing with fields, like the field of rationals, where we have addition and multiplication and everything behaves nicely.

And now, here is an amazing theorem, which we're going to use later on when we do more crypto. What I claim is that any polynomial of degree d, at least one, with coefficients in my field can have at most d roots in F. So what this means is you can consider some polynomial, p of x, which is p0 plus p1 x and so on all the way up to pd x to the d.

And the crucial thing is that all of these pi's belong to my field. And I'm only going to consider values of x that belong to my field. So a notion of a root is can I find an x that belongs to my field so that P of x evaluates to 0. And now, it turns out that you can show this basic statement we know and love about real valued polynomials, that if they have degree d, they have at most d roots, well, that's actually a property of a general field. So it's not just a property of the reals. It generalizes even in this sense.

So just to give you a hint about what we're going to do later-- we're going to do an application where we talk about crypto called secret sharing, where what we want to do is we want to take a secret and divide it up among a group of n people so that any k people can recover the secret, but any fewer than that cannot recover the secret. So they can together share their piece of the key and find out the answer.

And what we're going to be interested in is the way we're going to build up our secret sharing schemes is we're going to be working with polynomials not over the reals, but over finite fields. So roughly, what's going to happen is that the secret will be something like the constant coefficient of my polynomial. So it's some finite value from some finite field.

And then what each of the parties are going to have is they're going to have the evaluation of this polynomial at some point. And together, when they have enough points and evaluations, they can interpolate the polynomial and find the secret. And otherwise, they cannot. So I'll just give you the hint of that for later.

But stay tuned. When we do cryptography, we're going to see how this identity, namely Euler's theorem, is the basis for modern cryptography and the RSA encryption scheme. And we're going to see how this statement about polynomials and the number of roots is the basis for modern secret sharing. But we'll get there. So let's stop here. And I'll take any questions you guys have offline. Otherwise, have a great spring break. Hope you guys do something fun.