

YUFEI ZHAO: We're about to embark on a new chapter in this course where I want to tell you about Szemerédi's graph regularity lemma. Szemerédi's graph regularity lemma is a very powerful tool in modern graph theory, developed back in the '70s. Today I want to show you the statement and the proof of this graph regularity lemma. And next time, we'll see how to apply the lemma for graph theoretic applications. And we'll also use it to give a proof of Roth's theorem.

The idea of Szemerédi's regularity lemma is that if you are given a very large graph, G . And it's a fairly robust theorem, so any large, dense graph. And here, "dense" means, let's say, positive x density. Then it is possible to partition the vertex set of this graph G into a bounded number of pieces so that G looks random-like between most pairs of parts.

So for instance, I might produce for you a partition of the vertex set into some number of parts. I'll draw five here. So you give me a graph G . I manage to produce for you this vertex partition so that if I look at between a typical pair of parts, you see here maybe the edge density is close to 0.2 but otherwise, the bipartite graph looks like a random graph in some precise sense I will describe in a bit.

And if you look at what the graph looks like between another pair of parts, maybe now it's a different x density. Maybe it's around 0.4. And again, looks like a random graph with that density.

So in some sense, Szemerédi's regularity lemma is a universal structural description that allows you to approximate a graph by a bounded amount of information. So that's informally the idea. And you can already sense that this can be a very powerful tool. It doesn't matter what graph you input. You apply this lemma, and you get an approximate structural or as later on we'll see, it's also, in some sense, an analytic description of the graph.

So the first part of today's lecture will develop just a statement of this regularity lemma. I'll show you what exactly do I mean by "random-like." Well, first let me give some definitions.

I denote by the letter e if I input a pair of vertex sets, x and y . Here I might, later on, draw the subscript G if it's clear that I'm always talking about some graph G . So this is basically the number of edges between x and y .

And I say "basically" because even though I will draw and depict everything as if x and y are disjoint sets, and that's the easiest case to think about, I'm also going to allow x and y to overlap, and also allow x and y to be the same set, in which case you should read a definition as to what this means. But it's fine to think of it as disjoint sets. So you're looking at a bipartite graph between x and y .

We're also going to look at the edge density between x and y . And this is simply the number of edges divided by the product of the sizes of the sets, so what fraction of the possible pairs are actual edges. So from now on, I'll refer to this quantity as "edge density."

So now, here's the definition of what "random-like" means for the purpose of Szemerédi's regularity lemma. So we define a notion of an epsilon regular pair to be as follows-- throughout, and later on, I will omit even saying this-- G will be some graph. And we're going to be looking at subsets of vertices of G .

And we say that this pair of subsets of vertices is epsilon regular, again, in G , but later on I will even drop saying in G if it's clear which graph we're working with. So we say x and y is epsilon regular if for all subsets A of X , all B subsets of Y that are not too small, so each at least an epsilon proportion of the sets that they live in, we find that the x density between A and B differs from the x density between X and X by no more than epsilon.

Let me draw you a picture. I have sets A and B . So I have sets X and Y in my graph G . And I want to say that the edges between X and Y are epsilon regular, so it's random-like, if the following holds-- that whenever I pick a subset A in the left set and a subset B of the right set, the edge density between A and B is approximately the same as the overall edge density between X and Y .

So in particular, this bipartite graph, for instance, is not really dense in one part and really sparse in another part. Somehow the edges are evenly distributed in this precise manner. So that's the definition of epsilon regular. Yes, question.

AUDIENCE: What is the epsilon for the size of A the same as epsilon for [INAUDIBLE]?

YUFEI ZHAO: The question is here, why are we using the same epsilon here, here, and there? And that's a great question. So that's mostly out of convenience.

So you could use different parameters. And they do play somewhat different roles, but at the

end, we'll generally be looking at one type of epsilons. So we just make our life easier. So you could extend the definition by having an epsilon comma eta, if you like, but it will not be necessary for us, and mostly for simplification. Any more questions? All right.

Now if you have a pair x, y that is not epsilon regular, I just want to introduce a piece of terminology. So you can read from the definition what it means to be not epsilon regular. And sometimes I will say "epsilon irregular," but to be precise, I'll stick with not epsilon regular. Then we can exhibit this A and B that witnesses the irregularity.

So if x, y is not epsilon regular, then their irregularity as, we say it's "witnessed by" some pair A in X and B in Y , satisfying-- basically, you read the definition, and such that the density between A and B differs quite a bit from the density between X and Y . So when I say "to exhibit" or "to witness irregularity," that's what I mean.

Now, there's a bit of an unfortunate nomenclature in graph theory, where previously, we said "irregular graphs" to mean that every vertex is degree D . And now we say "epsilon regular" to mean this. Sorry about that. These are both standard, so usually from context, it's clear which one is meant.

So this is what it means for a single pair of vertex sets to be epsilon regular. But now I give you a graph. And I give you a partition of the vertex set. So what does it mean for that partition to be epsilon regular?

And here's the second definition. So an epsilon regular partition, we say that a partition-- and generally, I will denote partition by curly letters such as that, P . So the partition will divide a vertex set into a bunch of subsets.

So we say that that partition is epsilon regular if the following is true-- if I sum over all pairs of irregular or pairs of vertex sets that are not epsilon regular, so over V_i, V_j not epsilon regular, and sum over the product of their sizes, then what I would like is for the sum to be at most epsilon times the number of pairs of vertices in G . In other words, a small fraction of pairs of vertices, not necessarily edges, but just pairs of vertices, lie between pairs of vertex parts that are not epsilon regular. So for instance, if you do not have epsilon-- if all of your pairs are epsilon regular, then the partition is epsilon regular. But I do allow a small number of blemishes. And that will be necessary.

Just to clarify a subtle point here, here I do allow in the summation i equals to j , although in

practice it doesn't really matter. You'll see that it's not really going to come up as an issue. And one of the reasons that it's not going to come up as an issue is usually when we apply this lemma, we're going to have a lot of parts.

In fact, we can make sure that there is a minimum number of parts. And if none of the parts are too big, then having i equals to j contributes very little to that sum anyway. In particular, if all the set sizes in this partition are roughly the same-- so if they're all roughly $1/k$ fraction of the entire vertex set-- then that statement up there being ϵ regular partition up to changing this ϵ is basically the same as saying that fewer than ϵ fraction of the pairs V_i, V_j are not ϵ regular. And here, if k is large enough, I can even let you make i and j different. It's not going to affect things after small changes in ϵ .

So when it comes to-- so for people who are seeing Szemerédi's regularity lemma for the first time-- I think that's maybe all of you, or most of you-- I don't want you to focus on the precise statements so much as the spirit of the lemma. Because if you get too nitty gritty with is that the same as that ϵ , you get very confused very quickly. So I want you to focus on the spirit of this lemma. I will state everything precisely, but the idea is that most pairs are not ϵ regular. And don't worry too much about if you are allowed to take i equals to j or not.

So now we're ready to state Szemerédi's regularity lemma. And it says that for every ϵ , there exists some constant M depending only on ϵ such that every graph has an ϵ regular partition into at most M parts. You give me the ϵ , for example 1%, and there exists some constant such that every graph has a 1% regular partition into a bounded number of parts. In particular-- and this is very important, make sure you understand this part-- that the number of parts does not depend on the size of the graph.

Now, it's true that for some graphs, maybe you do need very many parts. But the number of parts does not get substantially bigger, or does not exceed this bound, even if you look at graphs that have unbounded size. So it is really a universal theorem in the sense that it's independent of the size of the graph. Any questions about the statement of this theorem? Yes.

AUDIENCE: So in the informal statement at the beginning, you said G was a large, dense graph.

YUFEI ZHAO: That's right.

AUDIENCE: Is the dense condition appropriate anywhere in here?

YUFEI ZHAO: So the question is, why did I say that G is a large, dense graph? And that's a great question.

And that's because if G had a sub-linear number of edges, then I claim that all-- if you look at the definition of epsilon regular pair, and your epsilon is a constant, and if your edge densities are sub-linear, then all of these guys, they are little o of 1. They go to 0.

So trivially, you will satisfy the epsilon regular condition. So if your graph is sparse-- sparse in the sense of having sub-quadratic number of edges-- then you trivially obtain epsilon regularity. And so the theorem is still true.

It's just not meaningful. It's just not useful. But there are settings where having sparse graphs-- and we'll come back to this later in the course-- it's important to explore what happens to sparse graphs. Yeah.

AUDIENCE: So that M is independent of G .

YUFEI ZHAO: Yes, M is independent of G . M depends only on epsilon.

AUDIENCE: M is really large, but there's not enough vertices in the graph.

YUFEI ZHAO: OK, question is, what happens when M is very large, but there are not enough vertices in the graph? Well, if your M is a million, and your graph only has 1,000 vertices, what you can do is have every vertex be its own part. Every vertex is its own part, a singleton partition. And you can check that that partition satisfies the properties. Every pair is a single edge and it's epsilon regular. Yeah.

AUDIENCE: So in the definition, is it sort of like all or nothing? You can either [INAUDIBLE] epsilon regularity [INAUDIBLE]. Do you get anything where if you, like, say, make this more continuous, so you allow for it to be-- you quantify how irregular it is, and then can you make [INAUDIBLE]?

YUFEI ZHAO: OK, so my understanding what you're asking is in the definition up there, the sum is-- we put the pair in the sum of this epsilon regular, and otherwise don't put it. Is there some gradual way to put some measure of irregularity into that sum? And there are versions of regularity lemma that do that, but they are all, in spirit, morally the same as that one there. Yeah.

AUDIENCE: In the informal definition, what does "random-like" mean?

YUFEI ZHAO: So in the informal definition, what does "random-like" mean? This is the formal definition of what "random-like" means. So actually later on in the course, one of the chapters will explore

what pseudo-random graphs are. So pseudo-random graph, in some sense, means graphs that are not random, but behave in some sense like random.

So "random-like" generally just means that in some aspect, in some property, it looks like a random object. And this is one way that something can look like random. So a random graph has this property, but random graphs also have many other properties that are not being exhibited in this definition.

But this is one way that graph can look like random. So that's a great question. And we'll come back to that topic later in the course.

All of these are great questions. So Szemerédi's regularity lemma, the first time you see it, it can look somewhat scary. But I want you to try to understand it more conceptually. So please do ask questions.

Before diving into the proof, I want to make a few more remarks about a statement. It is possible to-- we will prove this version of the regularity lemma. But as I mentioned, it is the spirit of the regularity lemma that I care more about.

And it's a very robust statement. You can add on extra declarations that somehow doesn't change the spirit. And the proof will be more or less the same, but for various applications will be slightly more useful. So in particular, it is possible to make the partition equitable.

And "equitable partition" sometimes is also called an "equipartition," meaning that it has such that all the A_i 's, all the B_i 's have sizes differing by at most 1. So basically, all the parts have the same size up to at most 1, because of divisibility. So let me state a version of regularity lemma for equitable partitions.

So for every ϵ in $(0, 1]$, for every m_0 , there exists a big M such that every graph has an ϵ -regular equitable partition of the vertex set into k parts, where k is at least m_0 , so I can guarantee a minimum number of parts, and at most some bounded number. Again this bound may depend on your inputs ϵ and m_0 , but it does not depend on the graph itself. And you see the slightly stronger conclusion for many applications is more convenient, to use this formulation.

And I will comment on how you may modify the proof that we'll see today into one where you can guarantee equitability. And you see that for this m_0 , m_0 too small, for example, if it's somewhat larger than $1/\epsilon$, when you look at the definition of ϵ -regular

partition, it suffices to check that at most ϵk^2 pairs, V_i, V_j is ϵ regular over $i \neq j$, again up to changing ϵ , let's say, by a factor of 2. So all of these definitions are basically the same up to small changes in the parameters.

Next time, we'll see how to apply the regularity lemma. And we will apply it in the first form, but you see the second form guarantees you a somewhat stronger conclusion, and sometimes more convenient to use. So for example on the homework problems, if you wish to use the second form, then please go ahead. Just make your life somewhat easier, but it essentially captures all the spirit of Szemerédi's regularity. Any questions so far?

I want to explain the idea of the proof of the regularity lemma. And this is a very important technique in this area called the "energy increment argument." Here's the idea. We start with some partition, so for example, the trivial partition-- and by that I mean you only have one part.

All the vertices are in one part. You're not doing anything to the vertex set. It's one gigantic part.

Or if you're looking at some other variant, you can easily modify the proof. So for example, you can also look at an arbitrary partition into little m_0 parts, if you wish to have that as your starting point. So or I'm saying is that this proof is fairly robust.

And we're going to do some iterations. So as long as your partition is not ϵ regular, we will do something to the partition to move forward. And what we will do is look at each pair of parts in your partition that's not ϵ regular. Well, if they're not ϵ regular, then I can find a pair of subsets which are denoted by the A 's that witnesses this non regularity, that witnesses the irregularity.

And we start with some partition. So now let us refine the partition into a partition in even more parts by simultaneously refining the partition using all of these A_i, j 's that we found in the step above. So you start with some partition. If it is not regular, I can chop up the various parts in some way.

So I start with some partition over here. And what we are going to do is, let's say between these two, it's not ϵ regular, so I can find some pairs of vertex sets that exhibits the irregularity. I chop it up.

And I can keep further chopping up the rest of the parts. If these two parts are not ϵ

regular, then I chop it up like that. And I can keep on doing it.

And originally, I have three parts. Now I have 12 parts. And this is a refined partition. And now I repeat until I am done. I am done when I obtain a partition that is epsilon regular.

Now, the basic question when it comes to the strategy is, are you ever going to be done? When are you going to be done? And if this process goes on forever or goes on for a very long time, then you might have a lot of parts. But we want to guarantee that there is a bounded number of parts.

So what we will show is that-- to show that you have a small number of parts, in other words, why does this process even stop-- and in particular, we want it to stop after a small number of steps, after a bounded number of steps. And to do this, we will define some notion called an "energy" of a partition. And this energy will increase.

So first of all, the energy is some quantity that we'll define that lies between 0 and 1. It's some real number lying between 0 and 1. And each step, the energy goes up by some specific quantity.

Therefore, because the energy cannot increase past 1, this iteration stops after a bounded number of steps. And once it's done, we end up with a epsilon regular partition. So that's the basic strategy. And what I want to show you is how to execute that strategy. Any questions so far? Yes.

AUDIENCE: Just to clarify [INAUDIBLE] a bit, if some V_i 's into non-epsilon regular partitions, is it possible for $A_{i,j}$ and $A_{i,k}$ to overlap somehow, right? Just kind of make those into three partitions?

YUFEI ZHAO: So if I understand correctly, you are worried about between different pairs, you might have interactions.

AUDIENCE: Yeah.

YUFEI ZHAO: So you have seen the proof, but I think this is actually a very important and somewhat subtle point, is that I do not refine at each step, I find a pair of witnessing sets. I find all of these witnessing sets all at the same time, and I refine everything all at once.

AUDIENCE: OK, so it's like if you do have overlap between two witnessing sets, that's OK?

YUFEI ZHAO: That is OK, because this step doesn't care. If you have two witnessing sets that overlap, that is

OK. We'll see the proof. Yes.

AUDIENCE: Do you just find one pair of witnessing sets for each V_i, V_j , even though there might be more?

YUFEI ZHAO: Question is, do we find just one pair of witnessing sets even though there could be more? And the answer is, yes. We just need to find one. There could be lots. So if it's not epsilon regular, it might be very not epsilon regular.

And in fact, being a witnessing set is a fairly robust notion. If you just take out a small number of vertices, it's still a witnessing set. Any more questions? Great. So let's take a quick break and then we'll see the proof.

Let's get started with the proof of Szemerédi's regularity lemma. And to do the proof, I want to develop this notion of energy which you saw in the proof sketch. So what do I mean by "energy?"

First, if I-- let me define some quantities. If I have two vertex subsets, U and W , let me define this quantity, q , which is basically the edge density squared. But I normalize it somewhat according to how big U and W are.

I'm going to use the letter n to denote the number of vertices in G . So this is some cube. And for partitions, if I have a pair of partitions, P_U of U into k parts, and the partition P_W of W into l parts, I set this q of P_U and P_W to be the quantity where I sum over basically all pairs, one part from U , one part from W of this q between U_i and W_j .

So this is the density squared. And I'm taking some kind of weighted average of the squared density. So here is a weighted average. If you prefer to think about the special case where this partition is an equipartition, then it is really the average of these squared densities. It's a mean square density.

And finally, for a partition P of the vertex set of G into m parts, we define this q of this partition P to be q of P with itself according to the previous definition. Or in other words, I do this double sum, i from 1 to m , j from 1 to m , q of V_i, V_j . And this is the quantity that I will call the "energy" of the partition.

It is a mean squared density, some weighted mean of the edge densities between pairs of parts in the partition. You might ask, why is it called an energy? So you might see from this formula here, it's some kind of a mean square density, so it's some kind of an average of

squares.

So in particular, it's some kind of an L2 quantity. And there's a general phenomenon in mathematics, I think borrowed from physical intuitions, that you can pretty much call anything that's an L2 quantity an energy. And so that's, I think, where the name comes from.

So this is the important object for our proof. And let's see how to execute a strategy, the energy increment argument outlined on the board over there. So we want to show that you can refine a partition that is not epsilon regular in such a way that the energy goes up.

And to do that, let me state a few lemmas regarding the energy of a partition under refinement. And the point of the next several lemmas is that the energy never decreases under refinement, and it sometimes increases if your partition is not epsilon regular. So the first lemma is that if you look at the energy between a pair of partitions, it is never less than the energy between the two vertex sets.

So for instance, if you have U and W like that, and I partition them into P_u and P_w , and I measure the energy, just basically the squared density between U and V versus summing up the individual squared densities after the partition, the left side is always at least as great as the right side. So this is really a claim. It's a fairly simple claim about convexity, but let me set it up in a way that will help some of the later proofs.

So let me define a random variable, which I call Z , in the following way. So here's a process that I will use to define this random variable. I will select x , little x , to be a vertex uniformly chosen from U , from the left vertex set. And I will select a vertex y uniformly chosen from W .

x and y , they fall into some part in the partition. So suppose U_i is the part where x falls, and W_i is the set in the partition where y falls. So U_i is a member of this partition. W_i is a member of the other partition of W . Then I define my random variable Z to be the x density between U_i and W_j . So it's W_j .

So that's the definition. So pick x randomly. Pick y randomly. Suppose x falls in U_i . Suppose y falls in U_j . Then Z is the x density between these two parts.

So Z is some random variable. Let's look at properties of this random variable. First, what is this, it's expectation? It's a discrete random variable.

And you can easily compute all of these quantities by just summing up according to how Z is

generated. So I look overall, i and j . What's the probability that x falls in U_i ? It is the size of U_i as a fraction of U .

What's the probability that y falls in W_j ? It's the size of W_j as a fraction of size W . And then Z is this quantity here. So this is what I find to be the expectation of Z .

But you see the density multiplied by the product of the vertex set sizes, that's just the number of edges between U and W . And you sum over all the i, j 's. So that, which is simply the edge density between U and W . So that's the expectation of the Z variable.

On the other hand, what's the second moment? In other words, what's the expectation of the square of Z ? Again, we do the same computation. First part is the same.

The second part now becomes a d squared. And look at how we define energy. This quantity here is basically the energy q between the partition U and the partition of W , except there's normalization. That's not quite the same as the one we used before. So we will just put in that normalization.

So now you compare the expectation of Z versus the expectation of Z squared. And we know by convexity that the expectation of Z squared is at least as large as the expectation of Z , that quantity squared. But if you plug in what values you get for these two guys, you derive the inequality claimed in Lemma 1. You have to cancel some normalization factors, but that's easy to do.

So that's the first lemma. So the first one is just about a pair of parts, and identify partition, each part, what happens to the energy between this pair. And the second one is a direct corollary of the first one.

It says that if you have a second partition, P prime that refines P , then the energy of the second partition, the refinement, is never less than the energy of the first partition. And it is a direct consequence of the first lemma, because we simply apply this lemma to every pair of parts in P . Between every pair of parts, the energy can never go down. So overall, the energy does not go down.

And finally, so far, we've just said that the partitions can never make the energy go down. But in order to do this proof, we need to show that the energy sometimes goes up. And that's the point of the third lemma.

The third Lemma tells us that you can get an energy boost. So this is the Red Bull Lemma. You can get an energy boost if you are feeling irregular.

So if U, W is not epsilon regular, and this epsilon regularity is witnessed by U_1 in U and W_1 in W , then I claim that the energy obtained by chopping U into U_1 and its complement against W_1 , against the complement of W_1 and W , so here again U and W , I find a witnessing set for their irregularity. And now I partition left and right according to-- chop each part into two. So this energy between this partition into two on both sides is bigger than the original energy plus something where we can gain. And this something where we can gain turns out to be at least epsilon raised to the power 4 times the size of U , size of W , divided by n squared.

Can you prove it? Let's define Z the same as in the previous proof, as in the proof of Lemma 1. In Lemma 1, we just used the fact that the L_2 norm of Z , the expectation of the square, is at least the square of the expectation. But actually, there are differences in it. It's called a "variance."

The variance of Z is the difference between these two quantities. I know that it's always non-negative. So if you look at how we derived the expectation of Z and expectation of Z squared, you immediately see that its variance we can write as, up to a normalizing factor, the difference between this energy on one hand and the energy between U and W , namely the mean square of the normalization.

On the other hand, a different way to calculate the variance is that it is equal to the expectation of the deviation from the mean squared. So let's think about the deviation from the mean. I am choosing a random vector in, on the left, U , and another random point, random vertex on the left, and a random point on the right.

In the event that they both lie in the sets that witness the irregularity, so in the event where x falls here and y falls here, which occurs with this probability, see that this quantity here is equal to the density between U_1 and W_1 minus the density-- and this expectation of Z is just the density between U and W . So interpreting this expectation for what happens when x falls in U_1 and when y falls in W_1 , ignoring all the other events, because the quantity is always non-negative everywhere.

But now from the definition of epsilon regularity, or rather the witnessing of epsilon irregularity, you see that this U_1 is at least an epsilon fraction of U . W_1 's at least an epsilon fraction of W . And this final quantity here is at least epsilon inside, so at least epsilon squared. So here we're

using all the different components of the definition of epsilon regular. Yes.

AUDIENCE: What happens if we're dividing with more witnessing sets?

YUFEI ZHAO: So you're asking what happens if we divide with more witnessing sets? So hold onto that thought. So right now, I'm just showing what happens if you have one witnessing set. Any more questions?

So here we have epsilon to the 4th. And if you're putting the normalization comparing these two interpretations, you'll find the inequality claimed by the lemma. So now we are ready to show the key part of this iteration.

I'll show you precisely how this iteration works, and show that you always get an energy boost in the overall partition. So I'll call the next one Lemma 4. And this says that if you have a partition P of the vertex set of G into k parts, if this partition is not epsilon regular, then there exists a refinement called Q where every part $V_{sub\ i}$ is partitioned further into at most $2k$ parts, and such that the energy of the new partition Q increases substantially from the previous partition P .

And we'll show that you can increase by at least epsilon to the 5th power, some constant in epsilon. So if you look at the strategy up there, if you can do this every step, then that means that the number of iterations is bounded by $1/\epsilon^5$. So to prove this lemma here, we will use the three lemmas up there and put them together.

So for all the pairs i, j such that $V_{sub\ i}, V_{sub\ j}$ is not epsilon regular, as outlined in the proof in the outline up there, we will find this $A_{i, j}$ in V_i , and $A_{j, i}$ in V_j that witness the irregularity. So do this simultaneously for all pairs i, j , where the V_i, V_j is not epsilon regular.

Now what we're going to define Q as is the common refinement. So take all of, just as indicated in that picture up there, simultaneously take all of these A 's and use them to refine P . Starting with P , starting with a partition you have, simultaneously cut everything up using all of these witnessing sets. Now, we only have witnessing pairs for pairs that are not epsilon regular. If they're epsilon regular, you don't worry about them.

One of the claims in the lemma now is-- this is the Q that we'll end up with. We'll show that this Q has that property. So one of the claims in the lemma is that every V_i is partitioned into at

most 2^k parts. So I hope that part is clear, because how are we doing the refinement?

We're taking V_i . It's divided into parts using these $A_{i,j}$'s, one j coming from each pair that is irregular with V_i . So I'm cutting up V_i using at most k sets, so one coming from each of the other possible. Maybe fewer than k -- that's fine-- but at most k sets are used to cut up each V_i . So you have at most 2^k parts once you cut everything up.

But the tricky part is to show that you get an energy boost. So let's do this. How do we show that you get an energy boost?

We're going to put the top three lemmas together. First, we want to analyze the energy of Q . So let's write it out.

So the energy of Q is the sum over this energy of individual partitions of the V_i 's. And by this $P_{\text{sub } V_i}$, $P_{\text{sub } V_j}$, I mean the partition of V_j given by Q . So what happens after you cut up the V_i -- that's what I mean by $P_{\text{sub } V_j}$, or rather I should call it $Q_{\text{sub } V_i}$, $Q_{\text{sub } V_j}$.

By Lemma 2, we find that-- so let me separate them into two cases. The first case sums over i, j such that V_i, V_j is epsilon regular. And by Lemma 1, so here we're using Lemma 1, we find that this quantity here cannot be less than the Q of V_i, V_j . So take those two parts.

Before and after the refinement by Q , the energy cannot go down. So I don't worry too much about pairs that are epsilon regular. But no let me look up here that are not epsilon regular.

So what we will do now is even though-- so let's look at that picture up there. So let's focus on what I drew in red. So let's focus between 1 and 2. So suppose the shaded part is the witnessing sets.

The witnessing sets got cut up further by other witnessing sets. But I don't have to worry about them because Lemma 2 or Lemma 1, really, tells me that I can do an inequality where I go down to just comparing the energy between this partition of two parts, this single witnessing set and its complement, versus what happens in its partner.

So in other words, over here, the Q of this pair, I am saying that it is no less than if I just look at what happens if you only cut up these two sets using the red lines.

Let's go on. Applying Lemma 3, the energy boost lemma, the first part stays the same. So this first part stays the same. And the second part, now, because I'm looking at witnessing sets for

irregularity, I get this extra boost. So this goes back to one of the questions asked earlier, where in Lemma 3, I don't have to now worry about what happens if you have further cuts, because I only need to worry about the case where I only have a single cut between the epsilon irregular pairs.

So putting it together, we see that the previous line is at least, if you sum over the Q 's of all the pairs plus this extra epsilon to the 4th term for all pairs that are not epsilon regular. I'm applying monotonicity of energy for the types, for pairs that are epsilon regular, an energy boost for pairs that are not epsilon regular. And for the latter type, I obtain this boost.

Now remember what's the definition of an "epsilon regular partition." Unfortunately, it's no longer on the board, but it says that this sum over here, if it is an epsilon regular partition, it is at most epsilon. So if it is not epsilon regular, we can lower bound it. And that's indeed what we will do.

The first sum here is, by definition, Q of the partition P . And the second sum, by the definition of epsilon regular, is at least epsilon to the power 5. So here we're using the definition of epsilon regular partition, namely, that a large fraction, so at least an epsilon fraction, basically, of pairs of vertex sets are not epsilon regular, but in this weighted sense.

And that finishes the proof of Lemma 4 up there. Any questions so far? All right, so now we are ready to finish everything off, and prove Szemerédi's regularity lemma.

So let's prove Szemerédi's regularity lemma. Let's start with the trivial partition, meaning just one large part. And we are going to repeatedly apply Lemma 4 whenever the partition at hand is not regular, whenever the current partition is not epsilon regular.

So let's look at its energy. The energy of this partition-- so this is a weighted mean of the edge density squared, so it always lies between 0 and 1, just from the definition of energy. On the other hand, Lemma 4 tells us that the energy increases by at least epsilon to the 5th power at each iteration.

So this process cannot continue forever. So it must stop after at most epsilon to the minus 5th power number of steps. And when we stop, we must result in an epsilon regular partition, because otherwise, you're going to continue applying the lemma and push it even further. And that's it. So that proves Szemerédi's graph regularity lemma. Question.

AUDIENCE: It's going to be some really big value of M .

YUFEI ZHAO: OK, let's talk about bounds. So let's talk about how many parts. So how many parts does this proof produce? We can figure it out.

So we have some number of steps. Each step increases the number of parts by something. So if P has k parts, so then Lemma 4 refines P into at most how many parts?

AUDIENCE: 2 to the k

YUFEI ZHAO: Yeah, so k times 2 to the k . And I have many iterations of this guy. So some of you are already laughing, because it's going to be a very large number. In fact, because it's going to be so large, it makes my calculations slightly more convenient.

It really doesn't change the answer so much if I just bound k to the 2 to the k by 2 to the 2 to the k . So the final number of parts is this function iterated on itself ϵ to the minus 5 times. So it's a power of 2 of height at most 2 to the ϵ to the 5.

It's a finite number, so it depends only on ϵ and not on the size of your graph. And this is the most important thing. It does not depend on the size of your graph. It is quite large.

In fact, even for reasonable values of ϵ , like 1% or even 10%, this number is astronomically large. And you may ask is it really necessary, because we did this proof, and it came out fairly elegantly, I would say it, how the proof was set up. And you arrived at this finite bound.

But maybe there's a better proof. Maybe you can work harder and obtain somewhat better bounds. So you can ask, is it possible that the truth is really somehow much smaller? And the answer turns out to be no.

So there is a theorem by Tim Gowers which says that there exists some constant. The precise statement, again, is not so important, but based on what I just said, you cannot improve this bound given by this proof. So for every ϵ small enough, there exists a graph whose ϵ regular partition requires how many parts?

So the number of parts at least this tower of 2 of height some ϵ to the minus c . So really it's a tower of exponentials of size, essentially polynomial in $1/\epsilon$. So maybe you can squeeze the 5 to something less. Actually, we don't even know if that's the case, but certainly you cannot do substantially better than what the proof gives.

So Szemerédi's regularity lemma is an extremely powerful tool. And we'll see applications that are basically very difficult to prove. And for some of these applications, we don't really know other proofs except using Szemerédi's regularity lemma.

But on the other hand, it gives terrible quantitative bounds. So there is a lot of interest in combinatorics where once you see a proof that requires Szemerédi's regularity lemma, or that is first proved using this technique, to ask can it be used using some other technique? In fact, Szemerédi himself has worked a lot in that direction, trying to get rid of the uses of his lemma. Any questions?

AUDIENCE: How could you modify it for equipartitions?

YUFEI ZHAO: OK, great. Question is, how can we modify it for equipartitions? So let's talk about that. So it's a fantastic question. So look at this proof and see what can we do if we really want all the parts to have roughly the same size, let's say differing by at most 1.

So how to make the epsilon regular partition equitable? Any guesses? Any attempts on what we can do?

I mean, basically it's going to follow this proof. As I said, the spirit of Szemerédi's regularity lemma is what I've shown you. But the details and executions may vary somewhat depending on the specific purpose you have in mind. Yeah.

AUDIENCE: Can we just add-- [INAUDIBLE] add things to the smaller part because we know that-- by the fact that it's not [INAUDIBLE] that parts aren't too small?

YUFEI ZHAO: OK, so you're saying we're going to add something or to massage the partition to make it epsilon--

AUDIENCE: Add vertices to the smaller parts of the partition.

YUFEI ZHAO: Add vertices to the smaller parts of the partition, now when are you going to do that?

AUDIENCE: When they're-- like so you do the refinement, then when they're not [INAUDIBLE]

YUFEI ZHAO: So you want to do this at every stage of the process.

AUDIENCE: Yes. [INAUDIBLE]

YUFEI ZHAO:

I like that idea. So here's what we're going to do. So we still run the same process. So we're going to have this P , which is the current partition. So I have current partition.

And as before, we initially have it as either the trivial partition, if you like, or m arbitrary equitable parts. Start with something where you don't really care about anything except for the size. And you run basically the same proof, where if your P is not ϵ regular, then do what we've done before, so basically exactly the same thing.

We refine P using pairs witnessing regularity, same as the proof that we just did. And now we need to do something a little bit more to obtain equitability. And what we will do is right after-- so each step in iteration, right after we do this refinement, so after we cut up our graph where maybe some of the parts are really tiny, let's massage the partitions somewhat to make them equitable.

And to make our life a little bit easier, we can refine the partition somewhat further to chop it up into somewhat smaller resolution. And this part, you can really do it either arbitrarily or randomly. Some ways may be slightly easier to execute, but it doesn't really matter how you do it. It's fairly robust.

You refine it further. And basically, I want to make it equitable. Sometimes, you can just do that by refining, but maybe if you have some really small parts, then you might need to move some vertices around, so I call that "rebalancing." So move and merge some vertices, but only a very small number of vertices, to make equitable.

So you run this loop until you find that your partition is ϵ regular. Then you're done. Whenever you run this loop, because we're doing the second step, your partition is always going to be equitable.

But we now need to control the energy again to limit the number of steps. And the point here is that the first part still is exactly the same as before, where the energy goes up by at least ϵ to the minus 5. But the second part, the energy might go down, because we're no longer refining, just refining. Because we're doing some rebalancing.

But you can do it in such a way that the amount of rebalancing that you do is really small. You're not actually changing the energy by so much. So I'll just hand wave here, and say that we can do this in such a way where the energy might go down, but only a little bit.

So you're only changing a very small number of vertices, very small fraction of vertices. So if

you change only an epsilon fraction of vertices, you don't expect the energy, which is something that comes out of summing pairs of vertex parts, to change by all that much. So putting these two together, you see that the energy still goes up by, let's say, at least $1/2$ of epsilon to the 5th power.

And so then, the rest of the proof runs the same as before. You finish in some bounded number of steps. And you result in an equitable partition that's epsilon regular.

I don't want to belabor the details. I mean, here, there's some things to check, but it's, I think, fairly routine. It's more of an exercise in technical details.

But the thing that actually is somewhat important is there's a wrong way to do this. I just want to point out that what's the wrong way to do this, is that you apply regularity lemma, and you think now it has something that's epsilon regular. Then I massage it to try to make it equitable at the end.

And so if I don't look into the proof, I just look at a statement of Szemerédi's regularity lemma, and I get something that's epsilon regular, I say I'm just going to divide things up a little bit further, that doesn't work. Because the property of being epsilon regular is actually not preserved under refinement. So look at the definition.

You have something that's epsilon regular. You refine the partition. It might fail to be epsilon regular.

So you really have to take into the proof to get equitability. So just to repeat, a wrong way to try to get equitability is to apply regularity lemma, and at the end, try to massage it to get equitable. That doesn't work. Next time, I will show you how to apply Szemerédi's regularity lemma.