[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: All right. Well, we have two more classes. And this is day two of trying to grapple with the recent big breakthrough in projection theory. So this is sharp projection theorems part two. In part two, we're going to talk about the AD regular case. So let's first remember our goal and remember what is the AD regular case.

OK. So our goal is the Furstenberg set conjecture, which is a theorem of Orponen, Shmerkin, Ren, and Wang. Says if E contained in R2 is a delta TC set, and then for every point in E, there are a bunch of delta tubes going through it. So T sub x is a set of delta tubes through x and the directions, if you like, of this set of delta tubes is a delta SC set.

OK. And parenthetically, we can assume a couple of things. We can assume that this is uniform, and also that the number of them is like delta to the minus S. And S should be bigger than 0. OK. So we put all of these tubes together to get one big set of tubes. And the question is, how many tubes are there? So then the conclusion is that the number of tubes is at least a little fudge factor.

So this whole thing, I'll abbreviate like that. And then the minimum of a few options. Which we'll call option A, option B, and option C. And these exponents are hard to remember, but I would prefer to remember-- for everyone to think of them in terms of three different scenarios.

So option A, that corresponds to the-- I call it the stars scenario. Have a bunch of points, and through each, you have a bunch of tubes, and they're just unrelated to each other. Option B is more intricate. It's the integer grid scenario. And option C is a kind of random scenario. So E is arbitrary and T is random.

OK. Now it's really a good idea to think about this theorem in these three different scenarios. So it's good to separate out the case where each one of those things is the minimum. So let me do that and recall for you how much we know about each case so far. And I said something about this last time, but what I said wasn't the best thing to say, so I wanted to do it again and get it right.

All right. So there's the case where A is the minimum option. And if you do a little algebra, that corresponds to S being at least as big as T. And in this case-- so if S is at least as big as T, then the theorem is true. So Furstenberg is true. And it follows just by double counting. So we did it very early in the course, like on the second day.

OK. Then I'll come back to B, which is the most interesting. C is the minimal. If you do a little algebra, it's if S plus T is at least 2. And in this whole range, where S plus T is at least 2, the theorem is also true. And it's true by the Fourier method. So it follows from the theorem we proved in the second or third week using Fourier analysis. So the whole case where C is minimal, we also already understand that.

OK. And that leaves a case where B is minimal. And that's like otherwise, if neither of these things are true, then B is minimal. And so that's the content of this theorem. So this is the hard new thing in the OSW thing. OK. Cool. Another little comment is that the number of tubes is not my favorite parameter to keep track of.

My favorite parameter to keep track of is I'll call it R of ET. Which, in words, is like the typical number of delta balls of E on a typical tube of T. Let's make a picture. So we have some delta balls, and then we have some delta tubes going through them. And in an interesting scenario, each tube will have many balls in it. And R is this number of balls in a tube.

So algebraically, by double counting, the number of balls in a tube would be the number of points times the number of tubes going through a point divided by the number of tubes. So that would be the algebraic definition of R. And then it's easy to convert bounds for T into bounds for R. So our theorem is equivalent to that R is upper bounded by the maximum of three things. Scenario one is A.

So scenario one, the number of balls per tube is one, which is always easy to arrange. B, you have to do a little computation. It works out to this. I don't really care what the exponent-- it's not important what the exponents are. And C is delta to the negative T plus 1. OK. And this delta to the negative T plus 1 has a nice interpretation which might clarify a little bit what we mean by this random case. So if E is a set-- is a delta TC set, so the size of E is delta to the minus T.

Let me say it this way. Suppose that E is a set of delta to the minus T delta balls in the unit ball. And then I take the average over any choice of delta by 1 rectangles of the number of delta balls that the rectangle covers. You would get delta to the minus T plus 1. So it's not difficult to locate rectangles with that richness. Almost any rectangle would do the job. And scenario C is when the rectangles just have this average richness. There's really nothing special about them.

OK. So next, let's remember-- so our goal today is to talk about the AD regular case. So let's remember what is the AD regular case. OK. So we're going to suppose that E is uniform. So that means that our delta is some capital delta to the M, where M is a large number.

And the uniformity is that if I take E and I intersect it with a cube of side length delta to the j, and I ask about its delta to the j plus 1 covering number, that's around some beta j for all the dyadic Q delta j that intersect. OK. So these beta j are the branching numbers. We drew pictures last time to illustrate what they mean.

And so definition would be that E is C alpha is to be regular. So actually let's say it's a delta TC AD regular set if has this branching structure like above and if I take the product j equals 1 to capital J of beta j. So that would be the delta j covering number-- delta to the j covering number of our set.

It is at most C times delta to the j to the minus t, and it is at least 1 over C times delta to the j to the minus t. OK. All right. So let's label all these hypotheses in our theorem by H, our hypotheses. And what we'd like to understand is given these hypotheses, how big could R be? And for today, we're going to talk about the AD regular case.

So let's give that a name. So RAD of ST delta C. That is the maximum over E and T obey our hypotheses, and E is a delta TC AD regular of r of E and T. How big could it be in this case? OK. It's really a mouthful of parameters, so let me just say, to reduce a little bit for today, we'll have C equals 1.

So that means it's really AD regular, and everything we say actually will work for small c that are bigger than E. That will allow us to not write the letter C. And then S and T will not be changing through our argument, so I won't write them every time. So now I'll just abbreviate it as RAD of delta. OK.

All right. So the theorem that Orponen and Shmerkin proved is that they proved this thing in the AD case. So RAD ST delta is indeed bounded by that. I'm not sure I copied it right. Got the minus signs wrong. So that's our goal for today is to understand this theorem, which is a really important special case of that theorem.

Cool. So it's clearly a special case of that theorem. And at first it might seem like a very special case, because these branching function basically-- there's a huge range of what the branching function could be, and AD regular is just one particular choice in this huge range. But it's a particularly important case, and we'll see tomorrow or next class that this is really a crucial case. It plays a big role in understanding the theorem in general.

OK. Cool. Yeah. It was a really important realization in thinking about this theorem that we should break it into cases. The different cases have their different things that are useful and their different things that are difficult. And so that recognition is key. So the next thing I want to try to talk about is what is special about AD regular sets that we could use to get a handle on them that we couldn't use just for the general setup.

All right. So this proof is quite difficult. We won't do all of it. But I would like to share some key ideas about it, which I think are really cool and are worth taking with you away from this class. All right.

OK. So in a nutshell, the important special thing about AD regular sets is in this lemma that has a one board proof, which is called the submultiplicative lemma. OK. So it says that RAD of delta-- let me put it this way. If delta is delta 1 times delta 2, then RAD of delta is bounded by-- and with a little bit of fudge-- RAD of delta 1 times RAD delta 2.

That is submultiplicative. And it follows from understanding, thinking a lot about a picture of an AD regular set. OK. So remember that E is a set of delta balls and T is the set of delta tubes, and we're going to imagine thickening them. So in this setup here, delta is smaller than delta 1 and delta 2, which are smaller than 1. So delta 1 and delta 2 are each bigger than delta.

So I'm going to thicken them and I'll get E1, I'll call it, which are the delta 1 balls around E, and T1, which are delta 1 tubes. All right. So what does it look like? Let's think about E and thickening E. So here is E1. It is a bunch of circles. Each circle has radius delta 1. And inside of those circles, there are some delta balls.

And those guys are my delta balls E. OK. Now, T1 are some delta 1 tubes. So they might look like this. Let's put one more ball up here. That looks like that. So those are my fat tubes, T1. And then inside of them, there are some delta tubes T2. So let me draw that. I've got too many, but I'll try to give some idea.

So through this point here, there might be a couple of red tubes that are inside of that orange tube. And then there might be a couple of red tubes that are inside of that orange tube. OK. I just drew the ones through this point. Experience has taught me that if I draw more, this picture will get too messy. But you should imagine that this is happening at every point. OK.

All right. OK. So the first observation is that E1 and T1 obey the hypotheses of our theorem. Because actually, if you take a delta T set of balls and you thicken them up to radius delta 1, you get a delta 1 T set of balls. At least, because if it's uniform. And if you look at who's going through this point, we have a delta S set of directions leaving that point.

And if you thicken them up a little bit, you get the directions corresponding to the orange tubes. And that will be a delta 1 S set of directions. OK. Great. Yeah. OK. Actually, let me put the punch line, and then we'll think through why this works. So RAD of delta. That's the number of white balls in a red tube. But that's bounded by the number of delta 1 balls in a delta 1 tube.

So I look in this orange tube, I see how many of these balls there are, times the number of delta balls in a delta tube within 1 delta 1 bar. So if I'm looking at this red tube here, I want to know how many delta balls are in it. It's bounded by the number of big balls that it enters times the number of little balls in each of the big balls. OK. So we'd like to relate these two things to our problem at a coarser scale.

And this one here, well, E1 and T1 obey our hypotheses. So this one here is bounded by RAD delta 1. Now here, if you take E and T and you restrict them to a delta 1 ball, and then you magnify it to bring it up to scale one, this also obeys our hypotheses. So if you look inside of here, you see a little delta TS set.

This and these red tubes, that's the same situation on a smaller scale or on a coarser scale. So when you blow this up, if you blow it up, then you get delta 2 balls and delta 2 tubes in the unit ball. So this guy is bounded by the RAD of delta 2. OK. Let me call this claim one and claim two. All right. So this picture is perhaps the most important thing to take away from this proof in this class. So take a moment and see if you have questions about this proof.

AUDIENCE: [INAUDIBLE]

LAWRENCE GUTH: Yeah. OK. So the question is, do we need to check that everything is AD regular. Yeah. OK. So let me add to these claims, and then if we want to, we can add details to the proof of these claims. E1 and T1 obey the hypotheses, and they're AD regular. And we need that this blow up here obeys the hypotheses and it's AD regular.

And if you believe those two things, then you can plug-- you just plug in and you get that. Yeah. And we can talk a little bit more about those two things if you want. Are there other questions or comments? OK. So I'll probably add a little more, but actually, I wanted to ask you all a question. In this argument, where did we take advantage of the fact that the original E was AD regular? Yeah, Jonathan?

AUDIENCE: Is it like picking one delta 1 ball and magnifying it?

LAWRENCE GUTH: Yeah. So it's related to when we picked this one delta 1 ball and magnified it. Yeah. OK. So let's talk about one and two a little more carefully. All right. So let's discuss one and two. OK, so point one. If E and T are uniform-- which we can pretty much assume from the-- which we can assume from the beginning with a little pruning.

Then E delta T implies that E1 is delta 1T and Tx delta S implies that T1x is delta 1x. And also, if E is AD regular, that implies that E1 is AD regular. Because this is saying all the branching numbers of E are the same as each other. And this is saying that all the branching numbers of E down to the depth of delta 1 are the same as each other.

OK. We didn't need E to be a regular to do this or this, but it comes in in part two. So because E is AD regular, then we get that if you take E and you intersect it with a small ball and magnify it, then this is a delta-- I guess delta 2T and AD regular.

So let me illustrate this with pictures instead of equations. So what does it look like? Let's draw an AD regular set and a not AD regular set, and we'll look at both of them. So if E is AD regular, what do I do? I pick, say, four of them. Let me do this in the color. Then for each one that I picked, I subdivide it.

And I pick four of them. OK. I'll stop drawing. You get the idea. OK. Now, this is a self-similar construction. If I take this thing in here and I blow it up, it's just like what I started with. So it will be AD regular and it will be-- in particular, in delta, there will be a delta T set for the appropriate delta.

OK. What would it look like if E was not AD regular? Well, for example, I might do the following thing. I might cut it into a four by four grid, and I might take all of the boxes. And then I could cut each one of them into a four by four grid and pick just one box. So this was the well-spaced example that we drew last time.

OK. So this whole thing is a delta T set. But if I took just one of these boxes, what I would see is much more concentrated. This is not a delta T set. So if I were to take this guy and take it out here and magnify it, then inside of here I would just see this. This here is a blank T set, but this here is not a blank T set.

Whereas in the AD regular case, if I were to take one of these and bring it out and magnify it, then what I would see in here would be very similar to the original set. This is similar to E. So this here is a blank T set and it's AD regular. And that thing in there is still a blank T set and it's still AD regular.

OK. So this is the special feature AD regular. The word regular refers to the fact that all the different scales look the same. And so when we look inside this ball here, we have another example of our original problem at a coarser scale. And that feeds into an inductive argument. For example, we can say this submultiplicative lemma. Yeah?

AUDIENCE: The first equality over there, do you need some of uniformity on the number of the balls from A to C? Because RAD is like the average, and you're saying something like the average is less than or equal to [INAUDIBLE] the averages.

LAWRENCE GUTH: Yeah. That's right. That's right. Yes. That is a good point. So the comment is that RAD was defined as an average number of delta tubes per delta ball. And here I was pretending that all the tubes are average. So let's call this a proof sketch. You do need some uniformity. If you think back to the original setting, you have this set E and then you draw all these tubes. You can sort them by how many delta tubes are in each one.

And you take the category that has the most delta tubes, and that will still you can work with that. So you can do-- you're correct. We have to uniformize it, and it's possible to uniformize it. OK. So let me describe that. So we can do some uniformizing so all tubes have the same number of balls. OK. Other questions or comments? OK. Cool.

All right. So this little lemma is very important. For people who have seen some other things in harmonic analysis, there's a submultiplicative lemma that's sort of like this in decoupling theory. And it's the beginning of what's special about decoupling. If those words don't mean anything to you, don't worry about it. So let me show you some corollaries, some applications of this lemma. And then we'll go in earnest into using it to prove this theorem. All right.

All right. So the first application is sort of tongue in cheek, but also I think it's philosophically important. I'm going to call it-- so applications. I'm going to call it a brute force proof of our theorem, of theorem of the AD regular case. So here's how it goes. So for some specific delta 0, we check by brute force that RAD ST delta is bounded by the right-hand side. We won't be able to get it perfectly, but we could get something like this.

OK. So how could you do this? Once you fix a particular delta 0, like 1 over a million, there are really only finitely many different ways you could arrange the delta balls, and there are really only finitely many ways you could put the tubes. So there's a giant brute force calculation that you couldn't really do but we could imagine, and you could just check this.

Once you have this, then you use the submultiplicative lemma many times. And so when you used it once, you would get RAD ST delta 0 squared. That would be bounded by right-hand side times delta squared to the minus epsilon. Because the right-hand side is just the power of delta 0, so it multiplies. Everything fits and you keep going.

OK. So this brute force computation is so large that it's out of the question to do it even with a supercomputer, et cetera. Nevertheless, it is worth noticing that this very difficult theorem has a one blackboard proof up to a brute force computation, up to a finite computation. Most difficult theorems in math. It's not obvious how to reduce to a finite computation. I wouldn't say this is obvious either, but for most difficult theorems in math, there's not a half page argument to reduce it to a finite computation. For this one, there is. So it shows that this is quite important.

OK. Let me give another application. Second application is the general case. General AD versus the projective AD. So the course was called projection theory. And at the beginning we talked about problems about orthogonal projections. And this theorem is really more general. I guess I've erased the theorem. The theorem we're talking about today is more general than just orthogonal projections.

And if you had some orthogonal projections, the way you would get all these tubes is you would pick a direction from your direction set, and then you would get a bunch of tubes in that direction that cover your set. So I make a definition. Let's say that ET is projective if for every x in E, the direction set of the tubes Tx-- let me put it this way. For every x1 and x2 in E, the direction of the tubes through x1 is the same as the direction of the tubes through x2.

In other words, you have just one set of directions, and Tx is all the tubes through x in that set of directions. OK. So if you think about the problem that we wrote down, this is a very special case. So it used to be at every point you picked your own independent set of tubes, but now there's an extra rule that the directions have to be the same at each point. So a very special case.

OK. So let's say RAD projection of delta is the maximum over the usual stuff. So ET obeys our hypotheses. E is AD regular, but also, ET is projective. All right. So in general, we're taking a maximum. And this is a restricted maximum, so we clearly have that RAD projective is smaller than RAD.

But it turns out that they're basically equal to each other. So let's make that a lemma. So RAD of delta is less than here. There is some fudge factors, so I'll write something like this C epsilon delta to the minus epsilon. Actually, let me-- I won't quite say that that's true. But let me just say something.

OK. So here's a cool observation based on that diagram and that proof. So notice that in the submultiplicative lemma if delta 1 equals delta 2 equals square root of delta, then the inner-- the small ball problems are all projective.

OK. So in this problem, we have a delta scale Furstenberg problem, and we broke it up into two pieces, a coarse piece at the large scale and a small piece inside each small ball. So I claim that this piece inside the small ball is projective. Why is it? Well, so take one of these points, and I'm going to look at the directions of the red tubes going through this point.

Now I'm working in this small ball so I can only distinguish them up to an angle of the square root of delta. So what directions will they be? They'll be the directions of these orange tubes going out through here. So there could well be two long red tubes going through here that are in the same orange tube. But inside of this small ball, they'll just be on top of each other. I won't see any difference between them.

So the directions that I'll see within the small ball, they just correspond to these orange tubes. And if I look at different guys inside the small ball, I have the same set of orange tubes. OK. So that gives me an inequality. RAD delta is bounded by RAD delta to the 1/2 times RAD projective delta to the 1/2. Yeah?

AUDIENCE: I thought that you were writing projective [INAUDIBLE].

LAWRENCE GUTH: Right. Only the inner one is projective. The outer one is arbitrary. However, we can repeat this. So now that's bounded by RAD delta to the 1/4, RAD projective delta to the 1/4, RAD projective delta to the 1/2. And now you can keep going.

OK. So you do this many times. There will be a T appendix left here that still has no projective on it, but it will be very minor. And the rest of them all have projective. And so the conclusion is that to prove the theorem, it suffices to check the projective case.

So that's a cute application of the submultiplicative. OK. So those are a couple of cute applications of this lemma. Yeah?

AUDIENCE: Is that an extra RAD delta to the kind of 1 over epsilon where the delta T epsilon error term comes from? Or is it a conditional thing as well?

LAWRENCE GUTH: There are several places where we might incur a delta T epsilon in the final thing, and this is one of them. Yeah. Yeah. OK. So let's say something big picture about trying to prove this theorem. So we were saying yesterday that for a while in this subject, there were only kind of-- there were only epsilon improvements on some classical things. And what was very striking to me is to go from just making epsilon improvements, to go to proving the sharp theorem.

And while Pablo Shmerkin was visiting, I talked with him a lot about all this stuff. And he said to me something one day that really stuck with me. He said, philosophically that-- yeah, here's how I would like to put it. That in order to prove this theorem, what we need to do is make an epsilon improvement on the submultiplicative lemma. So let me write that down, because I think it's worth taking with you out of this class.

So philosophy. To prove the theorem, what we need at epsilon improvement on the submultiplicative lemma. OK. So in other words, we'd like to say either one possibility is that RAD square root of delta already obeys the bound. OK? Well, then we're happy. And by the submultiplicative lemma RAD of delta also will obey the bound. So we'll be happy.

Or we have-- so I'll write down the submultiplicative lemma. So that's the submultiplicative lemma. Or we have not this lemma, but a little tiny improvement on it.

OK. So in this case, of course, we're happy. And in this case, we can iterate. And so you can think of RAD of delta as being some power of delta. And then every time we do this, that power goes down, or some negative power of delta. Every time we do this, that power is going to go down. It will keep going down until we get into this case and we're happy.

All right. So let's illustrate that. So if, say, we have RAD of delta-- let me write it this way. It'll be easier. It's delta to the minus gamma. Then we'd have RAD of delta squared would be bounded by-- so this would give us delta to the minus 2 gamma plus epsilon.

Then we would do this again. RAD delta to the 1/4 would be bounded by delta to the 1/4 to the minus gamma plus epsilon plus epsilon. We keep going and this exponent keeps getting better and better, until we can't make an epsilon improvement anymore. So if we have this dichotomy, then we can get the theorem. OK. Cool.

All right. So then you might reflect, well, why is it easier to make an epsilon improvement kind of theorem than it is to prove a really sharp bound straight out? Well, there's a thought process for that, which is we already have a result and then we want to make an epsilon improvement. Then we say to ourselves, what would it look like if the result we have is sharp? And maybe there's a lot of specific things that have to happen for the result we have to be sharp, so we get a lot of structure.

Then we have to get a contradiction, and that shows you can do just a little bit better. So that's why that epsilon improvement result is easier than just getting a great bound straight away. So if we go back to Bergen's projection theorem and we tried to make it an epsilon improvement in a projection theorem, well, we didn't say this at the time, but it might have been good to say. What would it be like? What would it mean if the bound we have so far is sharp? It would mean we have this structure and that structure, and eventually, we have this subset that is almost closed under addition and almost closed under multiplication.

OK. So the idea is we are going to try to do the same thing. We're going to think about what would it mean if the submultiplicative lemma was sharp. We're going to see that it just means the set has to have a lot of structure, and eventually, we'll get down. Maybe not today, but eventually Orponen and Shmerkin got down to the structures, got down to the Bergen projection theorem and got a contradiction. OK. OK. So let's talk about what would it mean if the submultiplicative lemma is sharp.

OK. So let me draw a couple pictures for you. And one of the pictures corresponds to the submultiplicative lemma is sharp and the other one, it's not. So we'll have some fat balls in a fat tube. And then in each fat tube, in each fat ball, we have some little balls. And the little balls are in little tubes.

So there's your typical little tube OK, in each of these balls, there are some-- in each of these fat balls, there are some little balls. OK. So in this scenario, in this picture, would the submultiplicative lemma be sharp or not sharp? OK. Stand next to the picture. This picture, would the submultiplicative lemma be sharp or not sharp?

AUDIENCE: I don't think it's sharp because there's just one tube for the coarser scale. There's several tubes for the fine scale.

LAWRENCE GUTH: OK. So let me clarify what this is a picture of. So there are many coarse scale tubes, but I just drew one of them, and I'm just drawing what's happening inside of one of them. So inside of that one coarse tube, this is what we see. Based on this is the--

AUDIENCE: So there's only one fine scale tube?

LAWRENCE GUTH: There could be more fine scale tubes, but this is a typical one. Yeah, so there are many orange tubes and many red tubes, but I'm drawing only one of each to focus our attention. OK. So let's review for a second the proof of the submultiplicative lemma, because it's really important. And maybe we'll go in a little deeper, and then I'll try to show it again.

OK. So our goal is to estimate RAD of delta, which is the number of delta balls per delta tube. So I picked a delta tube like this red one here. I'll have to estimate the number of delta balls in it. First I estimate the number of fat balls that it intersects. Then I multiply that by the number of delta balls that it intersects within each fat ball. And that was my bound. And then each one of those things, I bounded by induction and I got those two numbers.

All right. So if I take a particular fat tube, and inside of it a particular thin tube and I see this, would that argument be sharp? Right. OK. So the suggestion is that this is not sharp because inside of this red tube, most of the fat balls are contributed nothing. Only this one fat ball is contributing.

And it would be sharp if the number of delta balls in the fat ball was the number of-- sorry, the number of delta balls in the delta tube was the number of fat balls times the number of delta balls in one of them. So let me show. So this is not sharp, and let me draw the one that's sharp. You'll see the difference. This is not sharp. Delta tube should look like this.

OK. So the delta tube should intersect a bunch of fat balls. It should intersect that many fat balls, as many as possible. And inside each one of them, it should intersect a lot of delta balls, as many as possible. That's what it takes for the submultiplicative lemma to be sharp. So it should look like this. So this is sharp. Cool. And this is a kind of striking picture.

So in this ball, these points are in just a few tubes in a certain direction. You could also imagine we're in the projective case. And in these other balls, the points are in those same tubes. It goes all the way across. Of course, the balls don't have to look exactly the same as each other, so maybe a better picture might be like so. But the tubes are exactly the same. OK.

OK. So let me pause. I'm going to erase a little bit to make some space around here. Is that OK with everybody? That if the submultiplicative lemma is sharp, it has to look like this, not like that. OK. All right. So let's erase around here a little bit. OK. So we already reduced to the projective case.

So let's say D is the set of directions, which is the same everywhere of our projections. And so D is a delta S set in S. And here I fixed a typical fat tube. So here's a typical root delta tube. And I'm going to look at the angles that are within this fat tube. So here's our circle. And there's a little range of angles here. This is a delta to the-- there's theta, which is a delta to the one half arc. That corresponds to the fat tube.

OK. And so D is some delta set, delta SC set in here. And let's say that C is D intersected with theta. So it's just these directions. OK. So the red tubes correspond to one of the directions in theta. So this one there. Red tube. And so when I project this set in the red direction, I just get a small set. But it's not just the red direction. I could draw some other directions inside of the tube also.

Let me do one more. I'll take here a blue direction inside of here. So I draw the tubes in the blue direction. What would it look like? I think I need to make this picture come out nicely. I need to add a few more points. So that'd be one more red tube. Oops. And now I'll draw a typical blue tube. The blue tube also manages to intersect our points a lot, so it should look like this.

I don't think I can draw any more without making a complete mess, but you get the idea. So the blue tube passes through all these balls also, and each time the blue tube passes through a ball, it manages to line up with one of the red tubes. All right. OK. So we have a kind of special situation here with a product structure. So what is our set?

So let's say over here these dots, the projection in the red direction, let's call this A. And then down here, we can make a set that illustrates where the fat balls are in the fat tube. Let's call that B. And so what we see is that for every C in C, the projection in the C direction-- or I'll write it this way. A plus CB-- that's the projection in the C direction-- is not much bigger than A.

The projection in every one of these directions is small. They're all about the same as each other. A was the projection in the red direction. The blue one is just as good. So A plus CB should be not much bigger than A. OK. OK. So let me summarize what we have seen here.

OK. So f submultiplicative lemma is sharp. Then we get sets A and B and C, which are subsets of real numbers. And in the main case, A is a delta A A set. It's rho A set. Size of A is like rho to the minus A. B is a rho B set. Size of B is like rho to the minus B. C is a rho C set. Size of C is like rho to the minus C. And for every C in C, A plus CB is not very much bigger than A.

OK. Now the size of A and B and C can all be traced back to the numbers in our problem. And we find that if A is smaller than B plus C, that's equivalent to saying that RAD is square root of delta is good.

OK. Sorry. A is larger than B plus C, then that means you do a boring computation. So there's an intermediate result. Maybe I'll call it a key lemma, which is called the ABC sum product lemma, which just goes into this paper of Orponen-Shmerkin. And it says if you have all of this stuff then, indeed, A is bigger than B plus C. Yeah?

AUDIENCE: What does good mean?

LAWRENCE GUTH: Good means that it obeys the theorem. So the theorem is up there. So it's smaller than the maximum of 1 and root delta to the minus T over 2 plus S over 2 and root delta to the 1 minus T. So if that's the case, we're happy. And otherwise-- yeah, OK.

OK, so let's go back to our general strategy, our big picture strategy. We would like to prove this theorem, that RAD of delta is always smaller than that. And we considered the submultiplicative lemma, which doesn't give us the theorem yet, but it's an important tool. And we'd like to say if we can have an epsilon improvement in the submultiplicative lemma, then that's great. We go to the next scale and we have a better bound. So you keep doing that until you stop.

And we must stop in a situation where there's no more, where the submultiplicative lemma becomes sharp. Maybe that's a better way to say it. We have this submultiplicative lemma. We use it. It's either sharp or it's not. If it's not, we get a nicer exponent at delta squared than we did at delta. We keep doing that. And the exponent is not going to keep getting better forever, so we stop at some moment where the submultiplicative lemma is basically sharp.

But the fact that the submultiplicative lemma is sharp gives us a lot of structure. It gives us the structure of being projective, and it gives us this kind of product structure. And it leads us to this problem, which looks a lot like additive combinatorics and which is very special problem compared to our original setup.

And then the bound we need to prove for this problem, if you just unwind the computation, we need to prove this bound for this problem. If you unwind it, it will tell us that RAD obeys the desired bound. Yeah. I guess rho is also-- rho Here is square root of delta. Yeah. OK.

OK. So this situation is now a lot more structured than what we started with. And it looks reminiscent of additive combinatorics, maybe of Bergen projection theorem and so on. And indeed, they were able to deal with it with that toolkit. I could say a little-- I don't want to write too many equations, so I won't say the whole thing. I could say a little bit about it. But I think that what we've done so far is what I feel is most worth remembering about the proof of this theorem. So let's pause here and see if people have any questions or comments about it.

Let me summarize it. So the AD regular case is special because all the scales are related to each other. And so we can break it up into problems at different scales and put them together. And by doing that, we get problems that are somewhat nicer. So first of all, if we look at the problem in a small ball instead of the original problem, it becomes nicer because the directions through any point are the same.

For any two points, the directions are the same because they correspond to the directions of the fat tube through there. So this already made it a little bit nicer. Then, moreover, in order for the submultiplicative lemma to be sharp-- so if the submultiplicative lemma is not sharp, then you just keep going down scales and you keep getting better exponents. And eventually that has to stop when the submultiplicative lemma is sharp.

When the submultiplicative lemma is sharp, there's even more structure because these different balls have to talk to each other. And so if you see in this ball, you see a tube here, then you have to see a tube in the same place in that ball. And that means that this set of points has kind of a product structure. And in particular, what's literally true is if you take these-- well, sorry.

So if you take the set of red tubes in this ball and then in this ball and then in this ball, that is a product which is very helpful. OK. Cool. So the AD regular case allows us to restrict to a situation which has a lot more structure, and it restricts to this lemma about projections of a product set. OK.

AUDIENCE: So in the argument, I can continue until the submultiplicative lemma is sharp. How do you know that kind of sharpness corresponds to the delta T epsilon factor, rather than just kind of something much smaller? Because it seems like-- I don't understand how you can get a quantitative bound out of the more qualitative proof by contradiction or-- yeah.

LAWRENCE GUTH: Yeah. OK. So this question is about keeping track of the small parameters. So every time I apply the submultiplicative lemma, I was using the word sharp or not sharp. But of course, there's some kind of a gray zone and I have to decide where I'm going to cut. What am I going to consider to be sharp and what is not sharp? And then, in this analysis-- yeah, OK.

Actually, so in order for this to stop in a reasonable amount of time, I'm going to have to have a cutoff that's a tiny power of delta. Otherwise this would go on too long, wouldn't work. So I would say, so what would the definition of sharp be? Let's do that. All right. Let's get rid of this set. So we're going to use the submultiplicative lemma repeatedly.

And so if our AD of delta is bigger than, I need to put something here, RAD of root delta squared. And here I'm going to put delta to the epsilon 0 where epsilon 0 is a tiny parameter. Then I'm going to call this sharp and I'm going to stop and analyze.

And the other possibility is that RAD of delta is less than delta to the epsilon 0 RAD of root delta squared. And here I'm making progress, so I'll keep going. So this is the not sharp case. So now I'll keep going. OK. So that's if we're more careful about what is sharp and not sharp.

All right. So here, if I keep going every time I make a definite amount of progress, and so this actually cannot keep going forever because eventually the exponents would be negative instead of positive. So it will stop, so we will end up here. Once we end up here, we have to analyze the sharp case. And we do have to remember, it wasn't perfectly sharp, but we had a fudge factor of a tiny power of delta. And that complicates all the pictures that we have drawn.

So in this picture here, this tube goes through many fat tubes. It doesn't have to have many delta balls in all of them. It has to have many delta balls in a fraction, like delta to the epsilon fraction of them. So it makes the picture a little bit more complicated. And actually, because of that, in this theorem, if we were honest, it's a little bit more complicated.

This would actually be pi C of x, where x is a subset of A cross B and it's like a delta to the epsilon subset. So we wouldn't have to fill the whole product. OK. So that's how we would be careful about sharp and how we would have to be more careful here.

OK. Let me tell you just the first idea of the proof of this key lemma. So we have step proof. So we have A plus CB is only a little bit bigger than A. That reminds us of Plünnecke-Ruzsa. So an appropriate version of Plünnecke-Ruzsa. Well, OK. We have this big subset x. We might want to get rid of that by doing Balog-Szemerédi-Gowers.

So a Balog-Szemerédi-Gowers would reduce to A prime plus CB prime smaller than A prime. We actually have all of it instead of just a big piece of it. And then Plünnecke-Ruzsa would give us-- so this would be true for many C in C. And that would give us A prime plus C1, B prime plus C2, B prime plus C3, B prime would still be smaller than A prime. You could get that from Plünnecke-Ruzsa.

And so now we're not done yet, but you have the feeling that we keep using these tools from additive combinatorics, and we keep having stronger and stronger sounding information about A and B, more and more suspicious. OK. And it definitely takes some skill and some good ideas to finish from here. But we're in a setting where the tools of additive combinatorics are quite powerful and they're able to analyze what happens. OK, good. Any final questions or comments? Yeah?

AUDIENCE: How did you reduce to the projective case here?

LAWRENCE GUTH: We reduce to the projective case at the beginning. So after we had the submultiplicative lemma, we reduced to the projective case by saying that this ball is in the projective case. And so using that, we were able to show that it suffices to prove it for the projective case.

AUDIENCE: But then how do you chain that argument in the projective case?

LAWRENCE GUTH: What's that?

AUDIENCE: Like, I don't understand how you use that kind of argument and the reduction to projective at the same time.

LAWRENCE GUTH: Ah. The question is about how we put together reducing to the projective case with all the other things that we did. Yeah. OK. So using this lemma, we had an argument that if you could prove a good bound for the projective case, then combining it with this lemma and noticing that this is projective, then we get a good bound for the general case. So now we just imagine we're trying to prove the theorem in the projective case. So we can assume that we're in the projective case. And then we start applying the submultiplicative lemma.

Ah. So you're wondering, when we do this-- yeah. So you can think of in this picture where, over here we have not just some abstract thing but this would be the R in the small ball and the R for the fat tubes. So we look at those, and if the R is better than multiplying the small balls times the fat tubes, then we keep breaking things up. We can take the small balls and break them into smaller balls. And otherwise, if they're comparable, then we look at what's happening. Yeah. That's a good question. Yeah?

AUDIENCE: So I understand what's going on. Is this example supposed to be similar to the case B that we had at the beginning where-- for instance, if we started with tubes and points on a grid, we kind of immediately fall just into the situation. And what we're doing is blowing things up until we get there.

LAWRENCE GUTH: Yeah. OK. So the question is how these pictures and these examples relate to the three examples that we had in mind at the beginning. So the three examples we had in mind at the beginning were there was a simple one, the stars example. There was the grid that was interesting, and there was the random example. And a funny thing is that none of those is AD regular.

Now, the random one, we could fidget to make it AD regular if we wanted to. And the stars, we could make it AD regular if we wanted to. But the grid is actually is really not AD regular. And so that is a little bit funny that in the most important example it's not in this case, even though this case is very important in the final proof. So you can't really start with the grid.

At least until recently, it was not known what is the actual worst behavior in the AD regular case. So it may be that RAD is much smaller than R, but the argument so far doesn't prove that. And it's a little bit funny because it means there's not an example that you can keep in mind when you go through the AD regular case. Yeah?

AUDIENCE: What goes wrong with trying to construct an AD regular grid where you just kind of have a grid, and each point has some nested subgrids within it?

LAWRENCE GUTH: Yeah. So the question is, could you make an AD-- yeah, I guess that's true. You could make a roughly somewhat AD regular grid by taking a grid of balls, and then in each one of those balls, you put a grid of balls and you repeat like this. Since the grid itself is not AD regular, this is only-- it's only AD regular with the big C. Could you do that? Yeah, I guess you could do that.

Yeah. Yeah. So you could do that. And then what would happen here, the A, B, and C would be like sets of integers except they would have this AD regular structure. So A would be, take an arithmetic progression of intervals, and in each one of them, put an arithmetic progression of small intervals and so on. That's what A and B and C would look like.

AUDIENCE: Like a Cantor set looked?

LAWRENCE GUTH: Yeah. It's kind of like-- it's like a Cantor set, but in order to make the numbers come out well, you need to have a lot of intervals in an arithmetic progression in each step. Here, I'll draw it, and then I'm over, so I'm going to stop. So the example is a big line. You make an arithmetic progression with a lot of intervals.

OK. Then you look inside of each one of these intervals and you repeat what I just did. And for this to be a kind of sharp, interesting example, it needs to be a lot of intervals. OK. OK, good. Sorry for going a little bit over. Those are great questions. Let's stop there and I will see you next week.