YUFEI ZHAO: OK, we are still on our journey to proving Freiman theorem. Right? So we've been looking at some tools for analyzing sets of small doubling. And last time, we showed the following result, that if a has small doubling, then there exists a prime. It's not too much bigger than a, such that a big subset of a, of at least 1/8 proportion, a big subset of a is Freiman 8-isomorphic to a subset of this small cyclic group.

So last time, we developed this tool called of modeling lemma, so Ruzsa's modeling lemma that allows us to pass from a set of small doubling, which could have elements very spread out in the integers, to something that is much more compact, a much tighter set. That's a subset, a positive proportionate subset of a small cyclic group.

And remember, the last time we defined this notion of Freiman 8-isomorphic, Freiman isomorphism, in this case, it just means that it preserves partially additive structure. It preserves additive structure when you look at most 8-wise sums. All right. Well, this is where we left off last time. If you start with small doubling, then I can model a big portion of this set by a large fraction of a small cyclic group. All right.

So now, we're in the setting where we are looking at some space, a cyclic group, for instance. And now, we have positive proportion, a constant proportion subset of that group. And we would like to extract some additional additive structure from this large set.

And that should remind you of things we've discussed when we talked about Roth's theorem. Right? So Roth's theorem also had this form. If you start with Z mod n or in the finite field setting, and you have a constant proportion of the space, then you must find a three-term arithmetic progression. In fact, you must find many three-APs.

So we're going to do something very similar, at least in spirit, here. We're starting from this large proportion of some space. We're going to extract a very large additive structure, just from the size alone.

So let me begin by motivating it with a question. And we're going to start, as we've done in the past, with a finite field model, where things are much easier to state and to analyze. The question is, suppose you have a set a, which is a subset of f2 to the m. And a is an alpha proportion of the space where you think of alpha as some constant.

Question is, OK, suppose this is true. And must it be the case that a plus a the subset-- all right, so a itself, just because it's a large proportion of the space. So just because it's the 1% of the space, doesn't mean that it contains any large structures. It doesn't contain necessarily any large sub-spaces, because it could be a more random subset of f to the m.

But there's a general principle in additive combinatorics or even analysis where if you start with a set that is quite large, and it might be a bit rough, a is a bit rough, it's all over the place. If you add a to itself, it smooths out the set. So a plus a is much smoother than a.

And the question is, must A plus A contain a large subspace? And here, by "large," I mean the following. Because we're looking at constant proportions of the entire space, by "large," I would also want a constant proportion of the entire space. So does there exist some subspace of bounded codimension? So if alpha is a constant, I want a bounded codimensional subspace that lives inside A plus A.

It turns out the answer is no. So there exists sets that, even though there are very large and you add it to itself, it still doesn't have large subspaces. So let me give you an example.

And this construction is called a nivo set. So let's take A sub n to be the set of points in F2 to the n who's Hamming weight-- so Hamming weight here is just the number of 1's or it's a number of non-zero elements-- number of 1's in x or, in general, the number of non-zero elements among the coordinates of x, so the number of non-zero coordinates.

And I want the Hamming weight to be less than this quantity here. So visually, what this looks like is I'm thinking of the Hamming cube placed so that the all-zero vector is here, the all-ones vector is up there. And then it's sorted by Hamming weight. So this is called a Boolean lattice.

And I'm looking at all the elements, A, which are within a Hamming ball of the 0 vector. So this is the set. It's not too hard to calculate the size of the set, because I'm taking everything with Hamming weight less than this quantity over here by central limit theorem. The number of elements in the set is of the form a constant fraction of the entire space, where alpha is some constant if c is a constant.

So it has the desired size. But also, A added to itself consists of points in the Boolean cube whose Hamming weight is at most n minus c root n. And I claim that this sumset does not contain any subspace of dimension larger than n minus c root n.

So this is the final claim. It's something that's, again, one of these linear algebraic exercises

that we've actually seen earlier when we discussed the proof of a cap set, right, so the polynomial method proof of cap set. If you have a subspace of dimension greater than some quantity, then you should be able to find a vector in that subspace whose support has size at least the dimension.

OK. So you see, in particular, we do not have any bounded codimensional subspaces in this A plus A. So even though the philosophy is roughly right that if you start with a set A and you add it to itself, it smooths out the set, it should contain-- we expect it to contain some large structure. That's not quite true.

But what turns out to be true, and this is the first result that I will show today, is that if you add A to itself a few more times, that, indeed, you can get large subspaces. And this is an important step in a proof of Freiman's theorem. And this step is known as Bogolyubov's lemma.

So Bogolyubov's lemma, in the case of F2 to the n, says that if you have a subset A of F2 to the n of fraction alpha of the space, then 2A minus 2A contains a bounded codimensional subspace, so a very large subspace. so 2A minus 2A contains a subspace of codimension less than 1 over alpha squared.

Here, I write 2A minus 2A even though we're in F2. So this is the same as 4A. But in general-- and you'll see later on when we do it in integers-- 2A minus 2A is the right expression to look at. So 2A minus 2A, something that works in every group. But for F2 to the n, it's the same as 4A.

So the main philosophy here is that adding is smoothing. You start with a large subset of F2 to the n. It's large. Does it contain large structures? Not necessarily. But you add it to itself, and it smooths out the picture.

So it has a rough spot. It smooths it out. And if you keep adding A to itself, it smooths it out even further. You add it to itself enough times, and then it will contain a large structure and just from the size of A alone.

And there is a very similar idea, which comes up all over the place in analysis, is that convolutions are smoothing. So you start with some function that might be very rough. If you convolve it with itself and if you do it many more times, you get something that is much smoother.

And in fact, adding and convolutions are almost the same things. And I'll explain that in a second. So this is an important idea to take away from all of this.

So when we do these free analytic calculations-- so there will be some free analytic calculations-- the first time you see them, they might just seem like calculations. So you push the symbols around, and you get some inequalities. You get some answers. But that's no way to learn the subject.

So you need to figure out, what is the intuition behind each step? Because when you need to work on it yourself, you're not just guessing the right symbols to put down. You have to understand the intuition, like why, intuitively, each inequality should be expected to hold? And this is an important idea that adding is smoothing and convolution is smoothing.

All right. So let me remind you about convolutions. So recall, in a general abelian group, if I have two functions, f and g, on the group-- so complex value functions-- then the convolution is given by the following formula. So that's the convolution.

And it behaves very well with respect to the Fourier transform. The Fourier transform turns convolutions into multiplications. So this means, point wise, I have that.

Convolutions also relate to some sets, because if I-- and this is the interpretation of convolutions that I want you to keep in mind for the purpose of additive combinatorics. If you have two sets, A and B, then look at the convolution of their indicators. It has an interpretation.

So if you read out what this value says, then what this comes out to is 1 divided by the size of the group over the number of pairs in A and B such that their sum is x. So up to normalization, convolution records some sets with multiplicities. So the convolution tells you how many ways are there to express x in terms of a sum of one element from A and another element from B.

And in particular, this function here is supported on the sumset A plus B. So this is the way that convolutions and sumsets are intimately related to each other. So that's proof Bogolyubov's lemma.

We're going to be looking at this sumset, which is related to the following convolution. So let f be this convolution of indicators, AA minus A minus A. Of course, in F2 to the n, you don't need to worry about the minus signs. But I'll keep them there for future reference.

Here, by what we said earlier, the support of f is 2A minus 2A. It's not too hard to evaluate the Fourier transform of f, because Fourier transform plays very well with convolutions. So in this case, it is the Fourier transform of A squared, Fourier transform of minus A squared.

The Fourier transform of minus A, if you look at the formula, is the conjugate. It's a complex conjugate of the Fourier transform of A. So again, in F2 to the n, they're actually the same. But in general, it's the complex conjugate. So we always have this formula here.

So we always have that. So by the Fourier inversion formula, we can write f in terms of its Fourier transform. Here, I'm using that. We're now using that. We're in F2, so that's what the inverse Fourier transform looks like.

And so we have the following formula for the value of f in terms of the Fourier transform of the original set. OK. We want to show that f, whose support is the 2A minus 2A we're interested in-- we want to show the support of f contains a large subspace, a small, codimensional subspace.

So observe that if f has positive value, then-- if f of x is positive, then x lies in its support. So we just want to find a large subspace on which f is positive.

But we can choose our subspace by looking at Fourier coefficients according to their size. So what we can do is let R be the set of, essentially, Fourier characters whose corresponding value in the Fourier transform is large. So it's at least alpha to the 3/2. And that value will come up later.

So let's look at this R. And what we're going to do is we're going to look at the orthogonal complement of R and show that f is positive on the orthogonal complement of R. First, R is not too large. The size of R is, I claim, less than 1 over alpha squared. Why's that?

So this is an important trick that we've seen a few times before. The number of large Fourier coefficients cannot be too large, because they're Parseval, which tells us that the sum of the squares of the Fourier coefficients is equal to the L2 norm of the original function, which, in this case, is just the density of A. So that's alpha. So just looking at that, the number of large terms cannot be too many.

OK. So we have this small set, R, on which f has large Fourier transform values. Now, let's look at f of x. So let's look at f of x. We want to find out, when can we control f of x to make sure it is positive?

Well, for the values of r-- little r-- not in big R or 0, we see that the Fourier transform-- we would like to upper bound this quantity here so that this is negligible. This is a small term. Again, this is a computation that we've seen several times earlier in this course.

All of these terms are small. So I want to show that the whole sum is small. I don't want to bound each term individually and then sum up all the possible contributions. That will be too big. But we've seen this trick before where we just take out some subset of the factors.

So in particular, I'll take out two of the factors and get alpha cubed upper bound plus the remaining factors. And once again, use Parseval on this very last sum, keeping in mind that I'm throwing away some of the r's, including 0. So it will be a strict inequality.

OK. Yeah. So this step should be reminiscent of very similar computations that we did in the proof of Roth's theorem. So if x lies in the orthogonal complement of uppercase R, then f of x-- well, let's evaluate f of x from the Fourier inversion formula. We have this.

So I can now split the sum as the 0-th term, the large terms. Now, you see, for the large terms, because we're in the orthogonal complement of A, I can make sure that they all come with a positive sign. And finally, the small terms.

And you see that the main term is alpha to the 4. This term is always non-negative. And the error terms, the small terms, are strictly less than alpha to the 4th in magnitude. So as a result, this whole sum is positive. Yeah.

AUDIENCE: These 1's are also 1 sub A's, right?

YUFEI ZHAO: Thank you. The 1's are 1 sub A's. Yeah. So this is the very similar philosophy to when we proved Roth's theorem. We look at a sum like this, so some trigonometric series, some Fourier series. And we decompose it into several terms based on how large their Fourier coefficients are.

We can control the small ones using what essentially amounts to a counting lemma and show that the small ones cannot ever annihilate the large, dominant terms. So as a result, f of x is positive on the orthogonal complement of R. So thus R lies in the support of f, which is equal to 2A minus 2A.

And furthermore, the codimension of R is at most-- so it could be some linear dependencies--

is at most the size of R, which is strictly less than 1 over alpha squared. And that proves Bogolyubov's lemma. So if you have a large subset of F2 to the n, you add it to itself enough times so that it's a smoothing operation. And then eventually, you must find a large structure.

And we only start by assuming the size of it. If it's just large enough, then we can find a large structure within this iterated sumset. Any questions? Yeah?

AUDIENCE: Isn't R in support of R [INAUDIBLE]?

YUFEI ZHAO: Sorry, come again?

AUDIENCE: You got that the orthogonal complement of R--

YUFEI ZHAO: Sorry. The orthogonal complement of R is in the support. Yeah. So R lives in the character space. OK, great.

So this is the proof of Bogolyubov's lemma in the finite field setting, working in F2 to the n, which is fine. It's a useful setting as a playground for us to work in. But ultimately, we want to understand what happens in the integers.

So if you look at where we left off last time, we started in the cyclic group, Z mod n. So we would like to know how to formulate a similar result but in the cyclic group where there are no more subspaces.

We encountered a similar situation, although we didn't go into it, when we discussed Roth's theorem. In the first proof of Roth's theorem that we showed, in the first Fourier analytic proof in the finite field setting, the proof won by restricting to subspaces, to hyperplanes. And then we keep on iterating by restricting to hyperplanes. So you can stay in subspaces. And the finite field setting has lots of subspaces.

And we said that to get that proof to work in the integers, we had to do something different. And we did something by restricting to intervals. But I also mentioned that, somehow, that's not the natural analog of subspaces. The natural analog of subspaces is something called a Bohr set. And so I want to explore this idea further now.

So the natural analog of subspaces in Z mod n are these objects called Bohr sets. And they're defined as follows. So suppose you are given some R, a subset of Z mod n. We define a Bohr set, denoted like this, so Bohr of R and epsilon, to be the subset of Z mod n, so including

elements x, such that rx is pretty close to a multiple of n.

So here, we're looking at the R mod Z norm. So this is the distance to the closest integer such that this fraction is very close to an integer for all little r and big R. You see, this is the analog of subspaces, because in the finite field setting, the finite field vector space, even if I set epsilon to equal to 0 and turn this into an inner product, then Bohr sets are exactly subspaces-- namely, the orthogonal complement of the set R.

But now we're in the integers, where you don't have exact 0. But I just want that quantity, that norm, to be small enough. So let me give you some names. So given the Bohr set, which, technically speaking, is more than just the set itself but also includes the information of R and epsilon-- so it's the entire data written on the board-- we call the size of the R the dimension of the Bohr set and epsilon, the width.

Bogolyubov's lemma for Z mod n now takes the following form. If you start with a subset A of Z mod n, and all I need to know is that A is a constant fraction of the cyclic group, then the iterated sumset 2A minus 2A contains some Bohr set Bohr R of 1/4 with the size of R less than 1 over alpha squared.

So earlier, we said that if you have a large subset of F2 to the n, then 2A minus 2A contains a large subspace. And now we say that if A is a large subset of the cyclic group, then 2A minus 2A contains a large Bohr set of small dimension.

And so this terminology may be slightly confusing. The dimension corresponds to codimension previously. So if you do this translation, this dimension-- I mean, if R were a set of independent vectors and you have 2 to the n, then that'll be the codimension of the corresponding subspace. But this is the terminology that we're stuck with. OK. Any questions about the statement?

You see, even the bounds are exactly the same, 1 over alpha squared. And I mean, the proof is going to be pretty much exactly the same once you make the correct notational modifications. So we're going to do that. So I'm going to write on top of this earlier proof and show you what are the notational modifications so that you can get exactly the same result here but with Bohr sets instead of a subspace.

The thing to keep in mind is that we have a somewhat different Fourier transform. So let me now use different colored chalk. So the Fourier transform of a function f from Z mod n, so a

complex value, is a function also on Z mod n defined by f hat of r equal to expectation over x in Z mod n of f of x times omega to the minus rx, where omega is a primitive n root of unity.

And you also had the Fourier inversion formula. It's what you expect. I won't bother writing it down.

So we go back to the proof. And pretty much everything will read exactly the same. So f is still the same f. And the Fourier transform has the same property. So all of these nice properties of the Fourier transform hold.

For inversion, it's basically the same except that the formula is slightly different. So instead of minus 1 to the r dot x, what we have now is omega to the rx. So here, we have omega to the rx. OK. Great.

The next part is the same, where we define r. So now we define r to consist of elements of z mod n, whose Fourier transform is large. I can take out 0. OK. This part is still the same. It's the same calculation.

Now, it's the very last part that needs to be just slightly changed. Where does the 1/4 come in? So where does this come in? So observe that if x is in the Bohr set with width 1/4, then rx divided by n is-- OK, so by definition, all of these fractions are within the 1/4 of an integer.

And if you think about what happens on the unit circle, if you are within 1/4 of the integer, then that means the corresponding place on the unit circle is on the left half circle. So in particular, the cosine of 2rx over n is non-negative. So it has non-negative real part.

So now we go back to this part of the proof, where we're applying Fourier inversion formula to f of x. So we had the Fourier inversion formula up there. But because f of x is real, it's really the cosine that should come in play. It should be a cosine.

And now, for the next step, we have no negative sign here, because this step-- OK, let me just cross out this step over there. All of these terms, the terms that correspond to little r and big R, they have non-negative contribution. Whatever the contributions here, it's non-negative.

So I cross out this term. All I'm left with is the main term, corresponding to the density, and the error term, so to speak, the minor terms, which is less than alpha to the 4th in absolute value. OK. So it's positive. So basically the same proof. Once you make the appropriate modifications, it's the same proof in Z mod n.

OK, great. So this concludes our discussion of Bogolyubov's lemma. So it says that-- OK, so continuing our previous thread, we start with a subset of z mod n of constant proportion. Then 2A minus 2A necessarily contains a large Bohr set.

And the next thing I want to do is to start with this Bohr set. So that's the definition of a Bohr set. But what does it look like? So it's a bit hard to imagine. So what does it look like?

In the finite field setting, we know it's a subspace. But in the Z mod n setting, right now, it's just some subset of Z mod n. OK, so in the next step, we want to extract some geometric structure from this Bohr set. So we're going to show that this Bohr set will contain a large, generalized arithmetic progression.

So you asked something earlier about-- something seems a bit fishy about the general strategy. Seems like our goal for proving-- we want to prove Freiman's theorem, which says that the conclusion is that A is contained in some GAP, some fairly compact additive structure.

And we're already losing quite a bit. So we pass down to 1/8 of A. So it seems like even if you contain the rest, even if you can contain this fraction, this large fraction of A, what are you going to do about the rest of A? That's an unanswered question.

A second unanswered question-- so right now, what I've told you, the strategy is we're going to find a large GAP inside 2A minus 2A, which is not quite the thing that we want to do. We want to contain A in a small GAP. But at least it's some progress, right? It's some progress to find some structure.

I mean, the name of the game is to try to find additive structure. So in the theme of this whole semester course is trying to understand the dichotomy between structure and pseudorandomness. And when you have structure, let's use that structure. See if you can boost that structure.

So there will be an additional argument, which I will show you at the beginning of next lecture at the conclusion of the proof of Freiman's theorem, which will allow you to start with the structure on a small part of A, but not too small-- it's a constant fraction of A-- and pass it up to the whole of A. And we've actually already seen a tool that allows us to do that.

So I want to cover all of A. So last time, we did something called the covering lemma, Ruzsa covering lemma, that tells us that if you have some nice control on A and you can cover some

part of A very well, then I can cover the entirety of A very well. So those tools will come in hand.

I mean, so similar to actually how we proved Freiman's theorem in groups with bounded exponent. And so we're going to use the covering lemma to conclude the theorem. But now I want to get into the issue of the geometry of numbers.

OK. I want to tell you some necessary tools that we'll need to find a large GAP inside 2A minus 2A. Now, it will seem like a bit of a digression, but we'll come back into additive combinatorics in a bit. So the geometry of numbers concerns the study of lattices. So it concerns the study of lattices and convex bodies.

So this is a really important area of mathematics, especially about a century ago with mathematicians like Minkowski playing foundational roles in the subject. So number theorists were very interested in trying to understand how lattices behave. So I'll tell you some very classical results that we'll use for proving Freiman's theorem.

So first, what is a lattice? So let me give you the following definition of a lattice in R to the d. It's a structure on a group, if you will, as an integer span of d independent vectors.

So I start with v1 through vd vectors that are linearly independent. And I look at their integer span. I think this is best explained with a picture.

So if I have a bunch of-- so here, I'm drawing a picture in R2. And this picture extends in all directions. If I start with two vectors, v1 and v2, linearly independent, and look at their integer span, so that's a lattice. So that's what a lattice is. You can come up with all sorts of fancy definitions, like a discrete subgroup of R to the n. But this is what it is.

So just to emphasize this definition for a bit-- and also, one more definition that we'll need is the determinant of a lattice. So what's the determinant of a lattice? One way to define it is you look at these v's, and you construct a matrix with the v's as columns. And you evaluate the absolute value of this determinant.

More visually, the determinant of a lattice is also equal to the volume of its fundamental parallelepiped, which is a parallelepiped-- well, in the two-dimensional case, it's a parallelogram-- which is spanned by v1 and v2 or these v's, although you have more choices, right? So you could have chosen a different set of generating vectors. For example, you could have chosen these two vectors, and they also generate the same lattice.

And that's also a fundamental parallelepiped. And they will have the same volume. You can make some wrong choices, and then they will not have the right volume. So if you had chosen these two, so this is not a fundamental parallelepiped. Great.

So let me give you some examples. The simplest lattice is just the integer lattice, Zd, which has determinant 1. If I'm in the complex plane, which is viewed as two-dimensional real plane, then if I take, let's say, omega being the 3rd root of unity, I have a triangular lattice.

And the fundamental parallelepiped of this lattice, that's one example. And you can evaluate its determinant as the area of that parallelogram. If I take two nonlinearly independent vectors-- so for example, if I'm in one dimension and I look at the integer span of 1 and root 2, this is not a lattice.

Now, the next definition will initially be slightly confusing. But I will explain it through an example or at least try to help you visualize what's going on. So if I give you a centrally symmetric convex body-- "centrally symmetric" means that k equals to minus k. So centrally symmetric convex body, OK.

So here, centrally symmetric is x in k if and only if minus x is in k. And I'm in d dimensions. Let me define the i-th successive minimum to be lambda i. OK, so i-th successive minimum of k with respect to lambda to be the infimum of all non-negative lambda such that the dimension of the span of the intersection of lambda k and-- well, little o lambda k and the lattice-- has dimension at least i.

OK. So let me explain. I start with a lattice. So I start with some lattice. And I have some convex body. So this is 0, let's say.

So I have some convex body, a centrally symmetric convex body like that. I initially could be bigger, as well, but that's scale it so that it's quite small initially. And let's consider an animation where I look at lambda k where k lambda goes from 0 to infinity. This is k.

So initially, lambda k is very, very small. And I imagine it growing. It gets bigger and bigger and bigger. So it gets bigger and bigger.

And let's think about the first time that this growing body hits a lattice point, a non-zero lattice point. At that point, I freeze the animation. And I record this vector. I record this vector where I've hit a lattice point.

And now I continue the animation. It's going to keep on growing and growing and growing until when I hit a vector in a direction I haven't seen before. So it's going to keep growing. And then the next time I hit a vector in a new direction, I stop the animation. And I look at the other vector.

So I keep growing this ball until I hit new vectors, keep growing this convex body. So for example, if your initial convex body is very elongated, if that's your k-- so you keep growing, growing-- you might initially hit that vector. And then you keep on growing it.

And the next vector you hit might still be in the same direction. But I don't count it. I don't stop the animation here, because I didn't see a new direction yet. I only stop the animation when I see a new direction.

So I keep growing until I see a new direction. And I stop the animation there. So think about this growing body, and stop in every place when you see a new direction contained in your lambda k. And the places where you stop the animations, they're the successive minimum of k. Yeah?

AUDIENCE: Is this defined if i is greater than d?

YUFEI ZHAO: Is this defined when i is greater than d? No. So you only have exactly d successive minimum.

Now, sometimes you might see two new directions at the same time. That's OK. But once you exhaust all d directions, then there's no more new directions you can explore.

We also consider the vectors that you see. So let me also call these so that we can-- OK, so we can select these lattice vectors bi. I am going to use underscore to denote. So I'm going to use this underline to denote boldface.

So it's a vector bi, which is in, basically, this. You should think of bi as the new vector that you see. And it will have the property such that b1 through bd form a basis of Rd.

So I keep growing this convex body. When I see a vector in a new direction, I record lambda. And I record the vector bi. I keep on going, keep going, keep going until I exhaust all d directions. I call these b's the directional basis. OK. Any questions?

All right. So the result from the geometry of numbers that we're going to need is something called Minkowski's second theorem. So Minkowski's second theorem says that if you have

lambda, a lattice, in Rd and k, a centrally symmetric body, also in Rd, such that lambda 1 through lambda d are the successive minima of k with respect to lambda, then one has the inequality lambda 1, lambda 2. So the product of these successive minima times the volume of k is upper bounded by 2 to the d times the determinant of lambda.

For example, and here is a very easy case of this Minkowski's second theorem, if your k is an axis-aligned box-- namely, it is a box where the width in the i-th direction is 2 over lambda i-- so then you see that the successive minima of this box are exactly the lambda i's. And you can check that for-- this inequality is actually an equality.

OK. So actually, in this case, lambda, the lattice, is the integer lattice. Now, this is a pretty easy case of Minkowski's second theorem. But the general case, which we're not going to prove, is actually quite subtle.

I mean, the proof itself is not so long. It's worth looking up and trying to see what the proof is about. But it's actually rather counterintuitive to think about.

It's one of those theorems where you sit down for half an hour or an hour. You're trying to prove. You think you might have come up with a proof. And then on closer examination, it'll be very likely that you made some very subtle error.

So it's not so easy to get all the details right. And we're going to skip the proof. But any questions about the statement? OK.

We're going to use Minkowski's second theorem to show that a large Bohr set contains a large GAP. And specifically, we will prove that every Bohr set of dimension d and width epsilon-- epsilon is between 0 and 1-- in Z mod nZ contains a proper GAP with dimension at most d and size at least this quantity, which is epsilon divided by d raised to the power of d fraction of the cyclic group.

So just to step back a bit and see where we're going, from everything that we've done earlier, we conclude that 2A minus 2A contains a large Bohr set. Here, epsilon is 1/4. So epsilon is a constant. And R is also going to be a constant. It's depending on the doubling constant.

And this proposition will tell us that inside this 2A minus 2A, we will be able to find a very large, proper GAP. So "proper" means that in this generalized arithmetic progression, all the individual terms are distinct, or you don't have collisions. So you're going to find this proper

GAP that is constant dimension and at least a constant fraction of the size of the group, so pretty large GAP.

To find this GAP, we will set up a lattice and apply Minkowski's second theorem. Suppose the Bohr set is given by R where the individual elements, I'm going to denote by little r1 through little rd.

And let uppercase lambda be a lattice explicitly given as follows. It consists of all points in Rd that are congruent mod 1 to some integer multiple of the vector r1 over n, r2 over n, through rd over n, so congruent mod 1.

So for example, in two dimensions, which is all I can draw on the board, if r1 and r2 are 1 and 3 and n equals to 5, then basically, what we're going to have is a refinement of the integer lattice, where this box is going to be the integer lattice. And I'm going to tell you some additional lattice vectors. And here, it's going to repeat, or it's going to tile all over.

So I start with 1, 3. And I look at multiples of it. But I mod 1. So I would end up with these points and then repeat it.

And so you would have-- so that's the lattice. So you have this lattice, lambda. What is the volume? What is the determinant of this lattice? So the determinant of the lattice, remember, is the volume of its fundamental parallelepiped. So I claim that the determinant is exactly 1 over n.

There are a few ways to see this. So one is that, originally, I had the integer lattice as determinant 1. And now I put-- instead of one point, I have endpoints in each original parallelepiped. So the determinant has to go down by a factor of n.

Or you can construct an explicit fundamental parallelepiped like that. And then you use base times height. OK.

We're going to apply Minkowski's second theorem. And I will need to tell you-- I don't need the definition of Bohr set up there. So I want to tell you what to use as the convex body.

The convex body that we're going to use is k being this box of width 2 epsilon. So that's the lattice. That's the convex body. And we're going to apply Minkowski's second theorem. So let's let little lowercase lambda 1 through lambda d-- so n is d-- be the successive minima of k with respect to lambda and b be the directional vectors, the rational basis corresponding to those

successive minima.

I claim that the L-infinity norm of bj is at most lambda j epsilon for each j. And this is basically because of the definition. I mean, if you look at the definition of successive minima and directional basis, this is k. I grow k, grow it by a factor of lambda j. And that's the first point when I see b sub j. So every coordinate of b sub j has to be at most this quantity in absolute value.

So now let me denote uppercase L sub j to be 1 over lambda jd rounded up. And I claim that if little l is less than big L-- lj, Lj-- then little lj-- OK, so if I dilate the bj vector by factor little l, so if I plug it in and just look at these two inequalities, I obtain an upper bound of epsilon over d on the L-infinity norm of lj bj, so just looking at this bound here and the size of lj. And if this holds for all j, then summing up all of these individual inequalities, we find that the sum of these lj bj's is at most epsilon in L-infinity norm.

So the point here is that we want to find the GAP in this Bohr set. And how does one think of a Bohr set? So it's kind of hard to imagine, because the Bohr set is a subset of Z mod n. But the right way to think about a Bohr set is in a higher dimensional lift, because a Bohr set is defined by looking at these numbers for R different values, R different coordinates.

So we think of each r as its own coordinate. So we think of there being capital uppercase R many coordinates. And we want to consider the set of x's so that all the coordinate values are small. So instead of considering a one-dimensional picture, as we do in the Bohr sets, we're considering a higher dimensional or d-dimensional picture and then eventually projecting what happens up there down to this Bohr set.

So what does Minkowski's second theorem have to do with anything? Well, once you have this higher dimensional lattice, what we're going to do is find a large lattice parallelepiped, so a large structure inside this higher dimensional lattice, and then project it down onto one-dimensional Z mod n.

So this is the process of-- so you already see some aspects of a GAP in here. So these guys, they're essentially the GAP that we're going to eventually wish to find. And right now, they live in this higher dimensional lattice. But we're going to pull them down to Z mod n.

All right. Now, where do these b's come from? So each b sub j is congruent to some x sub j times this vector mod 1 where x sub j is an integer between 0 and uppercase N.

So this inequality up here, star. So the i-th coordinate for star-- "coordinate" meaning this is an L-infinity bound, so the i-th coordinate is upper bounded by epsilon. But the i-th coordinate bound implies that if you look at this sum over here times Ri divided by N, this quantity, whatever it is, is very close to an integer for each i.

So the i-th coordinates implies this inequality, and it's true for every i. Thus what we find is that the GAP, which you already see in this formula over here-- so the GAP is given like that. So this GAP is contained in the Bohr set.

So we found a large structure in the lattice. But the lattice came from this construction, which was directly motivated by the Bohr set. So we find a large GAP in the Bohr set. Well, we haven't shown yet it is large or that it is proper. So we need to check those two things.

To check that this GAP that we found is large, we're going to apply Minkowski's second theorem. Let's check GAP is large. So by Minkowski's second theorem, we find that the size of the GAP, which is, by definition, the product of these upper case L's-- so if you look at how the uppercase L's is defined, you see that this quantity is at least 1 over the product of the successive minima times denominator d to the d.

And now we apply Minkowski's second. And we find that this quantity is at least the volume of k divided by 2 to the d times the determinant of the lattice times d to the d. But we saw what is the determinant of the lattice. It is 1 over N. You have d to the d, 2 to the d. And the volume of k, well, k is just that box. So the volume of k is 2 epsilon raised to d.

So putting everything together, we find that the size of this GAP is the claimed quantity. It's a constant fraction of the entire group. The second thing that we need to check is properness.

So what does it mean to be proper? So we just want to know that you don't have two different ways of representing the same term in the GAP. So if I have the following congruence, so if this combination of the x's is congruent to a different combination of the x's where these little l's are between 1 and-- OK, so I want to show-- so to check that it's proper-- so we're in Z mod n-- we just need to check that if this holds, then all the corresponding little l's must be the same as their primes.

Well, if it is true, then setting-- let's go back to the lattice-- setting the vector b to be a vector originally that corresponds to the difference of these two numbers-- so if we set b to be the difference of these two numbers, we find that, first of all, it lies in Z to the d, because these two

numbers are congruent to each other mod n.

And furthermore, the L-infinity norm of b is upper bounded by-- I mean, each one of them has small l-infinity norm. And this is some number that is bounded. It's less than uppercase L.

So the whole thing, this whole sum, the L-infinity norm, cannot be larger than this quantity over here, where I essentially use the triangle inequality to analyze this b term by term. All of these numbers are very small, because if you look at what we saw up there, so the size of b, we see that this whole thing is at most epsilon. And epsilon is strictly less than 1.

So you have some vector b, which is an integer vector, such that all of its coordinates have L-infinity norms strictly less than 1. So that means that b is equal to 0. So b is the 0 vector. So b is a zero vector.

Thus this thing here equals to 0. So this sum here equals to 0. And since the bi's form a basis, we find that the li's and l prime i's are equal to each other for all i. And this checks the properness of this GAP.

Yeah. So this argument, it's not hard. But you need to check the details. So you need to wrap your mind around changing from working in a higher dimensional lattice setting to going back down to Z mod n. And the main takeaway here is that the right way to think about a Bohr set is to not stay in Z mod n but to think about what happens in d-dimensional space where d is the dimension of the Bohr set.

OK. So now we have pretty much all the ingredients that we need to prove Freiman's theorem. And that's what we'll do at the beginning of next lecture. We'll conclude the proof of Freiman's theorem.

And then I'll tell you also about an important conjecture in additive combinatorics called a polynomial Freiman-Ruzsa conjecture, which many people think is the most important open conjecture in additive combinatorics.