[SQUEAKING]

[RUSTLING]

[CLICKING]

PABLO SHMERKIN: So let's recall where we were last time. So we are trying to prove Bourgain's projection theorem. I guess I will state it again later. But last time, we hopefully accepted the following fact as true. OK, so maybe if a subset of 0, 1 is-- OK, so given S is 0, 1, there exists eta epsilon positive depending on S. And if we have a set, which is an S delta, delta to the minus eta set, then there exists a little A in the set such that the size of A plus little a, a grows by delta to the minus epsilon compared to the size of a.

So Larry pointed out, OK, so initially, instead of this, we have some polynomial q of A, such as aA minus aA. So Larry pointed out to me that if you follow exactly what you did in the finite-field setting to go from q of a to this, you need to use the non-concentration condition that we are trying to prove. But it is possible to do it in a way, which is not circular. So OK, so I said, I was not going to cover all of the details.

So using Plunnecke-Ruzsa-- so Plunnecke-Ruzsa is the main tool, and suitable double counting, one can really prove this without using something that we haven't proved yet. OK, and the goal is to improve this by showing that A plus aA satisfies a better non-concentration condition that will allow us to iterate. So the goal is to show that A plus aA actually contains because it is not true that it always is, but it contains a delta S plus epsilon delta to the minus maybe let's call it eta prime set.

So the eta can get worse when one iterates. But one iterates finitely many times. So this is not serious. OK. OK, maybe for a different. So the epsilon is also going to change.

OK, so I'm basically going back to what I did at the end of last time. But I was rushing a bit. So I'm going to be a bit more precise, hopefully. Suppose that A is delta S delta to the minus eta square set. Well, in particular, it is an S delta, delta to the minus-- so I guess delta goes first. In particular, it is a delta S delta to the minus eta set because eta is a very small number. But why did I square it?

Because then, it is OK. And as we saw last time, if we take a subset of a which is uniform, delta m uniform for some good choice of delta, delta will have to be chosen small enough in terms of eta, but nothing else. Then, so we may assume let's call it delta uniform. Or delta depends on eta. So that basically, this doesn't change. Maybe one has to multiply by 2. But essentially, it doesn't change. Well then, this set, because it's uniform now, is also rho, is delta to the minus eta squared z for every who.

And then this implies that it is also a row is rho to the minus eta. So the problem here is that now, we have rho here, but we still have delta here. But we replace eta by eta squared so that we can have also rho here. And this will be for every who between delta and delta to the eta, right? Because we started with delta to the minus eta squared.

OK, we don't go all the way to 1, but almost. Eta is very, very small. So morally, this is almost 1. So from here to here, we use the uniformity. So one consequence of uniformity that I explained last time is that a uniform delta z is also a rho set for every rho. And this is just trivial inequality using that rho is at most delta to the eta.

OK, so this means that we can apply what we've already know for every scale rho between delta and delta to the eta, because it satisfies the assumptions. And another thing that I explained last time is that even though a priori, we only know that there exists an A so that A plus aA grows, by applying that to the set of exceptions, we get that this is true for almost every A in a strong sense.

So if we put all of these together, what we get is that-- OK, so maybe let's recall what I just said. So in fact, the set of A. OK, so let's work with some finite set of rhos, which are powers of delta again. So we have delta, delta squared, up to delta, to some maybe m prime. So that this is equal to delta to the eta. And recall that little delta is delta to the m. So m prime is less than m, but it is close to m.

And again, this is a finite set of scales, but not really finite, because OK. It is finite, but it grows with delta. But it grows logarithmically with delta. So m is logarithm of little delta in base, big delta. Big delta is fixed. So it's logarithmic in delta. And m prime is smaller. So it's also logarithmic in delta.

OK, so for each rho, we know that the set of A's so that A plus rho A doesn't grow. Sorry. little a, but measurable scale rho doesn't grow by rho to the minus epsilon has size at most rho to the eta times the size of A. And why is this true? Because if this was not true, then we could apply the statement above to the scale rho to the set of exceptions to this set, and get a contradiction.

OK, and this is very small in terms of delta. So rho to the eta is, at most, delta to eta squared. Sorry. Here, OK. OK, so here, we are counting a set of A's. So the only way in which we can really count a set of A's is by measuring some scale. Otherwise, it could be infinite. So we always have to choose a scale. So scale rho. OK, rho to the eta is, at most, delta to the eta squared.

And this means that-- sorry. This is wrong. Let's go back and fix it. So delta is too large. So delta is a constant. So I want to be away from 1, not away from 0, but away from 1. So the rhos that we consider are delta to the m, which is little delta, delta to the minus 1. And then we stop at delta to the eta. Sorry about that.

OK, because we are only looking at logarithmically many values of delta, OK, one has to work a little bit more because here, we are measuring things at scale rho. But rho is far away from 1. And there are only logarithmically many values of rho. So putting all together, what we get is that for all A outside of a small set, maybe something like a set of size delta to the eta, I don't know, cubed times the size of A. We get that A plus little aA grows for every scale in this family.

OK, and now, it looks like we are getting close because, well, we want to show that this contains delta S plus epsilon something set. And a necessary condition is that we have growth at every scale. So if you have a delta S plus epsilon something set, then this will be true. And if we knew that this set is uniform, we would be done, because for uniform set, the condition of growing at every scale is necessary and sufficient. But we don't know that this set is uniform.

So we are not done yet. So we are getting closer, but we are not done yet. Any questions? So this is more or less where we were last time. OK, so in order to be able to finish, we still have to work a bit. And there was a claim, that I was explaining last time, that we can improve this.

So here, we are projecting A times A. So this is pi A of A times A. So we can improve this to project an arbitrary dense subset of A times A. So claim, if g is contained in A times A. And it has size at scale delta. Let's say delta to the eta squared, or something like this, relative to A times A, then Pa of G already grows. And in fact, we will also want to apply it at many scales. Or maybe not. We'll see.

OK, let's start with scale delta. Then we'll see what we actually need. Then this already grows, let's say by epsilon over 2. OK, and I guess at this stage, we claim that there exists such an A. But once again, if there exists an A, then it is true for nearly all A by applying the fact that there exists an A to the exceptional set. So we can play this game as many times as we want. And we want to play it many times.

OK, so last time, we were in the middle of proving this. So let's start again. Assume that 1 is in A. This is just to avoid to-- OK, just for notational simplicity. It's not important. Well, if A equals 1 works, we are done because we are trying to show that some element of A works. So if the element 1 happens to work, we are happy. Otherwise, 1 doesn't work. And the fact that 1 doesn't work means that there exists G, which is dense in A times A, and such that pi 1 of G doesn't grow too much.

And this, OK. So all of these together are exactly the assumptions of Balog-Szemeredi-Gowers or rather the delta covering number of Balog-Szemeredi-Gowers. But as I explained a couple of lectures ago, Balog-Szemeredi-Gowers works exactly the same for delta covering numbers. So maybe let's recall Balog-Szemeredi-Gowers for delta-covering numbers if A at scale delta has size N, and there exists G contained in A times A such that the size of G at scale delta is at least 1 over K times the size of A plus A at scale delta.

And if we project in the direction 1, this G, this doesn't grow a lot compared to the size of A. If all of this is true, then there exists a set A prime contained in A, which is fairly dense. So the size of at scale delta is at least K to minus a constant times the size of A. And A prime has a small subset. The doubling constant for A prime is, at most, K to a constant power times the size of A prime. OK, this is Balog-Szemeredi-Gowers.

Well, here, we can take well, N is just the delta covering number of A, and we can take K to be the minimum between eta squared, delta to the minus eta squared, and delta to the minus epsilon over 2, which is really delta to the minus eta squared. You should imagine that eta is smaller than epsilon, even before squaring.

OK, so we apply by Balog-Szemeredi-Gowers and let A prime be the corresponding set. OK, and then A prime is dense in A, and has small subset.

OK, so A prime is a dense subset of A. A is a delta S something set. So A prime is also a delta S something set, because we saw that subsets of delta A something sets or delta S something sets where the constant. So basically, you have to multiply the constant by this. But A was a delta S delta to the epsilon square set. So this is still a delta S delta to the minus O of epsilon squared set.

So we can apply what we've know to A prime. So A prime is a delta S delta to the minus O of eta square set. So there exists A-- an A prime, sorry, such that A prime plus alpha A prime grows. By applying the fact from the beginning to A prime, it still satisfies the assumptions. We have lost a constant here, but this is harmless. Any questions so far?

OK, and now, we apply something that you've seen in the finite-field setting. And again, this works exactly the same way for delta covering sets. What you saw in the finite-field setting is that if you have a set with the smallest subset. So small pi 1 projection. And large pi A projection, then in fact, the pi A projection remains large if we pass to a dense subset of A prime just times A prime. OK, so this is another recall from the finite-field story.

If A prime plus A prime, let's say again, it grows by, at most, K. And A prime plus little A prime, let's say, so A prime has small doubling. But if you apply projection by A, it grows by something. Then in fact, this implies that for every G in A prime times A prime, which is dense, let's say dense with threshold 1 over K, everything measured with delta covering numbers, the projection is still large. And I think it's K to minus some constant. And the L deal basically survives.

Does this look like something you've seen before? OK, does it look like something that-- OK, well, go back to your notes, or otherwise, just believe me, this is true. Larry told me you've seen something like this, maybe with different letters. Very likely with different letters. But the point is have a set of small doubling, but large pi A projection. Then the pi A projection is large, even after passing through a subset.

OK, so it looks like we've almost won because look at the claim, and look at what we have here. OK, have we really won? Not exactly, because the claim was for A, and we got it for A prime, which is a dense subset of A, but it is not A. So one has to work a bit to really win. But you will have to trust me that one can win from here.

So the claim was that, what we obtain for A prime is true for A. And we got it for A prime, not for A. So we'll briefly explain how to get it for A. But I will not do all of the details, because otherwise, we'll be here forever. Maybe let's keep the objective inside. OK, we can upgrade. A prime to A.

And the idea is, well, either A prime is A, in which case, we win. Or if A prime is not all of A, but it is almost all of A, in the sense that A minus A prime is really small, we still win because so if the difference between A and A prime is really small, you just apply trivial bounds on the very small part, and we still win. So if A minus A prime is very small, very small maybe, of size less than delta to the plus eta, something like this, times the size of A, we are fine.

Otherwise, do the same. So apply the same argument to A minus A prime. If A prime is not very dense in A, then A minus A prime is still a delta S delta to the minus eta square set. And we can do the same and get a new A prime. So maybe the original A prime, we can call it A prime 1. And if we apply the same thing, we get an A prime 2.

Now at the same time, I have to play the game that if something is true for 1A, it's true for nearly all A because we found that there exists an A, and there exists an A prime. But in fact, there are lots of A's. But the A prime could depend on A. So one has to be a little bit careful. But one can play this game. And eventually, one gets, so all of these sets have large size because they are dense in A prime. Or is it A prime is dense in A?

So at some point, and they are all disjoint, they are all disjoint because we take away the previous ones when we apply this again. So at some point, we reach a step where we have exhausted almost all of A. Let's say something like this. And then we stop. OK, but I'm cheating here because-- and then we stop. And then basically, we win because if we have a set G which is dense in A times A-- OK, we still don't really win yet.

But if I said G is dense in A times A, it will be dense in AI prime times AG prime, maybe a different prime. OK, so one still has to work. OK, plus some work. Plus more work. And I'm basically ignoring the little a here. But it's OK to ignore the little a because everything you prove from one little a is true for nearly all little a. Maybe forget what I said, but OK, this is one thing you have to do. Just iterate, so iterate this. Here, we get a dense A prime. We want to get all of A, not a dense A prime. We iterate, plus some double counting. We get that plane.

Any questions? So one has to work here. So I haven't proved it. So I don't expect you to see the proof. So one thing you have to do in the proof is iterate. But you have to do more things, and double count carefully. But eventually, one gets that. So one is able to upgrade A prime to A. OK, so that claim is true. So that claim says that there exists an A. And once there is one A, there are many A's, nearly all A's, so that the pi A projection of any dense subset of A plus A grows.

And now, we are really close to the end of that goal. So A plus little aA is not necessarily uniform, but we know that it contains a uniform subset. And when we pass to a uniform subset, we can make it large for any subadditive function, set function mu. OK, so this is the mu we are going to use now. So mu of B is going to be, OK. So fix little a for which the claim holds.

So this is the claim star. So we fix a little a so that this holds. And then mu of B, and here, B is a subset of 0,1, so it's a subset of A really, is the delta covering number not of B, but of pi A inverse of B.

So this is subadditive, right? Because what is mu of B1 union B2? Well, if you take primitives, you get something which is contained in the union of the primitives. And then the delta covering number is subadditive.

AUDIENCE: Isn't pi inverse of E intersecting cross A?

PABLO SHMERKIN: Yes. Sorry. Thank you. So this is finitely subadditive, because delta covering number is subadditive. And we are just taking delta covering number together with the preimage, which is really subadditive. OK, so by lemma from last time, by the uniformization lemma from last time, there exists on A prime contained in A, which is uniform, sorry, not in A. Ah. There exists in B contained in a plus little aA.

So we are going to apply uniformization to a plus little aA. We are trying to show that this set contains a delta S plus epsilon set. So this is our ultimate goal. And we know that if this set was uniform, it would be enough to show that it grows at every scale to reach the conclusion. So we want to make it uniform, but we need to know that the uniform subset grows at every scale.

OK, so by uniformization, there exists a B in A plus little aA, such that this mu of B is larger than delta to the minus eta squared times the measure of A plus aA. And what is the measure of A plus aA? This is just the delta covering number of A times A. So this is the delta covering number of pi A inverse of B. So this is a-- sorry. Again, I forgot to intersect with A times A. So this is the G to which we can apply the claim.

So this is a dense subset of A times A. And we know that for every dense subset of A times A, the pi A projection grows. In particular, the pi A projection of this grows. But this is by A inverse. So if we apply pi A again, we land in B. And in particular, we land in A plus little aA. But in particular, we land in B, which is uniform.

OK, so I apply claim star to G equals pi I inverse of B intersection A times A, that satisfies the density assumption. And this is because uniform sets can be taken dense. So again here, I have to choose maybe a different delta than before, although I guess the same delta works because it's the same numerology. So this is true. We can take uniform sets, which are as dense in the exponential sense as we want. This is exactly what's going on here.

OK, and the conclusion is that the delta covering number grows. Sorry. Not pi A of B. Ah. Pi A of A times A intersected with B, I guess.

OK, so we are one step closer because now, we have a uniform subset of A times aA. And the delta covering number-- well, I guess B is a subset of these. So here, one will have just B. So the delta-covering number grows. We know that if a set is uniform, and the rho covering number grows for every rho, and in fact, we don't need every rho, it's enough to consider rhos, which are powers of delta. Then they are delta S plus the growth set.

So we just repeat. So here, we have something for delta. But then, OK. Let's see what is the right order to do this. So I want to claim that what we did for delta, we can do for any scale. So there is nothing special about delta. Because A itself is uniform, all of these that we did for delta, we can do it for any other scale rho. But the uniform set that we get can depend on the scale.

So what is the right order to do this? So the goal is to use the lemma from last time that says that if it is uniform, and the rho-covering number is large for every row, and in fact, one doesn't need every, every rho, it's enough to consider powers of the delta base in the uniformity, then it is a delta S set, where the S comes from the size of the rho-covering numbers. And OK, by using this idea for every rho, so we did it for delta, but the assumption is called for every rho, at least for every rho which is not very close to 1. So we can do the same.

The only danger is that we could get a different B for different rhos. We could get different B's. I'm not very happy about this. So how do we deal with this? One thing one could try to do, but I think we lose too much, is OK, delta is delta to the m. Now we want to go to delta to the m minus 1.

So we could just replace A plus aA by B. But then we will have another subset of B, and we are going to use this factor logarithmically many times. And that is true many times. So yeah. I don't want to do that. It would lose too much. OK, so there is something I'm missing. And I'm not going to figure it out right now. So because I told you that-- yes.

AUDIENCE: Can you prove the claim simultaneously for all scales, replace that delta by rhos, all the different rhos? And then I think that would fix it.

PABLO SHMERKIN: Yeah, you want to have the same G for every scale?

AUDIENCE: Yes.

PABLO SHMERKIN: Yeah.

AUDIENCE: Like if you could prove that stronger claim, then it will solve it.

PABLO SHMERKIN: Yes. But why can't we get the same G for all the scales? OK, because I told you I wouldn't give a complete proof, and just a sketch of some of the ideas, but here is a sketch of some of the ideas. And it's something technical that-- so really, the ideas are what I explained. And we're missing something technical. It's not fundamental. I would say just to summarize what's been going on, I still have to explain once we prove that goal, how to continue with the rest of Bourgain's original theorem.

But this is really the most difficult step where things really change from the finite-field setting. OK, so let me sort of briefly explain, again, everything that's been going on. So we know that A plus aA grows. But we want to show that it also satisfies a stronger non-concentration assumption, that the S in the non-concentration assumption also grows.

OK, first, we can take A to be uniform. If A is uniform, then we know that A itself is a rho S set for every rho. And this allows us to reach conclusions for every scale rho, which is something that we have to use. In the step that I'm missing as well. We used it before, but clearly, we have to use it again.

And then we know that A plus little aA grows at every scale. If it was true that a plus little aA was uniform, then we would win, because for uniform sets, we saw last time that if they are large at every scale, then they satisfy the corresponding non-concentration condition. So we have to take a large uniform subset of A, but large with respect to what? Well, large with respect to this. Because this is what allow us to use the fact that if we project something dense in A times A, then we grow.

And how do we know that if we inject something dense in A times A, then we grow? We have to combine two things. The first Balog-Szemeredi-Gowers. So for Balog-Szemeredi-Gowers we just take an arbitrary element of A, in this case, we took 1, but it can be an arbitrary element of A, we apply Balog-Szemeredi-Gowers.

So either that element already works and we are done for the claim star, or if it doesn't work, then we are under the assumption of Balog-Szemeredi-Gowers. We apply Balog-Szemeredi-Gowers, and then we end up with the set that satisfies the assumption of something else that you've seen in the finite-field setting, which is that if you have a very small subset, an expansion under pi A, then this expansion under pi A is robust under pass into subsets.

So in either case, we get that pi A of dense subset of A times A is large. And this allows us to use uniformity with this mu. So that's a bit of a summary of the main steps to show the goal. OK, plus technical details. I can recall right now, sorry about that. This implies the goal. The goal? That goal.

OK, why did you spend so much time towards this goal? Because now, we can iterate. We knew that if A is a delta S delta to the minus epsilon set, then A plus A grows. But not only it grows, it contains the delta S plus epsilon delta to the minus eta prime set. And then we can iterate. So A plus little aA is our new A.

So iterating A goes to A plus aA for every epsilon-- Maybe let's not call it epsilon. For every tau, there exists a polynomial that depends on tau. It's a very large degree if tau is close to 0 such that Q tau of A is a delta 1 minus tau. So the S can become arbitrarily close to 1. Delta to the minus some eta tilde, that depends on how many times you have to iterate set. And by is, I mean, contains. Yes.

AUDIENCE: You said that the epsilon is going to change [INAUDIBLE].

PABLO SHMERKIN: Yeah, I think instead of epsilon, it's epsilon over 2.

AUDIENCE: So but then if the epsilons are shrinking, how are you getting arbitrarily [INAUDIBLE]?

PABLO SHMERKIN: Yeah, that's a good question. So OK. Epsilon, in some sense, is shrinking. But OK, so there is one epsilon, which is the epsilon for the size of A plus little aA. And there is potentially a different epsilon if we want to find a delta S plus epsilon set inside a plus little aA. Those two epsilons, I think one is epsilon, the other is epsilon over 2.

But the point is that they depend on S continuously, both of them. So they can be taken to be continuous functions of S. Because they are continuous functions of S as long as S is less than 1 minus tau, S is bounded away from 0. So we reach 1 minus tau in finitely many steps. Maybe let's write this here because it's important. So epsilon is, or can be taken, continuous function of S. So remains bounded below as long as S is bounded away from 1.

So we can achieve this by iterating finitely many times. OK, and why is this good? Because this allows us to do instead of A plus aA, now we can do x plus aX, where x is potentially much bigger than a.

So the next step, and I think this step is the same as in the finite-field setting. So I'm basically just going to state it. If x is a delta t-- OK, x is 0, 1 still. And x is a delta. Maybe let's call it u. So u is going to be t over tau soon. Delta t to the minus epsilon set on A is 0,1. It's delta S. Maybe let's write eta here for consistency. Delta to the minus eta set. So now, instead of one set, we have two sets potentially of very different sizes.

And OK, they are both strictly bigger than 0, and strictly less than 1. And the interesting case is when S is less than u. So this is the case that we don't already know. Then there exists little A in A such that x plus little aX grows. It grows by some epsilon that depends on S and u.

OK, so just briefly, so the idea, and again, I think you've seen this in the finite-field setting. But the idea is that it's true for maybe not for A, but it's true for Q tau of A. And here, you have to take tau so that Q tau of a is bigger than x. So we can take tau, for example, to be 1 minus u over 2. Something like this. Then this is much bigger than x in size. So these are sets. But it has size much bigger than x.

Once the set of directions has size much bigger than x, we can do double counting and get growth with this instead of this. But once we get growth for a polynomial instead of A, we apply Plunnecke-Ruzsa one million times, and we go back down to A. OK, then apply Plunnecke-Ruzsa triangle inequality, double counting. And then it is also true for A.

So this part is really exactly the same. So we already had to go from Q of A to little a before. So initially, we proved that aA A minus aA A is bigger than A. And we use that to show that A plus little aA is bigger than A. So it's the same reduction here. And I think you've seen something very similar in the finite-field setting and modulo technical details. This part is the same.

But in order to get here, it's important that we can make this set bigger than x. This is what we wanted to iterate. Potentially, x is much bigger than A, but we can expand A by a polynomial so that it becomes bigger than x. And for this, we had to iterate. This is why we worried so much about iterating. OK. OK, the next step is to prove Bourgain's prediction theorem. OK, I'm not going to do the full proof, but just basically recall what you did in the finite-field setting, and explain, again, where one has to be a little bit careful.

But hopefully, by now, you will not be so worried about the part where you have to be careful. So proof of Bourgain's projection theorem. OK, now I regret calling this xx. But it's too late. So let x be in the unit ball of R2 delta t delta to the minus eta set, where eta is very, very small. And t is between 0 and 2. Let D be a set of directions, which is a delta S delta to the minus eta set. And what is the goal?

So what is the claim of Bourgain projection theorem? It is that there exists some a and D such that pi a of G exceeds the trivial bound by some epsilon. And the trivial bound is the square root of the delta. OK. Sorry. In the version that I stated, the size of x is delta to the minus t. So it matches the non-concentration condition. So here, we can just write delta to the minus t over 2. This is the square root of the size of x, which is the trivial bound. And we exceed it by some delta to the minus epsilon. And this is true for every G, which is dense in A times A.

OK, so this is Bourgain's projection theorem. This is what we want to prove. OK, so in the finite-field setting, I think to prove a similar statement, what you did is first consider three fix-- well, just pick three directions in D. In order to apply Balog-Szemeredi-Gowers. So again, it's similar to what we did before. Either things already work, or we can apply Balog-Szemeredi-Gowers.

Here, we have to be careful about how we choose the three directions, because if the three directions that we choose are very close to each other, this is going to be bad news. We're going to lose a lot. So well, even though this is 0,1, we can sort of change coordinates so that 0,1 and infinity are in D.

And this can be done in a way where we don't lose too much using that D is spread out. Because of this non-concentration conditions, there are three points in D which are far apart from each other. And that means that if we send these three points that are far apart from each other to 0, 1, infinity, so these are really the slopes. They are not the angles. They are slopes. So this can be done with a linear change of coordinates of distortion.

So the distortion is the norm of the corresponding matrix delta to the minus O of eta. So this distortion is, because eta is much smaller than epsilon. If we have a gain, well, we have to decrease the gain by what we lose in this change of coordinates. But this is mild.

AUDIENCE: [INAUDIBLE] A to be pi, 0, union pi infinity of x?

PABLO SHMERKIN: We are. We are going to look at the vertical and horizontal projections of x and the pi 1 of x.

AUDIENCE: Yeah, I'm just saying there's A cross A in the statement of the theorem.

PABLO SHMERKIN: Oh, sorry. Sorry. In x. In x. Sorry.

AUDIENCE: X cross x?

PABLO SHMERKIN: In x. x is in R2. In x. x is a two-dimensional object.

AUDIENCE: And the point is you want to try to prove this using A equals sub pi 0, union pi infinity x?

PABLO SHMERKIN: Yes. Yeah, that's the point. But yeah. So sorry. There was a typo here. So x is already two dimensional. And what we have been doing, we have a Cartesian product. Or here, we have a Cartesian-- OK. But [INAUDIBLE]. So OK, we want to apply this, which unfortunately, is going to be hidden.

OK, so x, this x, is going to be some-- the projection of x, horizontal projection of the x in R2. So it's really similar to the finite-field setting. So I'm skipping details. I'm going to continue skipping details, mostly focusing on the differences. So one difference is that we have to be careful with the three directions that we take. If they are too close to each other, the distortion will be too large, and it can be so large that we lose all the gain.

But because the set of directions is spread out by the non-concentration condition, we can do it in such a way that the distortion is controlled. OK, so now, x is contained in x1 times x2, where this is the horizontal projection, and this is the vertical projection. So if, because we are assuming that also one is in D, if one works, we are done. Because we are trying to show that something works. So the claim is that there exists an A and D with this property.

So if we are so lucky that one has this property, we are done. If one doesn't work, so this is very similar to what we did a bit earlier, but somehow, we have to do it twice. Once to do the iteration for A plus little aA, and then again, to conclude the proof. If one doesn't work, then there exists some G, which is dense in x1 times x2. OK, sorry. Why is it dense in x1? Ah. OK, actually, yeah. This is more complicated.

So one has to be careful. So OK. So OK. So I want to assume that x1 has size roughly square root of the size of x with the idea that if it had size bigger than the size of x, we would win. But in fact, in order to reach this conclusion, we have to apply the electromagnetic hours already, because the conclusion we want is not only that the projection of x grows, but the projection of every dense subset of x works.

So OK. So one really has to do is apply value of electromagnetic hours three times. So maybe for 0, infinity, and 1, in this order, either it works or we apply Balog-Szemeredi-Gowers to replace x by x prime, and then x double prime, and x triple prime. OK, I'm not going to do all of the details, but did something like this happen in the finite-field setting? I think it also, this part has to be done in the finite-field setting. So--

AUDIENCE: We didn't do this strong version in detail.

PABLO SHMERKIN: Oh, OK. We didn't do the strong version in detail. OK, also, I'm not going to do the strong version in detail because the 20 minutes I expected to take are turning into the whole lecture again. But I'm leaving today. So now, I really have to finish it over the next 20 minutes. So OK. So OK, either 0, 1, or infinity, or I guess 0, infinity, or 1 work, or in each case, we can apply Balog-Szemeredi-Gowers and pass to dense subsets of x of x and the corresponding projections of x.

So Balog-Szemeredi-Gowers-- OK. So let's take 0. If the horizontal projection has the property we want, we are done. Otherwise, there exists a dense subset of A with a small horizontal projection. But you can imagine that the horizontal projection is another projection. So it's the same thing we did last time, we did a bit earlier. And then we have to replace x by some x prime, and then by some x double prime, and then by some x triple prime.

So I'm not going to do it in detail. I'm sorry. Yeah, I see many of you are confused. And that's OK because one has to do it carefully. It's not obvious how to do it, but it can be done. OK, so eventually, after renaming the set back to x, we are in the following situation, this situation where, OK, let me just write it here, where the horizontal and vertical projections have size, at most, delta to the minus epsilon, t over 2 plus epsilon.

OK, I guess let's say we apply Balog-Szemeredi-Gowers twice. So we don't apply it to pi 1 yet. OK, just to see how this goes, so suppose that pi 1, so suppose that one doesn't work. Suppose that one doesn't work. And that means that pi 1 of G doesn't grow for some dense subset of x. And if it's a dense subset of x, it is a dense subset of x1 times x2. This is x1. This is x2. So this is x1. This is x2.

OK, then we're going to apply Balog-Szemeredi-Gowers. So here, you see that we are in this setting to apply Balog-Szemeredi-Gowers. So we're going to apply Balog-Szemeredi-Gowers. So it looks like after applying Balog-Szemeredi-Gowers lots of times, and I skipped the details, but you have to trust me that this can be done, we are in the setting where we can apply this and be done. But not so fast. What's the issue? Why not so fast?

So all of the part that I'm not explaining here, you also have to do it in the finite-field setting if you want to get the stronger statement that the projection of a dense subset grows, not a dense subset of the given set. If you want to prove that in the finite-field setting, you have to do this multiple Balog-Szemeredi-Gowers. So this is not a difference, OK, you haven't done it in detail in the finite-field setting. But this part is not a difference between the finite-field setting and the Euclidean setting.

But there is something which is a difference. So what is the issue? Why can't we just apply, oh, OK. We can apply this. But why do we have to be careful about applying this? Yes?

AUDIENCE: So if we apply Balog-Szemeredi-Gowers, we get A prime plus A prime. And then [INAUDIBLE] y cross y. That covers a lot of G. And then if I want to apply what's on the left board, I would need to know that y is a delta u something set.

PABLO SHMERKIN: Exactly. So here, we have a Cartesian product. OK, so one small issue is that instead of x1 times x1, we have x1 times x2. But this is not important. The issue is that x1 is a projection of x. And we are assuming that the size is roughly the square root of the size of x. But why is it not concentrated? In order to apply this and conclude, we need to know that x is not concentrated with u equals to t over 2. So it's very similar to the issue that we had when we needed to iterate.

So we knew that a plus little aA is large. But we didn't know that the non-concentration condition also improves. It's a similar issue here. We know something about the size of these projections, but a priori, we don't know anything about the non-concentrated. We need these projections to be non-concentrated. OK, so maybe ignore all of this, and just say the following. Using Balog-Szemeredi-Gowers many times in a clever way, we can reduce to the case where the set x is a product set. You can even assume it's a product set of something, a self product set.

So using Balog-Szemeredi-Gowers as many times, we can assume that x is x1 times x2. And then using Balog-Szemeredi-Gowers again, so Balog-Szemeredi-Gowers allows us to go back and forth between dense subsets and just everything. So this is what allows us to do. So we can assume that we are in the setting of a product set. And then we want to apply this.

But not so fast, because we don't know that the projections satisfy the non-concentration assumption a priori. So we have this problem again. So it's very similar to the problem we had before. And it is solved in a similar way. But it has to be solved. So in some sense, so you can think of it in this way. We say that t over 2 is a trivial between [INAUDIBLE] bound for projections.

What we have to show is that this trivial bound actually holds in the non-concentration sense. So here, 0, 1, and infinity are just generic projections. In fact, if we pick three random points in D, this will work, because three random points will be separated. Three random points will not be too concentrated because of this. There are not too many points in a single ball.

So if you randomly sample three points in D, the linear map that makes these points horizontal, and vertical, and diagonal, will have small distortion. So we need to know that the random projection of x is a delta T over 2 S set. In some sense, the goal is to show that the random projection is a delta t over 2 plus some gain S set.

But to prove that, we need to show that it satisfies the same thing, but without the epsilon, but the non-concentration version of that. And this can be done. It can be done in several ways. One way is a similar thing to what we have done today. Well, today and last time. But basically, go from growth to non-concentration by uniformizing everything carefully. And it can also be done by some careful double counting, similar to, but not exactly the same that you've done before.

So there are several ways of doing it, because it's like the easy case. But it has to be done. OK, so maybe let's write that down. OK, so we need a last fact. If x and D are as before, as in Bourgain's projection theorem, then there exists some A and D. And once again, once we know that there is one A and D, that means nearly all A and D have this property such that pi of x contains a delta t over 2 delta 2, maybe O of epsilon, O of eta set.

Again, t over 2 is like the trivial bound. But it's not so trivial in this case, because we are claiming some concentration for the projection. But because at the same time, for covering number, so for the covering number version of this is sort of obvious, because again, if we have one set, you look at two projections in more or less orthogonal directions, and most one of them can drop by more than square root of the size of the set. And this is true at every scale.

So you see that if we knew that this is uniform, we would be done. So it's really very similar to what we have been discussing in the rest of the lecture. So this can be proved using a similar strategy as for A plus aA. Or it can also be done without using uniformization by double counting. So OK. OK, maybe not uniformization. But if one doesn't use uniformization, maybe one needs some of Ruzsa's lemma. But OK, so it can be done.

But I just wanted to point out that this is yet another place where we need to worry about the fact that, so the difference between the Euclidean setting and the finite-field setting. OK. OK, and now, once we know this, we are really done, because using Balog-Szemeredi-Gowers three times, we can assume that x is a Cartesian product. And we just need to know that the prediction in a given direction grows, and it is given by this.

And then we apply double star. OK, so since there are five minutes left, let me summarize what's been going on. So what is the idea, the general idea of the proof, both in the finite-field setting and the Euclidean setting, first, we work with the set in one dimension, and shows that if the set is not already everything, then it's expanded by some polynomial. Then we simplify the polynomial using Plunnecke-Ruzsa.

Then we iterate. In order to iterate, one has to be much, much, much more careful in the Euclidean setting in order to show that the hypothesis of the expansion are satisfied when we apply the growth once. So we can iterate. But eventually, we can iterate. That means that we can expand the given set A to almost everything by some polynomial.

Once we can expand by polynomial to almost everything, we get this growth. It is still of some product type, but the difference is that x now can be much bigger than A. And the fact that x is much bigger than A is not an issue, because instead of working with A, we work with a polynomial applied to A, which has size bigger than the size of x. And finally, we want to project something which is not a product, but by looking at three projections to begin with, which are far away from each other, and applying Balog-Szemeredi-Gowers many times, we can go back to this setting.

So that's the short version of the proof. And they are, again, in the Euclidean setting, one has to be careful to show that the projections satisfy the assumptions needed to apply this fact. So really, you need at least some concentration on x, because again, if x was a small interval, then this never grows for any A. So one needs some assumption on x. So one needs to prove something like that. But OK, one can prove it, and then eventually, one uses this. Any final questions?

AUDIENCE: How many times have we used BSG?

PABLO SHMERKIN: So either four or six. I'm not completely sure. No. I think the first time, we already have a product. Yeah, I think once, for A plus little aA, and then three times to reduce to this setting.

AUDIENCE: Is it for when we just want the projection of x itself is big in one direction? Or it's like for the proof that's [INAUDIBLE]?

PABLO SHMERKIN: Yeah, OK. So in the finite-field setting, as you are suggesting, it's much easier to prove that the projection of x itself is big than proving it for dense subsets. But in the Euclidean case, even if your goal is proving that the projection of x is large, and you don't care about dense subsets, you are still forced to care about dense subsets because you need the dense subsets to prove the non-concentration assumption when you iterate.

So how did we prove that A plus little aA satisfies the stronger non-concentration assumption? By proving that the projection of dense subsets of A times A are large. So even if your ultimate goal doesn't involve looking at dense subsets, you still need to look at dense subsets along the proof. And you cannot avoid that. Or maybe you can, but at least with this method of proof, you cannot avoid that. Yes.

AUDIENCE: Can you say a little bit more about how to handle the issue that the uniform subset at different scales might be different?

PABLO SHMERKIN: Yeah. So unfortunately, I didn't realize this issue would arise. So I didn't check carefully.

AUDIENCE: Oh, OK. That's fine.

PABLO SHMERKIN: But yeah, it's a very good question. Yeah, OK. So I'm not exactly sure, but let me say something that might work in this case, but it certainly works in other cases. So I explained an idea superficially of first getting an A prime. And if that A prime doesn't exhaust A, we take it away, and find another A prime. One can do something similar for uniform sets. So rather than taking one uniform subset, so one is given a set, and one knows nothing about that set.

And we can nearly exhaust that set by a finite union of uniform subsets. So one finds a dense uniform set, takes it away, finds another uniform set, takes it away. And in this way, we essentially cover the set with a very small error by uniform subsets. So this is one way to get around the issue that one uniform subset itself is maybe too sparse. And then a different scale is different, things could happen. So I think that could be one way. I don't think that's what we do in the paper. But I think that's one thing you could do to bypass this issue.

AUDIENCE: You said last time that there was possibly a way to simplify this [INAUDIBLE].

PABLO SHMERKIN: We are not sure yet. We are hopeful. Yes. Yeah, again. So Bourgain's original proof is different in the sense that he didn't construct an expanding polynomial at all. But it is similar in the sense that he proves this, and then uses the many Balog-Szemeredi-Gowers to get the full statement.

But to prove this, he did something different. But OK. Other claim is more difficult than what I explained. But I don't know. Yeah, now, we have this idea from last week that, yeah, might be a-- so basically, it would be a much easier way of doing this. To go from this to the full statement, the only way we know how to do it is how Bourgain did it, by doing the many Balog-Szemeredi-Gowers.