[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: So today is the last day of our unit about random walks on groups. And the goal of the day is to digest, think about the Bourgain-Gamburd theorem about mixing on SL2(Fp). So we stated it at the end of last time-- first, I'll remember to remind everybody what it says.

OK, so it's about mixing on SL2(Fp). So it says if mu is a probability measure, SL2(Fp). Actually, maybe it doesn't matter that much. But maybe we set it in terms of subsets. OK, so the question is, when does the random walk coming from A mix? So then the conclusion will be, the singular value of this walk is less than 1 minus C of epsilon, meaning that it does mix. OK.

And to see that this happens, we need to know that A doesn't lie in any subgroup. And so we have a hypothesis that's related to that. It says that for any Borel subgroup B, A intersect gB-- so any g in our group. This guy is bounded by K of epsilon p to the minus epsilon times A. So only a small fraction of A lies in any coset of a Borel group. OK. And we put in parentheses that A is symmetric. We don't really need it, but it also doesn't matter that much. And it will make things a little cleaner. Yeah.

AUDIENCE: I have a question about the difference between using a subset versus a measure. Is it possible to pass between the two cases by weighting by different subgroup sets?

LAWRENCE GUTH: Yeah, that's right. So the question was, what's the difference between taking a subset or taking a measure? I think there's not that much difference. I think if you can use this as a black box and prove things about measures, I think we also could just use measures through the whole proof and it would work fine. OK.

So the proof of this theorem breaks into two important pieces which are both interesting and which are both connected in different ways to themes of our class. OK. So the first part says that if you just have pretty much any set and you do this random walk, the only way it can fail to mix is if it gets stuck in something which is approximately a subgroup.

So it can fail to mix if it's literally contained in a subgroup. And a little bit more generally, it can only fail to mix if there's a set which is a lot like a subgroup which contains a lot of the mass of the measure. So let's try to write this down. So this is probably true for any group. I'm going to write it for SL2(Fp) for now. And then as we talk about it more, let's keep our eyes open of what we actually used.

This is symmetric. And sigma 1 TA is bigger than 1 minus-- sorry. Here I'll put 1 minus epsilon. And then there exists some kind of approximate subgroup called H. Yeah, hold on. Let me adjust it a little bit.

Let me also say that A is pretty big. It is not too tiny. Then there exists some H so that mu tensored to some power of H is pretty big. And this H is sort of like a subgroup in the following sense. So H, if you multiply it by itself many times-- for example, 5. You could do more. This is less than p to the O of delta times H. And H also is not too big.

OK, so if you had another group, you would have to adjust these things. But I think that that's minor. I think that's all you would have to change. OK. So if you have a set which is a decent sized fraction of our group-- and let's say it's symmetric-- then either that random walks mixes or there is some subset which is sort of like a group. And when you take K steps of your random walk, you have a high probability of being in this set. And this set is sort of like a group in this sense.

OK. All right. And so there's delta and epsilon. Maybe at the beginning I should say this. And delta depends on epsilon. And it's positive. All right. Any questions or comments about theorem 1? That's theorem 1, is Bourgain and Gamburd, essentially. Yeah.

AUDIENCE: Where does the A to the 2.5 come from?

LAWRENCE GUTH: Yeah. So what's significant here is just that this is significantly smaller than the whole group. And if we-- yeah.

AUDIENCE: Does it equal kind of a power on the side of [INAUDIBLE]?

LAWRENCE GUTH: Right. OK. So this comes last time from this-- so actually, it's also important that this is bigger than 2. So we had this important property of SL2(Fp), that it's-- ah yeah, so actually this wouldn't work for any group. So we have this important property for SL2(Fp) from last time, that every non-trivial representation has a large dimension.

And that had the property that if this thing had a small L2 norm, so it was kind of decently well mixed, then we would get a spectral gap. And then if you kept going, it would become really, really well-mixed. And that's crucially being used here.

And so here you could put anything bigger than 2. And the point is that it can only go wrong by getting stuck in something of size really smaller than p squared. Because we know once it's spread over something of size bigger than p squared, then it's going to keep going and spread all the way. Yeah, thanks. Yeah.

OK. OK. Well, this theorem raises a question. Who are all the subsets H that are like this? It becomes a much nicer theorem when we can write them all down, or at least say something meaningful about them. So theorem 2' says something about that. So this theorem is due to Helfgott. And it builds on Hrushovski.

OK. So it says if H is a subset of SL2(Fp) and it has these two properties, so all of this stuff is smaller than p to the delta H. And H is at least p to the epsilon. You can read that off from here. It's not going to get stuck in something smaller than A. And at most, p to the 2.5.

Then the conclusion is that H is concentrated in a Borel subgroup or Borel coset. So then there exists a Borel subgroup B of an element g in our group, so that H intersect B is bigger than some small constant times H. So our approximate subgroups basically have to sit in a Borel subgroup.

This isn't a complete classification because some of the subsets of B are approximate subgroups, and many of them aren't, but it's a meaningful piece of information. OK, cool. So you put these two theorems together, you can prove that theorem pretty quickly. Yeah.

AUDIENCE: This C, does it depend on anything, or is it a universe?

LAWRENCE GUTH: Does it depend on anything? Let's put this here. Let's see. I have all these epsilons and deltas. I don't think it depends on anything. So I guess here, there should be this. So we can do this. So this is getting stronger as epsilon is getting smaller. Epsilon is telling us that our thing is not too small. And then I think that this is just a constant. I guess that does suggest that we could have said something stronger up there. Yeah. Oh, this is definitely H, not B. I think it's right. Yeah.

AUDIENCE: It's a question about the choice of g. If the H is a symmetric subset, would that kind of imply that the optimal g is always going to just be the identity? Because you think of H is kind of being centered in some way around. There would be some shifting subgroup by kind of an identity element, wouldn't change the intersection?

LAWRENCE GUTH: Yeah. So the question is, maybe if H were itself symmetric, which in fact, I think will happen in this setup, then maybe you don't need this g. And maybe it's actually only actual subgroups. I think that may be. So let me just scribe that I'm not sure we need g.

Right. Yeah, so a good example to think about is if you just take a coset of a Borel subgroup. And in the abelian world, a coset of a subgroup will be very small when you add it to itself. But in the non-abelian world it's not so clear. So I'll make a remark. If H is a coset of a Borel group, then how big is just H times H? Well, it's g times B times g times B. And because it's not commutative, I can't bring this across. So it's not obvious that this is small. So I think this is, in fact, large.

So based on that, I think it's likely that we don't need this. But it makes the proof a little shorter. OK, cool. So I have a couple of big-picture comments about this. One comment is let's compare to the perhaps simpler setting of an abelian group. So maybe you have the group of integers or the group of integers modulo p. Suppose we had a subset in there that had this property and this property. So when you add it to itself, many times it expands, but only by a little bit by a small power of the size of the set.

Then what can we say about the structure of such things? That's a big open problem in additive combinatorics. That's the Freiman-Ruzsa problem. If you have something even smaller here, like log p or so, then there would be a classification that would have to look like a generalized arithmetic progression. But in this range, there's not a whole lot. There's nothing like a structure theorem that currently exists in this range.

So in a certain sense, this is stronger information in the non-commutative world than we have in the commutative world. Of course, you might point out that this is not a complete structure theorem either. But this is a really strong piece of information about our group. And I think in some ways, it's stronger than any of what we in the commutative work. Yeah.

Another comment. So I'm not sure. So Hrushovski and Heathcote both worked on this. Hrushovski came from logic. And I think that he was interested in the idea from the point of view of logic of an approximate subgroup, like a basic thing we have axioms for is that the axioms for a group and what happens if you relax those axioms a little bit? What can you say? So an interesting question from the point of view of logic.

OK, cool. All right. So the goal for the class is to discuss each of these things in moderate depth, probably without giving complete proofs. And they have quite different ideas. But in both cases, the ideas are related to stuff in our class. So first, let's talk about theorem 1. So theorem 1. So we start with the measure mu, which is mu A. And we notice that mu L2 squared is 1 over the size of a. So that's is at most p to the minus epsilon. Because it was built into our assumptions here that A is at least p to the epsilon.

Now, let's start convolving mu with itself. And we want to see if we convolve it with itself enough times is what will happen to the L2 norm. So let's say mu 2-- so that's called this mu 1. Mu 2 is mu 1 convolved with mu 1. And in general, mu k plus 1 is mu k convolved with mu k. So we're convolving with this thing with itself many times. And we're going to see what happens to the L2 norm.

So recall that if at some point mu k L2 squared were to be smaller than 1 over p to the 2.1, then we would get that sigma 1 of T mu k was bounded by this. And then we would get that sigma 1 of T mu. So mu k is mu convolved with itself 2 to the k times. So that would be p to the minus 1 over 10 times 2 to the k. So if k is not too big, we would get a spectral gap. So that's our plan.

We keep doing these convolutions, and we watch what happens to the L2 norm. At every step, maybe the L2 norm goes down significantly. In that case, we're happy. Or maybe it doesn't go down. In that case, we're stuck. The stuck situation, we'll try to analyze it carefully, and see that the stuck situation only arises because of this structure, this setup. So let me write that out.

So let's say G stands for our goal all, which is that mu k L2 squared is less than 1 over p to the 2.1. If we get there, we're happy. p will stand for progress. We're not necessarily there, but mu k plus 1 L2 squared is bounded by p to the minus delta mu k L2 squared. And then S stands for stuck, which means we're not at our goal. And we are not making progress. Not goal and not progress.

OK. Cool. So as long as we don't get stuck, we will reach our goal in about over delta steps. So if S never happens, we reach the goal and thus the over delta steps. And no matter what delta is, delta is just a large constant. And that is enough to give the conclusion. So then if you do a little calculation using this, you'll get a spectrum.

So suppose we do get stuck, let's think about what that means. So there was a comment at the beginning of the class that perhaps there's not a big difference between a set and a measure, because if you have a set, you can turn it into a measure. If you have a measure, you can look at the points where the measure is pretty big, and you get a set. And that's actually what we're going to do.

So let's say that A lambda of mu k is the set of group elements for our group SL2 Fp, where mu k of g is approximately lambda. So we can decompose our group into pieces where mu k has various sizes. And one of these sizes must have a decent fraction of the measure. So we choose lambda so that mu k of A lambda is around 1. So we've cut our whole group into maybe logarithmically many pieces. Their total measure is 1. So one of them has to be pretty big.

All right. Now let's say Ak is defined to be this set. So what we've learned is that mu k is roughly uniform on Ak. And Ak has a good fraction of the total measure of mu. OK, great. So more or less, you could imagine that mu k is just the uniform measure on Ak. All right. So now what would we learn algebraically about Ak from the fact that we are stuck? It actually has a nice interpretation in terms of a previous character in our course, which was the idea of energy. So let's recall what energy means and put it in a non-commutative context. And we'll see that it exactly fits here.

So let's say if we have two subsets A and B in our group, the energy of A, B is the number of a1, a2, b1, b2. These guys are in A, and these guys are in B, so that a1 times b1 is the same as a2 times b2. Energy A, B. Now this is the same definition that we saw before in the context of a commutative group. Our group is now not commutative, which means that it matters what order I wrote those things in. And so it matters what order I wrote those things in. So E A, B is not actually the same as E B, A. But at least, otherwise it's a very similar-looking definition.

We can relate this to this setup. We can notice that E A, B is-- I take the characteristic function of A, and I convolve it with the characteristic function of B. When I do that, If I evaluate this at G, it will tell me how many ways are there to write G as little a times little b. Now I'm going to square it because I want to have how many ways-- I have that many choices for this and that many choices for that, and sum over B. So this is this.

OK. So notice I took something. I convolved it by something else, and I took its L2 norm. That's the character that's involved in being stuck. So something we can say. Actually, before we do that, let me make a little normalization. So remark. E A is in between A cubed and A squared. The A squared comes from-- I could take these to be equal, and I could take those to be equal. And A cubed upper bound comes from once I've chosen these three, there's at most one choice for the last one.

So these ones have a lot of additive structure. And the typical example is an actual subgroup. And these have very little additive structure Now let's relate to being stuck. So if we are in the stuck case for mu k, then we could see-- all right, I don't want to write this. So mu k is essentially 1 over Ak, its characteristic function of Ak. And if we want to be rigorous, I guess we could say something like this.

If we say mu k convolved with mu k is larger than p to the minus delta mu k, this is equivalent to saying that the energy, Ak with Ak, is smaller than some p to the delta times Ak cubed. Claim if we unwind this, we'll get this. So we unwind this, we'll see an energy, the energy of Ak. This is almost the minimum possible value of the L2 norm. And so it's almost the-- sorry, it's almost the-- this goes the other way. So let me say that again. So this is encoding this energy. It's a lower bound for it. And the lower bound is almost the biggest it could be. So this energy is almost the biggest it could be. Yes.

AUDIENCE: There could be a k squared if it's a lower bound for energy.

LAWRENCE GUTH: Right. No, it's Ak cubed. And this is a minus sign. So this is possible. Yeah. So this just got copied there. All right. OK, so if you look at this L2 norm, these are probability measures. So this L2 norm is at most that L2 norm. So when you look at this inequality, this L2 norm is really pinched. It's at least a small fraction. It's at least a large fraction of mu L2 squared and it's at most mu L2 squared. So it's just about as big as it could be. So this energy is just about as big as it could be.

OK. Cool. All right. So if we get stuck, there must be a subset of very high energy. And that is a step towards being an approximate subgroup. But it's not as strong as the conclusion of theorem 1. So at this point, the story connects to some characters that we talked about when we were doing projection theory over finite fields. It relates to the Balog-Szemerédi-Gowers theorem and the Plünnecke-Ruzsa inequality.

So first of all, let me just recall what they said in the commutative case. And then this is a non-commutative case. So then we'll have to talk about what's actually true in the non-commutative case. So recall that if we have A contained in Z, a commutative group, then we have two important theorems, Balog-Szemerédi-Gowers. Balog-Szemerédi-Gowers says that if A has a lot of energy, then it has a large subset. And the large subset actually has a small subset.

So this set has added a structure in one sense, but it has a large subset which has added structure in a stronger sense. And then once we have this-- so another important theorem, the Plünnecke-Ruzsa theorem that says if A prime plus A prime, smaller than L times A prime-- so L will be this-- then you can also look at the sum of three of them, or four of them, or whatever. So I'll just write this guy is bounded by L to the 5th. You could put other things. You get the idea.

This sequence of theorems is important because it gives amplification of structure. So this assumption is that A has some additive structure and some fairly weak sense. And that implies that a big subset has additive structure in a stronger sense. And then therefore, it has additive structure in even stronger sense. So if we were allowed to just apply these two theorems, then from this moment, we would get the conclusion of theorem 1. Conclusion of theorem 1 would be that there is some set H, which could have been Ak, which obeys this.

We can't do that because this is for a commutative group. So we need to talk about what happens in a non-commutative group. But I thought it was-- we'll talk about that. But I also thought it was a nice moment to recall these theorems and really appreciate that they're important and useful and do something kind of neat. Halfway through the course, check-in email. When I asked people what things that we did, do you feel like they really sunk in and you understood them and what things would feel it was maybe hard to remember what actually was happening? this theorem in particular was cited as something where it was hard-- people did not feel like they remembered very well what actually happened.

OK, so we get to talk about it again. Cool. Yeah. Oh, I guess another thing you might wonder is, OK, this logically is stronger than this, but does it really matter. Later when we look at theorem 2, we'll have a good example. Theorem 2 really requires this as an input. So we'll see an interesting example of why you might want this. Let's talk about the Balog-Szemerédi-Gowers theorem.

So I will tell you-- yeah, I guess first I should tell you what is true in the non-commutative world, which really looks pretty similar to this. And then we'll actually remember the proof a little bit to see if we can get the proof to sink in a little bit deeper. OK. So I'm going to erase this. And what I want us to remember is just that to finish theorem 1, what we need is a good version of Balog-Szemerédi-Gowers and Plünnecke-Ruzsa in the non-commutative world.

All right. So theorem. Non-commutative Balog-Szemerédi-Gowers. It says if A is a subset of a group, G, and the energy A A is at least k inverse A cubed, then there exists a subset A prime in A so that it's pretty big. And there are a few different things we could write, but actually, let me call this A tilde. A tilde times A tilde inverse is not too big.

It looks very little different. But let's take this opportunity to remember how the proof worked, and see if we can get it to sink in a little bit more deeply. All right, so proof sketch. All right. So remember that the energy A A is the sum over all of our group elements of r A A of G squared, where r A A of G is the number of pairs a1, a2 squared so that a1 a2 is g.

That's the energy. So for the energy to be large, there must be a lot of group elements g, where this is pretty large. We'll call those popular products. So p, the set of popular products. This is the set of g, where r A A of G is greater than k to the minus something times the size of A. So the size of A is the biggest it could be. And so they're almost as big as it could be. So they're quite popular.

All right. So it's an exercise to see that if the energy is really big, most of the energy actually comes from popular products. Energy A A is pretty much as-- it's more than 1/2 comes from. So 1/2 of the energy of A is less than the sum over g in the popular products of r A A of G. If this is having popular products, that's almost equivalent to having a lot of energy. Those are basically the same thing.

All right. So then we'll do a little thought experiment. Suppose that every product was popular. Every product in the product set of A had the same number of representations. Well, that number of representations would have to be pretty large because there's a lot of energy. And that would tell us that A times A was small. So then we would be done, and A tilde could just be A. But that's not true because it's not true that every element can be written as a product in the same number of ways.

So what we want to do, we're hoping to find a subset, A tilde, so that all the products in A tilde-- A tilde times A tilde-- they all can be written as a product in lots of ways, or something in that sphere. So now let's make the popular product to graph. So the vertices are G cross G, there's a copy of G. There's a copy of G. And we put an edge from A to B. So edge from A to B if a times B is a popular product.

Actually, this is not a G by G graph. This is an A by A graph, taking elements of our set and looking at their products. Not every pair is a popular product. They might not all be popular, but there are by hypothesis, because there's a lot of energy, there's a lot of popular products. So the number of edges is at least-- there are lots of edges in this graph. A tilde is going to be some subset over here.

And the good feature of A tilde is that for every a1 and a2 both in A tilde, there are many edges-- there are many paths. So let me draw it. So let's say this is a1 and this is a2. And then we're going to think of paths of length 4 that go from a1 to a2. So there are many or more than something paths of length 4 that go from a1 to a2 in our graph.

And how many? Well, from a typical point, there are about maybe 1 over k times A edges out. And so we have A choices, A choices, A choices. And then we have to go to a2. There was a graph theory lemma that we can choose a substantial set, A tilde, here that has all of these edges. That was the hardest part of the proof, but it's exactly the same as before. It had nothing to do with anything being commutative. So at that part, I won't repeat.

So there was a lemma that we can choose A tilde. So this is true. And A tilde is pretty big. OK. Cool. So let's name everybody in the middle here. I name these. Let me actually rename these. I'll call these A. A and A prime. And then we go over here to a1 to a2 to a3. So we know a a1 is popular. We know a2 times a1 is popular. We know a2 times a3 is popular, and so on.

What this allows us to do is to write A over A prime in many different ways. a over-- I have to be careful because I'm in a non-commutative group. I think I want to write a. Let's call this a tilde a times a tilde inverse. So that's a times a1. And now I'll put a2 times a1 inverse. And now I'll put a2 times a3. And then I'll put a tilde times a3 inverse.

So if we do this out-- so this guy here is a1 inverse a2 inverse. You can see cancelation, cancelation. And this guy here is a3 inverse, a tilde inverse. Cancel that. So that's this equality. It's not difficult. But we've written it in terms of four guys that are popular. So this a2 times a1 is not just a2 times a1, it's a popular product. It can be written in many different ways.

So a times a tilde inverse can be written as z1, z2 inverse-- or maybe p1, p2 inverse, p3, p4 inverse, where pi are popular. This can be done in many different ways. And the number of choices for the pi is at least A cubed.

So the conclusion is that the size of A tilde times A tilde inverse is bounded by-- I'm going to choose four popular guys. So that's p to the 4th. And then I have overcounted, because each person has now been represented A cubed different ways, and ignoring some factors of k. And so that's A. OK.

So I remember the first time I taught it I struggled a little bit with the intuition of this. The first thing you might hope to do is to choose A tilde, so that every product in A tilde can be is popular. They all can be written in many ways, just as a product of two things in A. And then there are clearly not so many of those. But I think it's not possible to do that. It's not possible to choose A tilde where every one of these products is popular in that sense.

But by being a little bit more flexible, we've chosen A tilde so that every quotient can be written in many ways, not just as a quotient of two things in A, but many ways in this slightly more flexible framework. Any questions or comments? Yeah.

AUDIENCE: So why can't you get the same result for A times A if you just choose paths of length 3?

LAWRENCE GUTH: Yeah, I think you could get the same result for A times A if you choose paths of length 3 or A times B. Yeah. Sorry, I maybe was not transparent about-- so the thing with paths of length 3 still works. I think that might produce two different subsets. Actually, maybe in the commutative case, I actually should have had two different subsets. At least initially. It works fine. The reason that I have A tilde times A tilde inverse, the motivation for that is going to come in a moment. And actually, I'm going to try to argue that this is better than A times A.

AUDIENCE: And just to clarify anything, the way you stated it before the popular products here is the subset of A times A, with the small projection. Because it was stated in terms of a subset having a large projection, then that leads to-- you can choose A subset of such that the subset is small.

LAWRENCE GUTH: Yeah, that's right. That's right. So in the projection theory version, this would be the small projection. And there'd be some big piece of A cross A that projects here. Yeah, thanks. Yeah. OK. Cool. So this is the non-commutative analog of Balog-Szemerédi-Gowers. It really looks a lot like Balog-Szemerédi-Gowers. And this inverse is because we want it. We probably could also have had A tilde times B tilde if we wanted.

The situation with Plünnecke-Ruzsa is actually a little bit more complicated. So non-commutative. All right. So let's ask a question. Suppose we have is A as a subset of a group, and A times A is small. Does it imply that A times A times A is small? So this would be that natural analog of Plünnecke's inequality. And the answer to this is no.

OK, here's the example. So let's say H contained in G is a subgroup. And A is H union one other element, which is not in the subgroup. Now what happens when I take A times A? A times A is H times H-- that's good-- plus, or union, H times g union g times H union g times g. There's nothing special about these. But because g is only one element, all of this stuff is small.

So we can see that A times A is smaller than-- 1, 2, 3-- it's smaller than 4A, probably 3A. But now what happens when I take A times A times A? OK, now I have three factors. And so that concludes H times g times H. In the non-commutative world, I cannot slide this over. And this is not so good. There is no reason for this to be small. And so the size of H times g times H could well be the size of H squared. OK. So that didn't work very well. What to do about it. There is a theorem that if A times A times A is small, then A times A times A times A is small. Yeah.

AUDIENCE: This is the counterexample that you have an example in SL2 Fp?

LAWRENCE GUTH: Yeah, it does. So H could be any subgroup, say a Borel subgroup, and G is not in there. Yeah, this will happen. So this could happen in SL2 Fp. So here's a theorem that if you assume that a product of three copies of A is small, then the product of four copies of A is also small. But that theorem is not immediately useful to us because we don't have any input that a product of three things is small.

There's another fix to this-- let's just suppose think about this example and how we might fix it. So this example has a lot of algebraic structure. This is the algebraically structured part, and this is garbage. And it would be helpful to separate A into the structured part and the garbage. So how might we do that? So let's look at-- so this is the continuing in the example. Let's look at A times A inverse. You could do a times-- anyway, it works a little bit better for A times A inverse.

So what do we have? We have H times H. Actually, yeah, so we will eventually look at A times A inverse, but let's think about how we might separate this. So this set has different pieces. And if we want to separate the algebraic part from the garbage-- so we'd like to find this part and distinguish it from that part. How are they different from each other?

Well, let's think about how many representations there are for each product in this product set. Any product in here has many, many representations, and the products here have only one representation. So this guy here is the set of popular products. And actually, once we identified it, the popular product set then is an actual subgroup. It has a lot of algebraic structure.

Now let that motivate the following proposition, which is an analog of p. So it says if A times A inverse is less than or equal to KA-- so A inverse, that means take each element of and take its inverse. And then if Q is the set of popular quotients, so it's the set of g so that r A A inverse of g is bigger than K to the something A. Then you can raise Q to any power you like. Q, let's say, we do it L times.

Q has a lot of algebraic structure, and the proof has a similar idea of writing things in many different ways. So proof. OK. Oh, actually, I think it's a little bit tricky. I think it should be that. So the proof is let's look at A times Q times Q is Q. Say, I'll illustrate the proof, and L2 equals 4. Let's illustrate with L equals 3. That'll give you the idea, times A inverse.

So somebody in here can be written as little a. And now this is a popular quotient. So this is little a2 inverse little a3. And there are many choices. And then actually let's illustrate with L equals 2. And then here, we'll have little a4 inverse little a5. We have many choices here. How many choices we have, about A choices, of about A choices. And then at the end, I'll do this. a6 inverse a7 about A traces. And then a8 inverse.

All right. So to say more better what I'm saying, suppose that I have an element of this set. Then I can write the element this way. And so I can say that x belongs to A A inverse. So x can be written as-- all right. So x can be written in many different ways. And for each way of writing it, we'll notice that this guy lives in A A inverse, that guy lives in A A inverse and so on.

So it can be written in around a cubed ways is as b1, b2, b3, b4, where the bi are in A A inverse. OK. So the size of this set, this is bounded by the size of A A inverse to the fourth. So I get to choose these B's. And now I've overcounted because for each x, there are around A cubed different choices of how I could write it. Divided by A cubed. And that is smaller than A.

OK. There's one thing I hid at the beginning. Let me do it properly and then pause and see what people think. So we were given A times A inverse is small. And that implies that the energy of A A inverse is large. This is stronger than this. And now I'd like to switch these. And this is actually the same as this. So even though I don't know A inverse times A is small, I do know that A inverse A has a lot of energy. And the reason I can switch them is if I have a1 a2 inverse equals a3 a4 inverse, that's a quadruple that's being counted in this energy.

Let's just rearrange things a little bit divide by a3. a3 inverse a1, a2 inverse equals a4 inverse. Then I will multiply on the right by a2. So that's a quadruple that's being counted here. Yeah.

AUDIENCE: Did you know that A inverse A has to be small because it's just the inverse of element in A inverse or A A inverse?

LAWRENCE GUTH: Yeah, maybe that's true. Yeah, maybe I made this too complicated. So these are equal, but you claim that actually A inverse times A is A times A inverse. So over here, I have a1 a2 inverse. If I invert it, if I invert that, I get-- what's the inverse of this? So invert this one. I don't think this quite works.

AUDIENCE: No, you didn't.

LAWRENCE GUTH: So I didn't manage to switch the order when I did that, but I can switch the order when I look at the quadruple. OK, good try. This is large. So therefore, there are lots of popular quotients. So Q is large, which is what we used here. Both of these proofs have following high-level idea. There's a lot of energy around. So it frequently happens that you can write a product or a quotient in many different ways.

If you could arrange a set where every product could be written in many ways, then you could get a clean estimate for the size of the product set. By playing around with the sort of products that you consider and by some experience and craft and skill, we can get that to happen. OK, so this is the analog of Plünnecke's inequality. The reason that it was desirable to put A tilde A tilde inverse in Balog-Szemerédi-Gowers is that that thing is designed to fit into this thing. So this Q is the approximate subgroup that we were looking for, which is, the conclusion of theorem 1.

So this finishes the discussion of theorem 1. There's a pretty detailed proof sketch. And it is a somewhat different setup that shows the significance of the idea of Balog-Szemerédi-Gowers and Plünnecke-Ruzsa in these combinatorial methods. OK, so now in the 15 minutes that's left, I'm going to try to say a little bit about the proof of theorem 2. The proof of theorem 2, it also connects with themes of the class, and especially with how some estimates are different in the case of a prime field or the case of a non-prime field.

OK. So let's have some let's do discussion of theorem 2. All right, I'd like to put us in the following framework. We're going to look at SL2 of k, bar where k bar is an algebraically closed field. You could imagine it's the algebraic closure of Fp. All right, and we're going to look at subgroups of SL2 of k bar. Let's remember some easy ones. A subgroup like this where A and A inverse are in this field. We have the Borel subgroup. So that's A A inverse t. And we have the unitary subgroup.

So those are some subgroups. We can take conjugates of them. Actually, let me come back to that. So then we could do two general things. We can look at subfields. So I can take SL2 k, where k contained in k bar is a subfield. Or you could take this any of with a subfield. Those are subgroups. And you could take conjugates. So conjugate any subgroup you get, another subgroup.

So that is a bunch of subgroups that we can identify without too much trouble that sit inside of SL2 of k bar. There is a theorem that says that this is not exactly, but reasonably close to the fullest. Dickson in 1901 gave a full classification of these subgroups. It's a little bit messy. It includes a variety of finite subgroups in certain cases and so on.

His theorem implies that if H is an actual subgroup, then either H is contained in a conjugate of SL2 k, where k is a subfield, or most of H is contained in a conjugate of L. If we know from the beginning that we're talking about a subgroup of SL2 of Fp, a proper subgroup of SL2, Fp, it eliminates the first case. So we must be in the second case.

All right. But Dickson's methods are-- some of them, at least, are not very robust if you replace the hypothesis of being an actual subgroup by something like H times H times H times H times H is pretty small. So here's a common example of a technique that we use to classify subgroups that does not do very well if we replace it by approximate subgroups. Example is there's the order of the subgroup divides the order of the group. That's really useful in 18.701. You want to classify the subgroups of subgroup with eight elements. But something like that is not going to survive an approximate subgroup in any useful way.

OK. So the proof of theorem 2 is based on a rather different strategy for this classification. That's due to Larsen and Pink. So let me say part A, which is the Larsen and Pink strategy. All right. So we're going to think about how a subgroup intersects these subgroups. So if you take SL2 k and you intersect it with one of these, it will either be empty, or just be identity or something like that, or it will be a subgroup that's like T0 of k.

So we're going to study H intersected with T of k bar, H intersected with B of k bar, H intersected with U of k bar. And we'll see the following behavior. The H intersected with T of k bar is roughly H to the 1/3 or big O of 1. So if I take SL2 Fp and I intersect it with this, I'll get T0 of Fp. It'll be about p of those, which is about the cube root of the size of SL2 Fp.

If I take some crazy torus in the algebraic closure that may not intersect except at the identity. So these are the two things that could happen. And then similarly with these. So they first proved that if you take an arbitrary group and you consider how it intersects these kind of fundamental algebraic subgroups, then it behaves a lot like it were SL2 of a subfield.

And then building from this, they were able to see that actually, it would have to be SL2 of a subfield. So that's their approach to the rough classification. Yeah, so this was a cool paper in pure group. They weren't interested in approximate subgroups. They were interested in pure group theory problem, classifying subgroups of SLD of k, or other like groups over finite fields. The complete classification is clearly hard because any finite group embeds in a permutation group which embeds an SLD of Fp. So all finite groups occur as subgroups of SLD of Fp for some D and some p.

And so it's not really plausible to give a complete classification, or anyway, it would be at least as hard as classifying finite groups to give a complete classification, but they were able to classify, I think it's subgroups of SLD of Fp, which have at least some size, like p to the epsilon or epsilon times D or something like this, the ones that are decently large. That Was in around 2000. So this really is a hard thing.

But it also turns out that their method interacts well with just assuming that H times H times H times H is small. So here, this whole argument works if we just assume that H times H times H times H times H is equal to H. OK. I may or may not have two minutes to try to describe what they did, how this part works.

AUDIENCE: Significance of the algebraic closure. If we only care about subgroups of SL2 Fp, what information are you gaining from intersecting with--

LAWRENCE GUTH: Yeah, so the question is, Why bring into play the algebraic closure?

AUDIENCE: Also, what is the algebraic closure of Fp?

LAWRENCE GUTH: And what is the algebraic closure of Fp? Yeah. You do need this for the following reason that should be familiar to us even as analysts. So suppose you have a matrix in SL2 r, and you would like to diagonalize it. You'd like to look at its eigenvalues, or whatever.

Well, you can't necessarily do that in SL2 r, because the eigenvalues could be complex. So in order to diagonalize it, you might want to move to SL2 c. And diagonalizing matrices is useful because you can classify matrices by their Jordan canonical form. And this is a convenient way to organize things. That's the same thing that's going on here.

AUDIENCE: The intersections look the same, though. The intersection of if H is just a subgroup of SL2 of k, then the intersection with T k bar is the same as the intersection of T k.

LAWRENCE GUTH: Yeah, that's correct. That's correct. But OK, so the comment was that if-- so what did this-- so if H is in SL2 k, then H intersected with T0 of k bar is the same thing as H intersected with T0 of k. Yeah. OK. I think the problem with is that I forgot to say something about what's written on the board here. This T is not T0. The conjugates are really important. So T of k bar is a conjugate of T0 of k bar. Yeah.

AUDIENCE: Are the elements that are being diagonalized just elements of SL2 or bigger linear operators responding to SL2? I think if it's just SL2, you only need a degree-2 extension.

LAWRENCE GUTH: Yeah. So the comment is that if we just want to diagonalize matrices in SL2 Fp, we probably only need a degree-2 extension of Fp, which we might feel more comfortable with. I think that's probably right. I think we could probably work with the degree-2 extension. OK.

All right. So the output of the Larsen-Pink argument is after conjugating. Yeah, so first of all, we can arrange that our H intersected with some unitary guy, is pretty big. After conjugating, we can arrange that this is U0 of k. So H contains some guys like this, where T is in some subset that I'll call E. And E is pretty big. Moreover, this part should not expand too much when we take products.

So here, let's call this thing a unitary piece of H. UH times UH, also, this isn't completely obvious, but it's not that difficult. It shouldn't be that much bigger than. And that tells us that E plus E shouldn't be that much bigger than E. Also, we should have this U is in some B, and H intersect B of k bar. So we've already conjugated so that U is U0. So this is going to be B0. And this should be around H to the 2/3.

Inside of there, we'll have a T. And H intersect that should be like H to the 1/3. So H intersect T0 k bar. That's going to be some guys that look like this. And these A's live in some set that I'll call F. So then if you think about how these matrices act on these matrices, we should get-- so T0-- so let's called this guy the T part of H. And U, say, here.

Well, OK. This thing here is the B part of H. So the B part of H times the B part of H should be not that much bigger than the B part of H. And if you unwind what that means, you multiply these, you'll have things like a a inverse at, where a is an F, t is an E. So the conclusion is that E times F is also not that much bigger than E or F.

So in other words, once you have a unipotent piece where you have these guys, then this multiplication operation just gives you addition of the numbers in here. So you have a simpler thing. And then once you have a bunch of upper triangular guys, then this matrix multiplication induces this regular multiplication, and we have small products.

This is consistent with the possibility that E and F could both be just some subfield k or some approximate subfield. But the sum product theorem says that Fp has no approximate subfields. And so the subproduct theorem rules out the possibility that this could happen. So the conclusion is the sum product theorem applies that this can't happen.

What's on the board? Let me I should add one thing. So we're going to assume that H is not contained in a Borel guy. So this whole thing that wouldn't be true for the Borel subgroup. It's only true. Once we assume that our subgroup is not contained in Borel, and then we get this. And that tells us that the intersection sizes have to match SL2 of a subfield. And then if you look more carefully at the matrices that appear in the intersections, you see that the entries there have to be kind of like a subfield. And then it becomes a question of whether there are approximate subfields.

OK. Let me say it one more time, and then we will break for the week. So the theorem says that if you have a subgroup, which is not pretty much contained in a Borel subgroup, then it must be SL2 of a subfield. It's going to be true for subgroups, but it's also true for approximate subgroups. So imagine we have an approximate subgroup, we're going to gradually prove that it has more and more properties of SL2 over a finite field.

So the first property we will consider is how it intersects various algebraic subgroups of SL2 of k bar. The sizes of those intersections behave in a way that matches what you would expect from SL2 of a subfield. There's a substantial thing to say about this proof that we won't have time to do in this class. But once you know that, these subgroups are useful because the matrices in here are simpler than general 2 by 2 matrices.

So the unitary subgroup has matrices that look like this. And when you multiply them, you just add the upper right hand corner. And so now if you just look at these upper right-hand corners, you see that adding those numbers doesn't change very much. That's a property of a subfield. So it's behaving like a subfield. And the upper diagonal matrices, the multiplication rule is only a little bit more complicated. It's a pretty simple mix of multiplication and addition. So you could see that these diagonal entries, they should be almost closed under multiplication. And they interact with each other. And so you would get something like this. Anyway, all those things being consistent algebraically with being a subfield. OK, yeah.

AUDIENCE: Could we use 2.5?

LAWRENCE GUTH: Where did we use 2.5? Yeah, that's a good question. Actually, I don't see where we used that here. Yeah, we used that here. So at the end, we have to say E and F cannot be approximately a subfield. And that requires that they're significantly smaller than Fp. So if you take Fp and you remove two or three elements, it is almost closed under addition and multiplication. So this could happen then if you were allowed to take a really large subset of Fp, a really large subset of SL2 Fp. OK, cool. I apologize for going over. Thanks for your patience. Have a good weekend. I will see you next week.