[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: We have seen, I would say, three sets of tools applied to the problem of projection theory. So at the beginning of the class, first, we talked about double counting. And then we talked about methods from Fourier analysis. And then in the last week, we started our third set of tools which are tools from combinatorial number theory.

So this week we're going to finish our unit on tools from combinatorial number theory in projection theory. So the goal for these two weeks is this projection theorem of Bourgain, Katz, and Tao. It takes place over finite fields, and the cool thing about it is that it distinguishes what's true over prime fields from what's true over other fields, which none of our previous methods could distinguish.

This is our goal. It says that if x is a subset of Fp squared, and it has size p to the sx-- you can think of sx like a dimension-- and sx is in between 0 and 2. So it's not too tiny, and it's not everything. And we have a set of directions in Fp, a set of directions as size p to the sd, and it's not too tiny.

Then, if I take the maximum over all of the directions of the size of the projection in that direction, it's at least-- so the square root of x. So far, that's not difficult to do. But then plus-- times a little bit more, times p to the epsilon. I'll call it epsilon-- p to the epsilon.

So this epsilon depends on these dimensions, but it's positive. So that's our goal. And it's significant that this would not be true in other fields. So this is false in Fq squared if q equals p squared.

And the example is that x is Fp squared, and D is Fp. And so then this thing is always Fp. Cool. So even though this epsilon is a really tiny number, which we could make explicit if we worked hard, but we would then be disappointed. Even though it's really tiny, this theorem is qualitatively an important improvement, and it has lots of applications.

This and the version over r have lots of applications, and almost everything in the rest of the course will be based on that, kind of whatever most of the directions we might go in. All right, so last week, we introduced combinatorial number theory and some product theory. And we proved a theorem, which is in this direction.

It's kind of the special case when x happens to be a product set-- A cross A. So our Theorem 1 was last time, if A is in Fp, and the size of A is p to the sa, and D is the set of directions in Fp, and size of D is p to the sd. Then, if you take the maximum t and D, of A plus t times A, which would be like pi sub t of A cross A, then that's at least the size of A times p to the epsilon 1.

That's a little bit bigger. And this, again, would be false in other fields. So if our field had a subfield, you could choose these guys to both be the subfield. And this would stay in the subfield. Cool. All right, so our goal today is to get from here to there.

And we were talking about it last time, and I wanted to-- so at first this looks like a very special case. I'll try to argue that it's only a little bit special case, and we'll see what the issue is. We call this an attempt at BKT using Theorem 1. So we're going to do a proof by contradiction.

We're going to say epsilon is very small, and it's to be decided. And so then we'll assume that pi t of x is less than x to the half times p to the epsilon for all the ts in D. So proof by contradiction, so we can assume that. And also, without loss of generality, we can suppose that 0 and infinity are in our set of directions, D. So in other words, we include the horizontal projection and the vertical projection.

So then we can make a set, A, so that A times A contains x And x is most of A times A. So the set A is pi 0 of x, union pi infinity of x, or the size of A bounded by size of x to the half times p to the epsilon. And x is contained in A times A.

So x is a big subset of a product, and we can apply Theorem 1 to our product. So theorem one tells us that the maximum t and D of pi t of A times A is big. And even bigger than that, so this could be epsilon 1. And epsilon 1 is much bigger than epsilon.

So epsilon 1 came from a theorem that's already proven. It's some number. And we get to choose epsilon as small as we like, so we choose it much smaller than epsilon. But this is not yet the projection theorem that we're aiming for. And the issue is we would like pi t of x to be big.

So the issue is, maybe pi t of x is much smaller than pi t of A cross A. Even though x is a big chunk of A cross A. Let me make a picture of this enemy scenario. So here's the enemy scenario.

All right, so there's a set x which has a small projection. So we're going to be doing a projection, say, in this direction. Projection of x is small. And x is a big portion of this larger set, A cross A.

So in blue, we have A cross A without x. And when we project that down, then we see a lot of stuff. Pi of t, A cross A without x. So we're worried that we may be so unlucky that even though pi t of A cross A is quite large, and even though x is a big chunk of A cross A, nevertheless, pi t of x is small. We're worried about that.

So that's where we left off before vacation. And now what I'd like everyone to do is to try to think of an example where this happens. Does this ever actually happen that inside of A cross A, you have a large subset x, and the projection of x is small, even though the projection of A cross A is big? Can you think of an example where that actually happens? Yeah.

AUDIENCE: You could choose x to be like the diagonal-- like all the values like, I guess, pi I for I is in Fp.

PROFESSOR: Right. So here's an example. So x is going to be contained in A cross A. X is the set of all A comma A in A cross A. And so then if I take the projection, I guess, in the direction minus 1 of x, I just get 0. A minus A is always 0.

So this is extremely small But in this case, this x is quite a bit smaller than the whole A cross A. So let me make a more precise question. So let's say A has size N. And x is a subset of A cross A, and it's really quite big.

So I'm going to say x is bigger than-- well, we're going to call it for now K inverse times N squared. But then maybe K is order of N to the 1 over 1,000. So it's really a big portion of A cross A.

And pi t of x is like way smaller than pi t of A cross A. We could make that more precise too, if we wanted to. And that's the situation we have in our attempt.

A is only a bit bigger than x. And so let's say the size of A is N. It's just the definition of N. And then x would be bigger than like p to the minus a few epsilon size of A cross A. So it's really a good portion of that.

So could it still happen that the size of the projection of x is way smaller than the size of the projection of the whole thing?

AUDIENCE: This is like-- can you pigeonhole principle somehow, like the fibers of the projection of A cross A?

PROFESSOR: Yeah.

AUDIENCE: And then if x is big enough, you know it has to be at least some nice number. I don't know.

PROFESSOR: So the suggestion is we can look at all of the fibers of the projection of A cross A, and then we should pick x to be the union of the big fibers. Yeah. That's right. That's right.

So in this enemy scenario, there have to be some very big fibers and some much smaller fibers that have to be kind different from each other. Yeah. But let's see if we can think of an example where this happens. What I should erase-- probably still have some space.

So I'm just going to continue below the line here. So first of all, can we even think of an example where the projection of x or of something would be small at all? So an example, maybe x could be the numbers from 1 up to N cross the numbers from 1 up to N. And then pi 1 of x is pretty small.

It would be basically integers from 1 up to 2N, so this would be quite small. So you could make this A. So x would be all of A cross A. Then the projection of A cross A would be small.

But you could also inflate A to make it bigger than that. So we could take A to be the integers from 1 up to N union, A tilde. Maybe this size of A tilde is also N, but A tilde is unstructured or random.

So now if I take pi 1 of A cross A, it's at least as big as pi 1 of A tilde cross A tilde. And A tilde has nothing to do with the numbers from 1 up to N. I could build some set where that's big.

But pi 1 of x is only around N. So there could be a big gap between these guys. But this example is not so-- not such a problem for our strategy. Because in this example, we just should have chosen the set A smaller like A to B.

We could choose as we like. We could choose a smaller set A so that A actually with the smaller choice of A, just these guys, x would be A cross A. And then we would be happy.

So the key to the proof of Bougain, Katz, Tao is a theorem from combinatorics called the Balog-Szemeredi-Gowers theorem. And it says that this example is pretty much the only way that pi 1 of x could be much smaller than pi1 of A cross A.

So, this is called the Balog-Szemeredi-Gowers theorem. And it's another of the central tools in the toolbox of additive combinatorics, along with the Plünnecke-Ruzsa inequalities that we talked about last time. All right. So it says, if A and B are contained in an abelian group and they have size at most, N. And I have a subset of A cross B, which is significant-- a significant fraction.

And the projection of this subset in some direction is small. All right. So, here, I want you to imagine that K is a lot smaller than N. So, this is a big chunk of all the pairs, A and B. And the smallest this could possibly be is around N, because if you just fix one element of B, you'd have around N choices for this element of A. So that would already give you around N here.

So, this is almost as small as possible. And then, the conclusion is, it's not necessarily the case that pi T of all of A cross B is small. We just saw a counterexample to that. But we could find big pieces of A and B so that the projection of the big pieces is small. So, then, there exists A prime and B prime, which are subsets of A and B. A prime and A, B prime and B, so that, I'll say, x prime is defined to be A prime cross B prime, intersect, x.

And x prime is still a big fraction of N squared. So, K to the minus some constant power, times N squared. And if I take pi T of A prime cross B prime, that is less than K to some power times N. So, this is small.

So in this example here, where A was like this, this part here would be the A prime. And it would be symmetric. B prime and A prime would both be this. OK, take a moment to digest that statement and see if you have questions or comments.

OK. So, we will prove this probably starting later today, but most of it next class. But for the rest of this class, we'll explore why this is useful. And in particular, we'll use it to prove the projection theorem, OK? Cool.

All right. So, let's do the proof of the Bourgain-Katz-Tao projection theorem. So, let's say epsilon is small, to be decided. So we're going to prove, there's some small epsilon where this is true. And I'm going to write C less than tilde, tilde D. It means that C is less than p to the big O of epsilon, times D.

OK. And, now, we're going to do a proof by contradiction. So we can suppose pi T of x is less than or equal to p to the epsilon, times x to the half, for all the T and D. OK. Without loss of generality, I can suppose that 0 or infinity are in my directions. And I'll take A to be pi 0 of x union, pi infinity of x. And so, now, I know that x is a subset of A cross A.

I'll say, A, the size of A could be N. And that is tilde, tilde, the square root of x, up to some powers of p to the epsilon. OK. And, therefore, x is greater tilde, tilde, N squared. So x is a big fraction of all of A cross N. Cool. Next, I choose some other T1 in D.

And I know that pi T1 of x is less than tilde, tilde, N. And so, once I know this, the Balog-Szemeredi-Gowers theorem tells me that this comes from a product, A prime, A prime cross B prime. So, Balog-Szemeredi-Gowers says, this A prime in A, B prime in B.

And x prime is the part of x and its product, A prime cross B prime. So, first of all, x prime is still quite substantial. And pi T1 of A prime cross B prime is small. And this means, in other words, that A prime plus T1 B prime is small.

And this here is-- this information here is strong in a helpful way. It's stronger than this, because, here, we can use the Plunnecke-Ruzsa inequalities. So, once we know that A prime plus T1 B prime is small, then Plünnecke-Ruzsa inequalities tell us that lots of other things are small. So, for example, A prime, plus or minus A prime, is also small.

If I wanted to, I could do a triple sum, A prime, plus A prime, plus A prime, is also small, dot, dot, dot. You could do the same with B prime. OK. So, now, we see that A prime and B prime have a lot of structure, have a lot of structure. And, actually, I'll use the word, symmetry. All right. OK. Actually, I'll draw you a picture in a minute of how I visualize this, and why I use the word symmetry.

OK. So, next. Sorry, I have to make space. I'm going to erase the statement of the theorem. We're going to try to prove that one of the projections of x is big. All right. All right. All right. So, now, we have a lemma that says, if I have a subset y in A prime cross B prime-- the subset we're going to care about is just x prime. But the lemma is more general.

And A prime and B prime have these nice properties. So, I'll just copy these two. A prime plus T1 B prime is small. A prime minus A prime is small. And y itself is fairly large, although I don't need to write it in the statement. Then. pi T of y-- so this is for any T. So, T is just anybody in Fp, different from T1.

This guy is bigger than, tilde, tilde, y over N squared, times pi in a different direction, T over T1, of A prime cross A prime. So, we have-- the biggest y it could be is, it could have size around N squared, in which case this would disappear.

And so, I said, if we have a large subset of A prime cross B prime, then the projection of that large subset y is bigger than the projection and related angle of a whole product, A prime cross A prime.

This is good for us because we can apply theorem 1 to this. And, this inequality, you're kind of ruling out the enemy scenario. We're ruling out the possibility that there's a big chunk of A prime cross B prime that got compressed, even though the whole thing didn't. OK. Let me put over here, some intuition.

So, my intuition is that A prime cross B prime is very symmetric. And so, let me draw an exaggerated version. Suppose that A prime cross B prime does look like a lattice. And inside of there, we have some substantial but unstructured set, y.

So, if there were some angle where the projection of y got compressed a lot, I claim that compression of the whole lattice should get compressed a lot. Why? Because the lattice is symmetric. So I can cover the lattice by some translates of y. And if each of those gets compressed a lot, then the whole lattice should be compressed a lot.

So, we cover A prime cross B prime by a few translates of y. And then we would get that the projection of the whole thing should be bounded by a few times, the projection of y. So, now, let's carry out this intuition. We are going to cover y, cover A prime, cover something, by translates of y.

So, let's consider y plus T1, T2. These are some translates of y, a set of these guys, where T1 is an A prime cross A prime, minus A prime. And T2 is in this slightly weird-looking thing there. And so, these are a bunch of translates of y. And these translates of y cover somebody a lot of times. So, for every y1, y2-- let me do it this way.

For every A1, A2, and A prime cross A prime, and for every y1, y2, and y, there is a unique T1, T2, which is in this list, so that y1, y2, plus T1, T2, is a1, a2. So this is easier to prove than it is to write down. Just solve for T1. T1 must be a1 minus Y1. And these guys both live in A prime. So, T1 is in there.

And, ah, wait. Yeah, sorry. So this is going to have to be negative 1 over T1a prime. OK, now, how does y2 work? So, T2 is going to be a2 minus y2. And the a2 lives in B prime.

Yeah. So, the a2 lives here. And the y2 lives here. So the difference lives here, where it's supposed to. So, these translates of y cover A prime cross negative 1 over T1a prime. They cover it y times. They cover each point in here y times because, for each point in here, I have norm of y, cardinality of y choices for this guy. And then there's one in translation, a set of translations that lands here.

Does this sentence make sense? So we have rigged it, that we've taken a bunch of translations of this set, y. And we have covered a nice product many times. And the product, we actually could have arranged the product to be A prime cross B prime. That would have been a little easier than what we did. But we have arranged that the product that we've covered is A prime cross a dilation of A prime.

And that's a little bit better, technically, because it's going to be closer to theorem 1. Now, notice that this guy is small by hypothesis. And this set here is also small. This is just a dilate of this. So these sets are small. So we took a few translations of y. And we covered this product many times. And that means that, at least, typically, each of these translations was covering a lot of the product.

So that matches this intuition. So, then, we can say that pi T of A prime cross minus 1 over T, T1a prime, that guy is bounded by-- OK, so we're going to take A prime minus A prime, times A prime plus T1 B prime. This is how many copies of y we took. For each one of them, we have a projection. So that has the size, pi T of y. And after doing that, we've covered every point here y times. So we have this guy.

These are hypothesis bounded by N. So this is N squared over y, pi T of y. And that's basically our final inequality. Put the pi T of y by itself. Pi T of y is greater, tilde, tilde, y over N squared, times the size of this thingy, which is equivalent to pi negative T over T1, A prime cross A prime.

OK, take a moment with that. OK. So, now, let's finish the proof of the BKT theorem. All right.

So, we would like to estimate the maximum T in our direction set of pi T of x. That's at least the maximum T in our direction set of pi T of x prime. We're following up from over here. And here's x prime, which is a big chunk of A prime cross B prime.

And now, we can apply our lemma. The lemma says, that is at least as big as the maximum T in our direction set, x prime over N squared. And then we have pi negative T over T1, A prime cross A prime. Now, x prime, we already determined, has a size of about N squared. So this factor disappears.

And now, we are looking at projections in many different directions of A prime cross A prime. That's exactly the thing that theorem 1 told us about. So, by theorem 1, this is now greater [INAUDIBLE] N times p to the epsilon 1. And that's the proof. So, one of these is significantly bigger than the square root of x times p to the epsilon 1.

OK. Let me put up on here a little summary of what we did. And then I'll let you look down from above at the whole thing. So, we're trying to get from theorem 1 to theorem BKT. And the enemy was that we have a set x in A cross A, where x is around the same size as A cross A. But the size of the projection of x is much smaller than the size of the projection of A cross. A. That's our enemy.

And we learned two pieces of wisdom about this enemy, one in the BSG theorem, and one in the lemma. So, wisdom one, in the enemy situation, x lives in a smaller product. And the projection of the smaller product-- the projection of the smaller product, well, it doesn't live in that. But it has a-- so the enemy situation implies that x, intersected with a smaller product, is about the same as x.

So, a good chunk of x lives in a smaller product. And the projection of the smaller product is about the same as the projection of x. Then we look at that situation. And the second piece of wisdom was that, if the projection of a smaller product, of a product-- if this is small, so if this is around the smallest it could be, it would be one of the factors for one T1.

Then this product becomes very symmetrical. And then the conclusion is, for any subset, x prime, in this thing, which is substantial, then the projection of the subset actually is around the same as the projection of the whole product. So the enemy situation can only really happen because x was mostly in a smaller product that had a small projection.

And this situation is very symmetrical. So if a smaller product has one small projection, then when you project it in other angles, then the projection of a big piece of it can't be way smaller than the projection of the whole thing. All right. So, the whole story is on these boards. Take a couple of moments to look it over. Yeah.

AUDIENCE: Professor? So, should it be the case that the A prime there depends on T?

PROFESSOR: OK, so the question is whether A prime depends on T. So, let's see where A prime appeared. So we have our set with many small projections. And, OK. And then, in particular, the T1 projection is small. And so, the Balog-Szemeredi-Gowers theorem says there's A prime in A and B prime in B, so that the T1 projection of A prime cross B prime is small.

And this A prime depends on T1. Now, later, we're going to look at pi T of x for other T's. One could, I guess, apply Balog-Szemeredi-Gowers again to those other T's. But that's not what we're going to do. We use T1 to find A prime and B prime. And then we're going to run with these same A prime and B prime, even when we look at other T's. What did we learn?

Once we have this one small projection, we learned-- Balog-Szemeredi-Gowers told us this. And then we learned from it some other stuff. So, what we actually used is this. So, then, we used this stack. So this is telling us that A prime cross B prime is quite symmetrical. And that's what's helpful in the second part.

AUDIENCE: OK, cool. Thanks.

PROFESSOR: Good. Other questions or comments?

AUDIENCE: Why don't we know that A itself is also kind of symmetrical?

PROFESSOR: Yeah, the question is, why don't we know that A itself is also kind of symmetrical? That might depend on exactly what we mean by symmetrical. But an important thing, an important way that A prime is better than A is that, at this moment, A prime minus A prime is small. And we didn't know that about A. You can think of A prime as being A prime plus some garbage. And the garbage minus garbage is not small.

AUDIENCE: So, [INAUDIBLE] epsilon is different between the two because of the A prime plus minus A prime has-- there's a hidden factor there, right?

PROFESSOR: Yes, A prime plus, minus, A prime. Here, there's a hidden factor, which is p to the order of epsilon. And you're right. It's very important that the epsilon is different between the two theorems. So, epsilon 1 is the epsilon in theorem 1, which was already kind of small, if we had unwound the proof of theorem 1. And then, in theorem BKT, there's an epsilon. And epsilon in BKT is way smaller than the epsilon 1.

And the reason we need that is, OK, we started with this, where we have a factor of p to the epsilon. Then we did some stuff. Most important stuff is BSG. So, BSG, we'll lose-- there'll be a factor of 10 epsilon here that comes out of the proof of Balog-Szemeredi-Gowers. So, here, these things will be something like p to the 10 epsilon, if you did them carefully.

And then you go along. And we're following our proof. So, here, we did a little bit more work in the proof of the lemma. So, this is hiding a p of 10 or 20 epsilon. But, here, we have a gain of p to the epsilon 1. So, the p to the epsilon 1 needs to dominate the p to the 20 epsilon that we lost at various steps in this process.

So, epsilon should be 1/20 of the size of epsilon 1. Yeah, so we have to do a whole sequence of clever things to prove Bourgain-Katz-Tao theorem. And each time we do something clever, epsilon gets smaller by a factor of 10, until we're so clever that we really can't see the epsilon. Yeah.

AUDIENCE: Is there a conjectured boundary for what the largest possible epsilon could be?

PROFESSOR: Ah, that's a great question. Yeah, so, the question is, is there a conjectured bound about the right epsilon? Yes. So, yes. So, the conjecture, so we're working right now over finite fields. And the conjecture is that the worst case is a lattice. And we wrote it down. Let's write it down again. But we wrote it down at the beginning of the course.

So I think we're moving on. I'll put it here. So, suppose we have x, which is a subset of Fp squared. And then we're going to take the maximum t and D of pi t of x. And this is bigger than-- I think it was D to the 1/2, x to the 1/2. Maybe you should check it I don't remember these. I don't quite remember this off the top of my head. So, I think.

And the way to remember it is that the example is that x is a lattice. And then D is the set of rational numbers of small slope. That was the case with-- the case that we've seen so far, where the projections are small. So if you compute with this, you can check whether this is true or whether it needs to be modified a little bit. And then, if you rewrite it in the language of BKT, it gives you an exact formula for epsilon.

And that is a wide-open problem over FP. But the version over r has been proven. And I'll try, at least, to say at least some of the main ideas of the proof in the class. Yeah, OK. So, the ideas that we've just been talking about, going back and forth between whether a projection of the whole product is small, whether a projection of a big chunk of the product is small, that is a theme in additive combinatorics.

And I want to describe to you some of the other key words and ideas behind that theme to flesh out the lecture and see how it fits in, in general. OK. All right. So, this section is called additive energy versus the size of A plus B. So, these are two points of view about measuring the additive structure involved when you add numbers in A to numbers in B.

This one, we've already met. Let me tell you what additive energy is. So, the definition, if A and B are contained in an abelian group, the energy of AB is the number of quadruples. A1 and A2 are in A. B1 and B2 are in B, so that A1 plus B1 is equal to A2 plus B2. OK.

So, if you add numbers from A and B, how many coincidences are there? Does it frequently happen that you get the same number? So, another way of thinking about it is we could say, r, AB of z, so the number of representations of z as a sum of somebody in A and somebody in B. So this is the number of AB, so that A plus B is-- and the additive energy, AB, is also equal to the sum over z, rAB of z, squared.

So if this is large, it means there are a lot of numbers that can be written as A plus B in many different ways. There are many additive coincidences. So, for reference, if A and B both have size, N, then let's think about how big or how small the additive energy could be. It's at least N squared. That happens if-- so, there are N squared obvious solutions to this, which are, A1 equals A2. And B1 equals B2.

And if you have a random set of integers or a random subset of a big group, those might be the only solutions. So, here, we have only trivial quadruples, OK? And energy is at most N cubed, because once you've picked three of these, the fourth one is determined by this equation. So there's, at most, N cubed of these additive quadruples.

And this happens if A equals B, equals A subgroup of z. A and B are subgroup of z. And there's always a choice of B2 you can pick to solve this equation. OK, cool. So, having large energy is having a lot of additive structure. And, before, we were thinking that, if A plus B is small, then that would mean there's a lot of additive structure. How are those related to each other?

OK. So, there's a lemma that says, A times B, A squared times B squared, is, at most, the size of A plus B, times the energy between A and B. And the proof of this is just Cauchy-Schwarz. So, if I add up on z, rAB of z, I get the size of A times the size of B. This is by double counting. So, every time I add somebody in A and somebody in B, I get some z. I add this up and get all the pairs.

So, now, A times B is the sum, z, of rAB of z. And in this sum, I only need to count the guys in A plus B, because z cannot be written as A plus B at all. It means that this was 0. Now, I Cauchy-Schwarz. So I'm going to write it as the sum for A plus B of 1 to the 1/2 half, and the sum, rAB of z squared to the 1/2.

So that's the size of A plus B to the 1/2, times the energy to the 1/2. And that's the proof. All right. So we see from this that, if A plus B is small, then the energy must be large. We could bring the A plus B over to the other side. So, people usually write this. The energy is at least A squared times B squared, over A plus B.

So, for instance, if A and B have size, N, and A plus B is almost N, then the energy would be at least around N cubed, which is as big as it could be. All right. So, A plus B small implies that the energy is large. Now, do people think you can go the other way? If the energy is large, does that imply that A plus B is small?

AUDIENCE: No, because of course, you can have some garbage.

PROFESSOR: Good, no because you can have some garbage. Good. So, energy of AB large does not imply that A plus B is small. The example is that A equals B, equals numbers up to N, somebody with a lot of additive structure, union, some garbage. Good. Now, if you look at the definition of the energy, if I add some extra elements to A or to B, it can only increase the energy, perhaps not very much.

So the energy of this example is at least the energy of this example, which is large. On the other hand, the garbage can make this a lot bigger. And this is basically the same example that we've been grappling with all class. OK, cool. So, now, I'll leave that up. All right.

OK. Now, if the energy of A and B is small, that implies that the size of A plus B is large. That's just the contrapositive of what we wrote before, or you can see it from this formula. This is small. Then this one must be large. Right, OK. Now, all right. Now, suppose I have a set, x, in A cross B. And I look at pi1 of x.

So remember that pi1 of A cross B is just A plus B. So when I take pi1 of x, it's like A plus B. But I don't take all the pairs. I just take the pairs in x. All right. So, even if pi1 of x were small, if x was a large set and pi1 of x were small, that would still produce a lot of energy because that would still be a lot of numbers, a bunch of numbers, z, that have many representations of A plus B.

Let me just write it down OK, so, yeah. So, actually, so let me recall that A plus B large, and x being a big fraction of A times B, that does not imply that pi1 of x is large. And this is, again, the main thing that we were struggling with. But with energy, this is true. So, lemma, let me check exactly what it says.

All right. So, lemma, x squared is less than or equal to pi 1 of x, times the energy of AB. So if the energy is small, then pi 1 of x is pretty big for every subset, x, for every large subset, x. The proof is almost the same as the previous proof. So, proof, rx of z is the number of ways of writing A plus B equal to z, where the pair, AB, is in x.

And that's smaller than representations of z as a general sum of sum between A and B, because we're only taking a portion of the sums. OK. So, now, x is the sum on z of rx of z. And, here, in this sum, we only have to take z's in pi 1 of x. We only have to consider those z's that can be written as the form, A plus B, where A and B are in x.

So, now, we can Cauchy-Schwarz this. And we get the size, pi1 of x to the 1/2, times an energy thing, which is smaller than the standard energy. So we have the sum, rx of z squared to the 1/2. That's smaller than the sum rAB of z squared to the 1/2, which gives us the energy. OK, cool.

So to summarize, having small energy implies that the sum set is large. But it's stronger than that. And it also implies that pi1 of x is large for all substantial subsets of A cross B. So it's stronger, in a way, which is exactly what we needed in our discussion. So that suggests a better version of theorem 1 that we could have tried to prove.

Let's erase this lemma. So, remember, theorem 1 said that, if we take the projection of A cross A in many directions, one of them has a large cardinality. But we could have proved, we could have tried to prove something stronger, that one of them has a small energy. Let me write it down. It's true. And it's a theorem. Theorem 2 says-- all right.

So if A is a subset of Fp, size of A is p to the SA, 0 less than SA, less than 1 D is in Fp. Size of D is p to the SD. 0 is less than SD. Then I'm going to consider all the directions and pick the best one. And instead of looking at the size of A plus TA, I'm going to look at the energy of A, comma TA. And I'll take the minimum over all the T.

OK. So, the biggest this energy could be is p to the 3SA. It's the maximum possible energy. And we're going to say, it's a little bit better than this. Now, because of this lemma, we automatically get not only a lower bound on pi T of A cross A. But we get a lower bound on pi T of any big subset of A cross A.

So, this automatically gives the BKT theorem. So, in fact, this easily implies a strong version of the BKT theorem, which I think was also proven by BKT. And it says that-- so, sorry. x is p to the sx. And D is as above. So, not only you can say, the maximum T and D, pi T of x is large.

So, that's the regular theorem. And this follows from theorem 2, and using our lemma. But we could say that-- so, this is true. But, even better, there is some T and D so that, if y-- so if y is a subset of x and it's a substantial subset of x, then pi T of y is big.

OK. So, the proof from here to here, we observed earlier that we can reduce to the case. Well, OK. Yeah, I guess I should say, using Balog-Szemeredi-Gowers, we could reduce to the case where x is a big subset of A cross A. And then y is a big subset of x. So y is a big subset of A cross A. And then you use the energy bound to get this.

OK. So, then, we say, proof of theorem 2 plus BKT2 are similar to what we did in class. So, they could be good optional exercises, or some pieces of it might be required exercises.

And I thought I should mention this. One of the things that Pablo impressed on me when we were talking over the vacation week was that this slightly stronger version is important in applications. Most of the applications actually depend on being able to say the slightly stronger thing.

OK OK. So, now, you might think, you might ask yourself, maybe energy is actually a better thing to be talking about because a bound for the energy is kind of stable, undertaking large subsets.

And so, maybe, we should have been talking about energy all along. And it's reasonable to ask, if you take our proof of theorem 1 and you just try to change all the cardinalities to energies, could you adapt it and prove theorem 2, instead?

And you can't do that. And it's a good moment to recap the main idea in theorem 1, and also a way of reflecting on the different big tools and how they're related to each other. All right. Let's look at this. Maybe, now, I can erase the BSG theorem. OK. Actually, before I erase the BSG theorem, let me make a little comment about it.

So, there are a couple of closely-related forms of the theorem. And another way I often say it is that, if the energy of AB is large, then there must be A prime and B prime, so that the cardinality of A prime plus B prime is small. Maybe I should write it, yeah. Let me put it here. So, theorem BSG variant. This is the way it's most commonly said. So I would be a bit remiss not to show it to you.

So if A and B are in an abelian group, and let's say A and B have size, at most, N. And the energy of AB is large. Energy of AB is greater than A inverse N cubed. Then it may not be the case that A plus B is small because of the example with garbage. But we can get rid of the garbage. So there exists A prime in A and B prime in B, so that A prime-- yes, so that they are substantial.

So, A prime and B prime are greater than or equal to K to the minus order of 1n. And A prime plus B prime is less than or equal to K to the order of 1N. So if we can locate the correct structured parts, then their subset is small, after we remove some garbage. Yeah, OK.

All right. So, now, let me remind you of one of the key ideas in the proof of theorem 1. Key idea in the proof of theorem 1. So, the key idea was based on the Plünnecke-Ruzsa inequality. And it was an idea that I called, contagious structure. So, it's said, if A plus T1a is small and A plus T2a is small, then A plus T1 plus T2a is also small-- not quite as small, but still pretty small.

All right. So, this was the key thing. If A plus TA is small for a bunch of T and D, then it would have to be small for an even bigger set of D by adding D, or by multiplying D, or by taking negatives. And then there would be a huge set of T's, where A plus TA was small. And then that contradicts basic double counting. That was the proof.

OK. So, what if we tried to work with energy throughout? So, instead of trying to prove that A plus TA is big, we're going to try to prove that the energy is small. So, we're doing a proof by contradiction. So the analogous question would be, if the energy between A and T1 of A is large, and the energy between A and T2 of A is large, then does that imply that the energy between A and T1 plus T2A is large?

OK. The answer is no. And it is the same issue that we saw before, that when we say, the energy is large, it means there's at least a piece of A and a piece of T1A that are very friendly with each other. But it doesn't have to be all of them. So, one piece of A must be very friendly with T1A. And one piece of A must be very friendly with T2A.

But those could be two different pieces of A. And in that case, we can't really learn anything more. It doesn't mean there's any piece of A that's friendly with both of them. So, here's the example. The answer is no. Here's the example. So, A might be the numbers up to N, union T1 times the numbers up to N, union T2 times the numbers up to N.

OK. So, then, T1A would be T1 numbers up to N union, dot, dot, dot. T2A is T2 numbers up to N, union, dot, dot, dot. So, these two guys are very friendly, not just because they're the same as each other. Not everybody who's the same as each other is friends, but because they have a lot of additive structures. And, therefore, this guy is large.

And these two guys are very friendly. And therefore, this guy is large. But there's nothing going on with T1. Plus, T1 and T2 are totally arbitrary. T1 plus T2, we would not typically have this large. So, all right.

If a, T1 plus T2A is typically small. OK, cool. So it was actually kind of crucial in our proof of theorem 1 to talk about the cardinality of the subset instead of the energy because, in that setting, we have the Plünnecke-Ruzsa inequalities.

And we get a great deal of structure from a single subset being small. We get lots of other things. On the other hand, it is nice, eventually, to get to theorem 2, because theorem 2 is stable. If you replace the actual product, A cross A, by a subset, which is almost all of it. So, both looking at cardinality and looking at energy have good features. Energy estimates are stable, undertaking subsets.

And the cardinality estimates have this contagious structure thing. Those are both very helpful. And Balog-Szemeredi-Gowers, this is important because, even though having a big energy is not synonymous with having a small subset, Balog-Szemeredi-Gowers says they're very closely-related. And so, by using Balog-Szemeredi-Gowers, you can often get the best of both worlds.

So if you were to pretend that you could have all the good features of an energy bound, and all the good features of a cardinality bound to see what you could prove, very often, you can prove those things by using Balog-Szemeredi-Gowers to move back and forth. Pablo gave the analysis seminar a few weeks ago. And he gave an overview of recent developments in projection theory.

And he made the claim, or suggestion, that this theorem might be one of the most important or most used theorems in math in the last several decades. It was a little provocative. I thought it was interesting. And I have gradually been appreciating it more and more. I'm learning about it. Anyway, I've been gradually appreciating it more and more. And, yeah, probably, as we keep going, we'll see how this keeps being important.