[SQUEAKING]

[RUSTLING]

[CLICKING]

LARRY GUTH: In just the last year, the sharp projection theorem in R2 was proven. And the proof is fairly big. It's fairly complicated. But for the last three classes, I'm going to want to circle back to that and talk about the new ideas in that proof. So it's been a while since we talked about it, so I thought we could start by-- we'll start just by writing down what is the statement.

So this projection theory, I like to think of it as a geometric measure theory analog of the Szemerédi--Trotter theorem. So I'm going to start by writing down the Szemerédi--Trotter theorem. So sharp projection. Start by writing down the Szemerédi--Trotter theorem.

And there are a bunch of different versions. And I've picked out a particular version that fits into the way the proof is going to go for this sharp projection theorem. So it says the following thing. So if E is a finite subset of R2, and then L is a set of lines, let's say it this way. For every x in E, we'll have LxE is a set of S lines that go through x. And L is the union over x at LxE. So this set of lines L, there's at least S of them through every point of E.

So then, conclusion is that the number of lines is at least the minimum of a couple of options. There's the size of E times S, or there's the size of E to the 1/2 S to the 3/2. We'll call this option A and option B.

Option A, it's very simple to imagine how this could arise. I have a bunch of points, set E. And for each point, I have S lines through it. And they're just all different from each other, so I get E times S lines.

Option B is more special, more remarkable, although not that complicated. It's an integer grid. And we talked about it a couple times in the beginning of the class. So an integer grid, that would make some sense.

So we might have a lot of lines through each point. But still, each line is going through many points. So the total number of lines is a lot less than E times S. And if you work it out, you get this.

So that's the Szemerédi--Trotter theorem. We saw early on that it implies a sharp projection theorem for finite sets in R2. And we're going to try to do some analog for delta balls in R2.

When we talk about delta balls, it matters how they're spaced. And the spacing condition we're going to talk about is the one that Pablo introduced, the delta SC condition. So I'll make a definition.

So if E is a subset of R to the D, we're going to say that E is delta S, C if-- so first of all, I'm going to imagine here that E is contained in a unit ball or almost unit ball. And then, if you take E, and you intersect it with a smaller ball, it's basically a new ball.

If you take E, and you intersect it with a smaller ball and take the delta-covering number of that, then that should be smaller than C times little r to the S times the delta-covering number of the whole set. So for any x and for any r which should be at least delta. So that's a delta S, C set. It's a spacing condition on the set E that says that it's not too concentrated into any smaller ball.

So here is the statement of the sharp projection theorem. So the theorem of Orponen and Shmerkin, Ren, Wang from 2024. So if E in R2 is a delta S, C set, then every x in E will have a set of tubes T sub x. This is a set of delta tubes through x. And sorry, E is going to be-- this is just a notation choice, but if I don't match my notes, I'll go bananas. It's a delta t, C set.

And so for each point, there's a set of tubes through it. And this set of tubes is a delta S, C set in S1. So let's make a picture. So we have the set E, which you may as well think of as a set of delta balls. So there's E. And then for each of these delta balls, it lies in a bunch of tubes.

And we'll repeat for all of these other guys. Maybe this tube keeps going and goes through there, and this tube there. So if this is x, those tubes there, that's Tx. And they point in different directions. You can think of the directions as being parametrized by S1. And they are a delta S, C set in S1. Yeah?

AUDIENCE: What if you use multiple tubes that are in the same direction, slightly offset of each other that-- I guess, less delta [INAUDIBLE]

LARRY GUTH: Yeah, so if you-- right, so if I had this point, and then I had those-- oh, I see. Yeah, we'll think of those as the same tube. OK. So the question was, if I have my delta ball, could I have two delta tubes going through it that are going in the same direction or within an angle delta of each other? I'm going to think of those as being basically the same tube.

Then T is the union over x of Tx which could be smaller than the sum of them because one tube might be in Tx for two different x's. And the question is, in terms of S and T, how big does this need to be?

So then the conclusion is that the size of T is bigger than-- OK, there's a little complicated stuff-- C epsilon delta to the epsilon big C to the minus order of 1. And then there's the minimum of k. So the E times S corresponds to delta to the minus T delta to the minus S.

So this corresponds to A. There's somebody that corresponds to B, delta to the minus T over 2 delta to the minus 3T over 2. That corresponds to B. And there is a new thing, delta to the minus 1 minus S, that corresponds to C. Sorry, it doesn't correspond to anything. There was no C up there, I will call it C. So who is C?

So C is a new phenomenon that has to do with the fact that these tubes and balls have a thickness, which we didn't have up there. So here, E could just be any set. And then T is a random set of delta to the minus 1 minus S or so tubes.

And the probability that a delta tube goes through a given point is delta. So a typical point Tx would essentially be a random set of around delta to the minus S tubes. And a random set-- there might be a couple that are close together, but on the whole, that would be a delta S set of tubes. OK, so that's where this new option comes from. Cool.

AUDIENCE: Professor?

LARRY GUTH: Yeah?

AUDIENCE: In option B, is one of the exponents supposed to be an S instead of t?

LARRY GUTH: Yes, thank you. That's supposed to be an S. Correct. So let's put these on top of each other for a second. So the size of E corresponds to the delta to the minus T, and big S corresponds to delta to the minus little S. And so those match up. And then C is something new. OK, good.

Now, I realized I forgot one hypothesis that I'm going to add. And then we'll pause, and you can look at these statements. So we needed one more hypothesis in our list of hypotheses, which is that this S should be bigger than 1. You see, if I just had one line going through each of these points, then I could just have one line and a bunch of points on it. And 1 is not bigger than this. So that's a special case.

And there is an analogous condition here. In our list of hypotheses, we need little S to be bigger than 0. This corresponds to having not just one tube through each of these points, but pretty many. So that's the statement of Szemerédi--Trotter and the parallel statement of this is the Furstenberg set conjecture proven just recently. So let's pause there to digest and see if there are questions or comments about the statement of this theorem. Yeah?

AUDIENCE: I'm still having a hard time understanding what's going on in case C. Would it be possible to draw a typical example?

LARRY GUTH: Yeah. Let's draw case C So E could go anywhere. And then I am going to draw a random set of tubes, a bunch of them. So here I go. So I drew these tubes, or I tried to do it, without thinking at all about where the points are. There was nothing clever going on.

But I drew a whole lot of them. And 1 over-- a delta fraction, if I draw a random tube, the probability that it goes through here is delta. So if I draw M times 1 over delta random tubes, I'll have about M of them going through this point. So let's say I do that. So I have about M tubes going through this point.

Now, what directions will they point in? Well, they'll point in random directions. And a random direction-- a random set is a delta S, C set, where's where delta to the minus S is the number of points. So that's why we have this. OK, other questions or comments?

So now let's start talking about the history of work on this problem. And going through the history of work on the problem will also be an excuse to recap a big section, big chunks of our class. So history of work slash recap of the class. OK.

So first, I'll do what I'll call classical bounds, which are essentially due to Kaufman and Falconer. And they were written up in the context of this theorem. Kaufman and Falconer were talking about projection theory. This is a little bit more general. But they were written up. Kaufman and Falconer were in the '60s and '70s, so a pretty long time ago, and then adapted a bit by Wolff.

The first classical argument is double counter. There are actually two slightly different ways of doing double counting. I'll tell you what they give. The first way gives that, if S is at least T, then the number of tubes is at least option A. And this is sharp.

So let's think about what it means. So S is at least T. That means that the number of directions through each point is bigger than the number of points. So let me draw a little picture to the side here.

So I have a few points, not very many. And then through each tube, I'm going to draw-- through each point, I'm going to draw a bunch of tubes. And the number of directions is bigger than the number of points. So I have a picture like this.

So what you see is that most of these tubes don't go through any other point besides the initial point. And that gives an easy estimate for the number of tubes, has to be at least the number of points times the number of tubes through the point.

There is another way to do double counting. And it says that, in general, so just if S is bigger than 0, which we've been assuming, then the number of tubes is at least delta to the minus 2S. So let me just leave this one as an exercise. We've done some double counting.

Then there was a Fourier method that we did in class. And the Fourier method showed that T was bigger than the minimum of two things. One of them is option C, And the other one is delta to the 1 minus S times option A. And so this number is smaller than 1. So this is smaller than A.

So this isn't sharp, unless we're so lucky that S is equal to 1, in which case, this is the minimum of A and C, and that's good. So this is sharp if S is equal to 1. So I won't redo this one, but this is just the Fourier method that we did in the third week of our class. So the classical methods are occasionally sharp, but usually not.

So another comment about this story, about why this is difficult, is that some of these statements are field-dependent. They're true over R, but they're false over C. But others are true in all fields. So if we go back to the beginning, the Szemerédi--Trotter theorem is true in both R and C, although it's not true in all fields. It is the false in some finite fields. Well, as written, it's false in all finite fields.

This theorem here is true over R, but it is false over C. And that counterexample has kind of haunted all of us who have been working on this. And one feature of these classical methods is that they don't distinguish. They work over R or C.

So to prove this theorem, we have to distinguish R and C. So we'll have to do something which, in some way, is different from these. And the next phase in the history begins with Bourgain's projection theorem. And it's a phase that I'll label as epsilon improvements.

So for quite a while, actually, there really was no progress on these classical methods. These classical methods are from the '60s and '70s. Wolff adapted them to slightly more general setting, but basically, there was no progress for a long time.

And phase II is the phase of epsilon improvements. And the dates of this phase are from around 2000 to very recently, whereas the Bourgain projections theorem, it implies, with a little work but not that much work, an improvement over the classical methods in a very particular situation. So it says, if T is 1, and S is 1/2, then this number of tubes is at least delta to the minus 1 minus some epsilon for that. So if you just plug in the classical bounds up there, it would give this. And there's an improvement, although only for a rather particular-- for one particular choice of the parameters.

So that was in 2000. And the next improvement is a theorem of Orponen and Shmerkin from 2021. And you can think of it as a generalization of this. It gives an epsilon improvement, but in not just for this special value of the parameters, but for lots of values of the parameters.

So they proved that, if T is bigger than S, which is bigger than 0, then the number of tubes is at least delta to the minus 2S minus epsilon. So delta to the minus 2S is what you would get classically by double counting. It takes half a page. And then, after 50 years and 50 pages, you can get this.

So in our class, we said a lot about the proof of the Bourgain projection theorem. I actually wouldn't say that we proved-- did every step, every detail. We haven't proven this, but let me just say something about the proof. So this is based on Bourgain projection theorem, plus a bunch of ideas.

And so an emotional reaction that I had to learning this and going through it with you is that it is hard work. And also, epsilon is really tiny. It seems like the cleverer we are, the harder we work, we keep making epsilon smaller.

So this was the second phase of the history, where people were able to make epsilon improvements. I could also mention other theorems not about this exact problem, but in a similar spirit. So for many different problems, it was possible to make an epsilon improvement by using Bourgain's projection theorem and by using developing ideas related to the ones we put there.

So it's hard work. Epsilon was really tiny. But one positive thing about it is that it did distinguish R from C. So both of these theorems are false over C. So we did start to do that. Cool.

So this theorem represents a third stage in the history, where we didn't just improve by an epsilon over the classical results, but we improved substantially. And in fact, people proved the sharp result. So sorry, I should mention, this is sharp. This is sharp for any S and T. Yeah?

AUDIENCE: So [INAUDIBLE] really sharp up to the delta to the epsilon factor?

LARRY GUTH: Yeah, it's sharp up to this. So it's sharp up to the delta to the epsilon. That's right. In the future, I'm going to abbreviate that mess by this because we won't be at the level of detail to have all of this anyway.

So the story I wanted to try to tell, to explore over the last three classes, is how is it possible to get from these epsilon improvements to sharp theorems? That's our theme for the rest of the class. And if I had to say it in a couple of sentences, what happened was people found a way where you can use an epsilon improvement theorem, and you can use it over and over again. And each time you use it, you improve a little bit the exponent in your bound until you bring the exponent up to the sharp value.

Now, it's not at all obvious how to do that. And that's what we're going to try to explain, not with all details, but to make it seem like a plausible approach that you could carry out. So any questions or comments before we start? Yeah?

AUDIENCE: Does doing that process make the C [INAUDIBLE] constant kind of correspondingly, I guess, tiny compared to the epsilon constant?

LARRY GUTH: Yeah, so I think the question is, as we use our epsilon improvement theorem many times, does that cause this to become remarkably bad? I think it does. I think that this is remarkably bad. OK, great.

So today, my first goal is I want to describe the simplest example where you can get-- where you can use an epsilon improvement theorem to prove a sharp theorem. And then after that, if we have time, I will outline the proof of the big theorem. OK.

So simplest example where epsilon improvement leads to sharp out. So this is a theorem of Beck from around 1980, which is about the incidences of points and lines. So it's in the setting of the Szemerédi--Trotter theorem. So it says, if E is a subset of R2-- and it's not all on a line. And in fact, I'll say at most half of it is on a line. So for any line L, L intersects E is at most half of it.

Then, the conclusion is there exists some point x in E so that-- all right, so let me make a definition. So I'll draw a picture. So this is E. So there's our set E. Here's a point x.

Starting at x, I will draw all the lines that start at x and go through some other points at E. So what's that set? So I'll call it LxE. So this is the set of lines so that lines L-- so that x is in L, and L intersect E is not just x, so it has size at least 2. Conclusion is there is a point in E so that the size of LxE is comparable to the size of E.

This is clearly a sharp theorem because the largest possible size that this could be is the size of E. Let me do a proof sketch. And we'll-- so we'll see that this follows from Szemerédi--Trotter. But we'll also see that we don't really need the sharp exponents in Szemerédi--Trotter. We just need to have something that's an epsilon better than a simple bound. That's why I've erased it, but than the double counting kind of bound.

So here we go. So the sketchy part of my proof sketch is that I will suppose that we are in the uniform case. Namely, for every x in E, I will suppose that Lx of E has size around x. So it's not really the case that they would all be comparable, but I'll just explain that special case.

So now, how many lines are there in total? We can do a double count to find the total number of lines. So we're going to double count pairs xL where x is in E, and L is in our set of lines. So our total set of lines is going to be-- all right, let's do this.

So how big is LE? Well, we could count the number of these pairs in two ways. We could count-- look at all the points of E, and each of them is in S lines. Or we could say, look at each of these lines, and how many points are in each line?

Well, the number of the size of S-- the size of E intersect L for a typical L and LxE would be the size of E divided by the size of LxE. So I have all the points in E. They are divided among the lines of LxE. Again, assuming it's uniform, they would all have around this size-- E over LxE, which is S.

So this is the number of points times they're each in S lines versus the number of lines times they each have this many points. So when we solve, we see that the E actually drops out. And we conclude that LE is around S squared.

So our goal is to have a lower bound on S. So we want to have a lower bound on LE. And the Szemerédi--Trotter theorem exactly gives us a lower bound on LE, so we can use it. So the Szemerédi--Trotter theorem tells us that LE is at least the minimum of S times E and S to the 3/2 times E to the 1/2.

But this LE is S squared. So if you compare these, we can get a lower bound for S. In this case, we would have S squared is bigger than SE. So S is bigger than E. And in this case, we would have S squared is bigger than S to the 3/2 E to the 1/2. And with a little algebra, that is, again, S bigger than E. OK. OK. Right.

Now, I claimed that we didn't need to have this exact formula here. Let me write down all that we needed to know. So what we needed to know is just that, if E is much bigger than S, then the size of LE is much bigger than S squared.

So that's a feature you can read off of this formula. If E is much bigger than S, then either one of these is much bigger than S squared. And that's all we needed to know because we know that LE is only S squared. So by the contrapositive, we could see that S is not much smaller than ES. It should be at least around E.

So this thing that we needed to know is very similar to the Orponen-Shmerkin theorem. E is much bigger. So now let's look up at that Orponen-Shmerkin theorem, says, if T is bigger than S, which is bigger than 0-- so T is measuring the size of our set E, and little S corresponds to big S.

So that first hypothesis is saying that the set is much bigger than the number of directions. And then the conclusion is that the number of lines is much bigger than S squared. That thing is this thing. So this matches the theorem of Orponen and Shmerkin, the epsilon improvement theorem.

So that suggests an idea that we could use that epsilon improvement theorem by Orponen and Shmerkin and put it into an argument like this one and prove some kind of generalization of Beck's theorem that might have a sharp conclusion, even though we stuck into it just an epsilon improvement. And that actually works. So I'm going to write it down now for us

So this is called the continuum. Or we could call it in our class the delta ball version of Beck's theorem. So this is a theorem of Orponen, Shmerkin, and Wang, in '23, I think. So we're going to translate this theorem into the language of delta S, C sets.

So we're going to say, if E is a delta, u set-- I'm going to leave out the C. Things will depend on C. But E is a delta, u set. So that's our set. Now I need an analog of this hypothesis that E is not too much in the line. And the analog says that E will be not too much in any thin rectangle.

So we'll say if R is a rho by 1 rectangle, then if you take E, and you intersect it with R, take the delta-covering number of that, that should be kind of smaller than the total delta-covering number So it should be rho to the A to smaller than constant rho to the A to smaller than the total delta-covering number. Maybe we can put this back here also.

So who is eta? At the very beginning of this story, we'll say, for every eta bigger than 0. So eta should be a positive number. So this is not super strong condition, but it's saying that E is not heavily concentrated into a thin rectangle

Then the conclusion is there exists an x in E so that the size of LxE is greater than the minimum of delta to the minus u and 1. So this would be-- sorry, not 1, delta to the minus 1. So what's going on here? Our original theorem would say that E is at least the size of E. And that's the size of E, so that's what it's doing here.

On the other hand, if I had every delta tube-- if LxE consisted of all the delta tubes through x, there would only be delta inverse of them. So it couldn't possibly be bigger than this. So we have it. OK. Cool. So that's the statement of the theorem. And it is, again, sharp because as we just explained, this couldn't possibly be bigger than this, and it also couldn't possibly be bigger than this.

So this theorem is sharp. It's also false over C. So the example over C would be that the set E could be R2, which is sitting in C2. And if you imagine R2, if you take any point in R2, and look at all of the lines, complex lines that go through that point, there aren't-- each one of them goes through many of the points of R2. And the number-- there aren't that many of them. It would violate this.

So this theorem is in the general spirit of projection theory, it is a sharp theorem. And it distinguishes R from C. And I believe it is the first such theorem in the literature. So we were trying for 20-plus years to distinguish R from C. For a long time, we could only do it with epsilon improvements.

But this is not just an epsilon improvement on what was known before. It's not just an epsilon improvement on what's true over C And it's a sharp theorem. Any questions or comments about the statement?

AUDIENCE: So the [INAUDIBLE] depends on eta?

LARRY GUTH: Yeah, so if you unpack this, it must depend on eta. I think it is a constant depending on epsilon and an eta, delta to the epsilon, and big C to the minus big 1. So it depends on eta. It's conceivable that this depends on eta. I'm not sure. Other questions or comments?

So the idea in the first place is to imitate this proof of Beck's theorem, proof that converts an epsilon improvement on some simple estimate for Szemerédi--Trotter into a sharp estimate for this Beck kind of thing. But when we try to do that, we will see that there is an issue with it. And when we flesh out the proof a little more, it will involve using the Orponen-Shmerkin theorem not just once, but many times at many different scales.

So let's sketch the proof. So the first plan is to imitate the proof of Beck. So let's carry that out over here, I'll put the Beck's proof up.

So attempt one. So let's say that LxE is a delta S, C set. Now, that's related to the cardinality of LxE. This is a theorem about the cardinality of LxE. What is the relationship? Well, it implies that the cardinality of LxE is at least around delta to the minus S, but it could be much bigger.

So an example where these would be about the same as if LxE was like an S-dimensional Cantor set. An example where they would be very different is if LxE was very clumped. So here's a piece of S1, and maybe LxE is clumped like that. So these dots represent LxE.

So if LxE was really clumped into a much smaller interval, then the S would actually have to be 0. This is not a delta S, C set, except for S equals 0. On the other hand, the cardinality of LxE in this picture is large. And in particular, it's much larger than delta to the minus 0.

So this S is some non-concentration condition. And non-concentration forces are set to be pretty big. But it could also be much bigger than that if it was kind of clumped.

So let's try to imitate our proof. I'm going to suppose that we're in the uniform case. So all of the LxE have the same S and the same size of LxE We can do our same double counting. That part doesn't change at all. And it gives us LE is at least LxE squared. OK. Sorry, is approximately LxE squared.

And so to get a lower bound on LxE, it suffices to get a lower bound on LE. So what happens here? So now we can apply the theorem of Orponen-Shmerkin, the epsilon improvement, which I just erased in a suboptimal use of blackboards.

But what it says in this case is that, if S is smaller than the minimum of u and 1, but bigger than 0, then we get LE is at least delta to the minus 2S minus epsilon. So this would be the classical bound. And then we have this. And then we're interested in LxE. That tells us that LxE is at least delta to the minus S minus epsilon.

So let's take stock. If S were equal to the minimum of u and 1, then LxE would be at least delta to the minus S. That would be the bound we want to prove. So in that case, we would be very happy. Now suppose S is in this range. Well, then we get this bound. This bound is a tiny bit better than this bound. But that's all.

This argument didn't work so nicely as that argument, because in that case, there was no distinction between cardinality and the value of S. In that case, this was the same as this. So we would get a contradiction at this moment. And we could conclude that, actually, S is the minimum of u and 1.

So we did learned something. If LxE were kind of spread out, then you might hope that the cardinality would be about delta to the minus S. And in that spread-out case, then we could get-- it would follow that S the minimum of u and 1.

So we see this corollary if LxE is a delta, S set, and the size of LxE is around delta to the minus S. Then we would get that the size of LxE is greater than around the minimum of delta to the minus u and delta to the minus 1, and we would be happy. But it's a big if. And we don't understand yet very well what might happen if these sets were a little more clumped. Cool.

So let me put up a summary of what we've learned so far. So lemma-- if S is bigger than 0, and a typical LxE is delta, S, C, then size of a typical LxE is at least delta to the minus S minus epsilon-- or sorry, yeah, if 0 is less than S, is less than the minimum of u, 1.

So the proof of this theorem, we don't have it yet. We have this lemma. The proof of this theorem is, by bootstrapping this lemma, using it many times at different scales and locations. So we are, effectively, using the epsilon improvement many times, not just once, and at many scales.

So let's do a sketch of this bootstrap iteration. So we're going to suppose that LxE, first of all, is uniform and also that it is delta S, C for some S which is bigger than 0. It's the start of our bootstrap. So because it's uniform, and it's a delta, S, C set, then it follows that, if you take its rho neighborhood, then this is a rho, S, C set for every rho.

And now we can apply our lemma. And the lemma tells us that the cardinality of LxE, the rho-covering number of LxE, is at least rho to the minus S minus epsilon. So rho to the minus S is what we get just from being a rho, S, C set. But actually, it has to be a little bit bigger than that by our lemma. And that's true for every rho.

And now, since we have a uniform set, and since these rho-covering numbers are a little bit bigger than expected, it tells us that LxE is a delta, S plus epsilon, C set. Actually, here I needed to put 0 is less than S is less than the minimum of u and 1. All of these are hypotheses of our lemma, so we use them here.

Now this is our starting hypothesis. So now we can iterate. So S keeps getting bigger and bigger, as long as it's smaller than the minimum of u and 1. And so it will approach the minimum of u and 1. And then we're happy.

So unlike in the proof of Beck's theorem, we had to use this lemma many times. We used it at many different scales. And also, we had to use it over and over again.

And each time we used it, we improved this exponent a little bit until the exponent approached the exponent that we had in the classical Beck's theorem. And so this is sort of typical of this new paradigm, that you can take an epsilon improvement result and use it over and over again to keep epsilon improving something until the thing reaches the sharp value.

Now, you might be wondering how we started this argument. We needed to start with an S That was bigger than 0. And not every set is a delta, S, C set for any S. So to start the argument, we observe that we had this non-concentration in rectangles condition-- non-concentration in rectangles condition here, that our set is not too much in any rectangle.

We haven't mentioned that hypothesis yet. We use it here to get started. So to start with, non-concentration in rectangles easily implies that LxE is a delta, eta, C set. Because so if E looks like so, there's not too much of E in any rectangle. There's our x.

Therefore, there aren't too many of these rectangles. You can't have-- if you had a lot of LxE in a small angular range, it would mean, you had a lot of E in a small rectangle. And that says this doesn't happen. OK, cool.

So now we have sketched the proof of the continuum Beck's theorem, the first sharp theorem in projection theory distinguishing R from C. So let's take a couple of minutes to digest, see if you have questions or comments. Yeah?

AUDIENCE: I guess the second step of the proof would be LxE rho. Is that the intersection of LxE with a kind of rho-radius ball?

LARRY GUTH: Sorry. So LxE rho is the rho neighborhood of LxE. Actually, you can also erase the rho. Maybe that's the best thing to do. So LxE is a delta, S, C set. And if it's uniform, it's also a rho, S, C set for all the bigger rho's.

AUDIENCE: So rho greater than delta?

LARRY GUTH: Yeah, for all rho greater. Yeah. Yeah, thank you. Thank you. Yeah. Yeah, so recap, so we don't-- actually, yeah, we don't need to mention this.

So say it again. So LxE is uniform set. If it's delta, S, C, then it's also rho, S, C for all the bigger rho's. And so then we get these covering numbers and then get it there. OK.

So let me share a personal reaction. So when this paper came out, first of all, I was very impressed because it was-- I had been wondering about this for many years. And it's the first sharp projection-type theorem which distinguishes R from C.

But I also felt kind of unsure whether this would open the door to prove sharp results for things like the Furstenberg set problem. Maybe it's a sharp result, but maybe it's about a kind of thing that somehow is more approachable.

It was well known classically that an epsilon improvement of something implied a sharp version of the incidence geometry, something here. And that's just not the case for Szemerédi--Trotter. Szemerédi--Trotter is not known to follow from an epsilon improvement on some classical thing. And so it wasn't clear to me if this would open the door to prove sharp theorems about all the things we most cared about.

So in the rest of this class, what I want to do is to give an outline of the proof of the Furstenberg set conjecture and the different ingredients that go into it, of which this will turn out to be an important one. Cool. So the first thing to say is that this proof breaks into cases according to the way the set E is spaced. This delta, t, C set, that definition, on the one hand, it's a good definition.

But under the umbrella of delta, t, C sets, there are several different, quite different, ways that the set E could be spaced. And we have to pay attention to them. And we have to treat them in different ways. So let's explain that.

So we're going to split into cases according to how E is spaced. So first of all, let's assume that E is uniform. And now I mentioned the word "uniform" above, but now I'll remember, recall exactly what is uniform.

So first of all, our delta is going to be capital delta to the m, where m is a large number. And uniform means that, for every j going from 1 up to m and for every q which is a delta to the jq, then, when I take my set, and I intersect it with q, and I cover it with cubes or balls at the next smaller scale, delta to the j plus 1, that is around Rj.

So I'll draw some pictures in a moment. The smallest this Rj could be is 1, and the biggest it could be is big delta to the minus 2.

Actually, yeah, let me say a little bit more. So the delta covering number of E That's going to be the product j equals 1 to m of Rj. And another thing we can say is, if I were to take the product j equals 1 to capital J of Rj, that would be the covering number of E at scale delta to the capital J.

And that should be-- so let's say-- so if E is a delta, t, C set, then this should be at least 1 over C delta j to the minus t. So that gives some information about these products. And being a delta, t, C set is equivalent to this information about these products. And so there are many ways that they could look. Let me put up a couple of examples which also should, hopefully, clarify everything on this board, and then we pause for questions.

So the first example that I want to put up is that all the Rj's are the same. So Rj is delta to the minus t for every j. And then you'll see that this will be OK. And I'd like to visualize this in two ways.

So the first way I'd like to visualize is as a tree. So in my visualization, I picked that delta is 4. No, delta is 16. No, sorry, delta is 1/4, and t is going to be 1. So in my visualization, delta to the minus 1 is going to be 4.

I'm going to have a tree. So that would be R1 guys coming out of it. And then each of these have R2, which is also 4, branches coming up. I won't draw that, but they're all the same.

Now, to the next stage of our visualization, we started with a square. And we cut it into smaller squares of side length delta. So that looked like this. This is stage 1. And R sub 1 is 4, so we took four, so maybe these four.

Stage 2, inside of each of those four squares, I took each of them, and I cut them into a 4x4 grid of smaller squares. And I took four of them. So I'll do the cutting carefully here. So maybe I took those four. And then down here, I took four. And over here, I took four. And over here, I took four. So our set looks something like that at stage 2. And then you could keep going.

So this is the first example. It's called the AD regular example, Ahlfors-David regular. Ahlfors and David are two people who worked on stuff that is related to this. And the word "regular," I think, here refers to the fact that all the scales are the same as each other. Now let me show you an opposite extreme.

So this one is called the well-spaced example. So in the well-spaced example Rj is as big as it could be, its maximal value, for some range of J. And then Rj is 1 for larger J. So as a picture, it would look like this. So we're going to visualize with big J as 1, big M is 2, like over there. So we'll have two stages.

And so our first stage is that R1 is 16. I won't draw it up. And then R2 is 1. So from each point here, we have just one going down. So the tree looks like this instead of like that, a lot of branching at the beginning and then no branching after that.

What does it mean our set looks like in space? Well, first, stage 1, I take my square, and I break it into 16 squares of side length 1/4. And at stage 1, I select all 16 of them. So the picture looks like this.

Now, at stage 2, so I have cut my square into 16 squares of side length 1/4, and I look inside each one of them. I cut that into 16 squares, the smaller size. And I pick one of them, and I do that in all the ones I selected at stage 1, which is all of them. So this set looks like this.

So this is called the well-spaced set. So I could fix a t. I could make a delta t well-spaced. I can fix t and delta. I can make a delta t well-spaced set or I can make a delta Ahlfors regular set. And they look rather different from each other. And there's also a gray area in between.

So this one is distinguished by saying that the points are the most separated. So if these points were like electrons that were sitting in this ball, and they repel each other, they would come to a position that's like this. It would be well-spaced.

And this one is distinguished by being-- this is the most clumped which is allowed for a delta t set. So this isn't the most clumped possible set. I could take these delta balls, and I could smoosh them all there. That would be more clumped.

But if it's a delta t set, it's not allowed to do that. And this is the most clumped that it could be, subject to the definition of a delta t cell. OK, cool. So there's a range of spacing conditions. And this hopefully gives a sample of them. yeah?

AUDIENCE: Is there some sort of energy condition or a quantitative measure of the clumping that would allow you to separate the two cases? Because they're both delta, S, C.

LARRY GUTH: Right. So the question, is there a quantitative measure of clumping or concentrating or spacing that would allow us to distinguish them? I think the nicest way to do that is this list of Rj's, or perhaps this list of this list of these numbers here. So it's not just a single number, but it is quantitative. And if you're willing to have a list of numbers describing your set, then it does quite a good job.

It's really a modern innovation to focus on that. In previous times, people wanted to describe the spacing of their set with one number, which would be like the dimension of the set or the energy or something. And when they were describing with just one number, then it was lumping together these two rather different behaviors. And it turns out to be important to un-lump to prove this theorem, at least with the proof that we have. OK, good.

So now I can write down the outline of the proof. So here is the outline of the proof of the main theorem, or let me say, of the Furstenberg conjecture. OK. So first main ingredient is a theorem of Orponen and Shmerkin from '24. And they proved that the Furstenberg conjecture holds in the AD regular case.

So this is based on and builds on the stuff-- builds on the stuff we talked about today. So it builds on the continuum Beck theorem of OSW, which, in turn, built on the epsilon improvement of OS, which, in turn, builds on the Bourgain projection theorem.

It's not at all obvious given the Beck theorem, just an ingredient. And so there are also some great ideas in that we'll talk about next time. But also, it involves some special structure of the AD regular case. We'll talk about that a lot on Thursday.

The second ingredient is a theorem of Guth, Solomon, Wang. So I was involved a little bit. It's is a slightly older theorem. It's from 2019. And we proved that the Furstenberg conjecture holds in the well-spaced case. And it's based on the Fourier method.

This is much less difficult than this. It's the Fourier method that we explored for a while with a little bit of extra something something. How do you take advantage of being really well spaced instead of just AD regular? OK. Good.

So now, when I presented these two possibilities, the way I had thought about them is that these are just two opposite ends of a whole gray spectrum. And so if you have a method that works in one end, and you have a method that works on the other end, I guess you might hope to have a general thing that sort of combines them both and works for everything. But there's still all this stuff in the middle.

And it turns out-- and I think this is surprisingly important in its own right-- that the proof of the theorem actually reduces to just those two cases. So the third ingredient is a multi-scale argument which shows that the Furstenberg conjecture actually reduces to the AD regular case and what they call the semi well-spaced case. So this is a small generalization of this case that can be handled by, more or less, the same methods using Fourier analysis. OK, cool.

So in each of these, I think there is something worth talking about and worth remembering, an idea that could be useful certainly in other problems in this field and maybe in other parts of math. We have two more classes, so I'll [INAUDIBLE]. My goal is, in the two classes, to say something substantial about each of these three ideas which fit together to prove the Furstenberg conjecture. OK, good. Any last questions or comments? Mm-hmm?

AUDIENCE: Within each of these individual cases, how much variation is there in the structure of the sets? Because it's not really [INAUDIBLE] kind of rearranging [INAUDIBLE] points within one of the boxes does a lot or does a little.

LARRY GUTH: Yeah. OK, so the question is-- so we've taken all the delta, t sets. We have then divided them according to their branching function. And one extreme is the AD regular and one extreme is the well-spaced. And then the question is, if you look among different sets with the same branching function, like among different AD regular sets, how different are they from each other?

Well, I will try to answer that by drawing a second AD regular set in blue that we can compare to the one in white. So I should pick four big squares, and in each of them, I'll draw four dots. I would say that the sort of spacing feel of these dots is the same in all the examples that are AD regular. On the other hand, actually, let me draw another example. This would be-- here's another example.

So something else is a little different about this example. And so yeah, so there are some more delicate things that could be different about the different AD regular examples, even though-- so what's the same about them is you give me a radius, your favorite radius. You stick down a ball of that radius and count how many points are in there. That will be the same for any AD regular set.

But there are delicate things. Actually, another very important, delicate thing that would be different is-- so let's now think about the well-spaced set. An example of a well-spaced set is an integer grid. That has some special properties. It has translation symmetry and stuff like that.

A typical well-spaced set will still have one point in each of these boxes, but it won't have additive structure. It won't have translation symmetry. It won't have-- and it won't have very many lines that go through many points. So those kinds of things are different among the different sets with the same spacing condition.

OK, well, let's stop there for today. And we'll talk more about these things on our last two classes.