STOCHASTIC APPROXIMATION: THEORY AND APPLICATIONS

Dr. Gugan Thope

Department of Computer Science and Automation

Indian Institute of Science

Week 5

Lecture 22

Concluding the Convergence Proof: Internal Chain Transitivity, Connectedness, and Invariance

Hello and Namaste, everyone. Welcome to Lecture 22 of this NPTEL course on Stochastic Approximation. We are in Week 5, wherein we are combining all the ideas, including martingale convergence and ODE behaviors, to discuss the asymptotic behaviors of stochastic approximation algorithms. Right? And we stated some theoretical results that talked about this asymptotic behavior, and we started a part of the proof in the previous class. We will try to finish that proof in today's class.

Right? Okay. So, let us recall the technical details of what we have discussed so far. So, the theorem that we are trying to prove is the following. It states that suppose we have some generic stochastic approximation algorithm with this update rule, and further suppose these four assumptions from A1 to A4 hold. Then, the Xn sequence generated by this update rule converges to a non-empty, compact, connected, and internally chain-transitive invariant set of the limiting ODE X dot of t equals H of X of t.

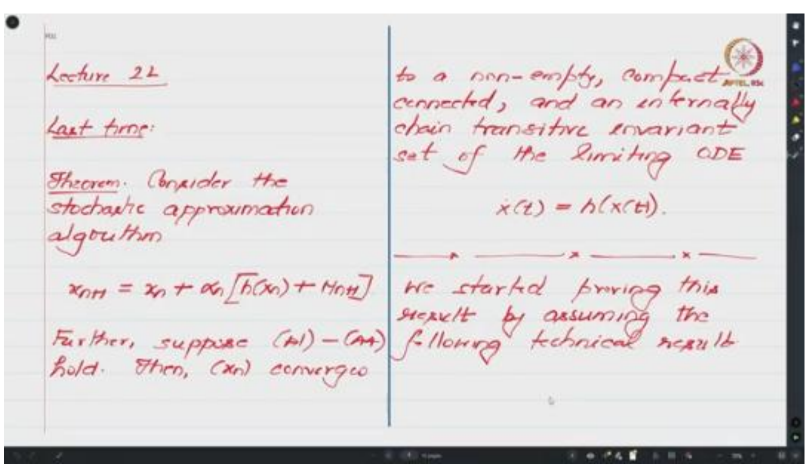

$$\dot{X}(t) = h\big(x(t)\big)$$

Again, I would like to highlight that this is a stochastic algorithm, which means the iterates generated on different runs of the algorithm are going to be different. However, the limiting ODE is the same across all these runs. And, towards proving this result, we assumed the following technical result, which said that if these assumptions A1 to A4 hold, then almost surely, for any t bigger than or equal to 0, the distance between the linearly interpolated trajectory and the solution of the ODE—which is started from the linear interpolated value of the trajectory at time s— the distance between them over this window goes to 0 as S tends to infinity.

$$\lim_{s\mapsto\infty} \sup_{t\in[s,s+T]} ||\bar{X}(t) - x^s(t)|| = 0$$

We also had an analogous result where, instead of looking at the window from S to S plus T, we looked at the window from S minus T to S. And we said that also goes to 0.

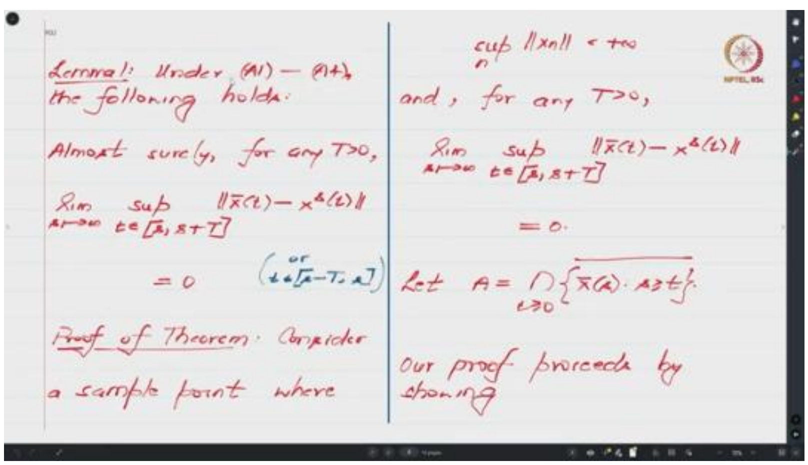

And then we started proving this theorem that we stated on the previous slide, right? And you know, in the first half of that proof, we said that let us consider a sample point where both condition A4 and the conclusions of this lemma simultaneously hold—that is, this is less than infinity, and for any t bigger than or equal to 0, this is equal to 0, and also the analogous, you know, conclusion where the window is s minus t to s, right? So all these three things hold, and then we said, now let us consider A, which is the intersection of these sets. So, in some sense, this is the tail of the linearly interpolated trajectory—in particular, this is the tail starting from time t—and we are looking at the intersection of the closures of these tails of this linearly interpolated trajectory. And we refer to this as set A, and we said that Xn actually converges to A. That is what we proved last time.

$$\sup_n ||x_n|| < +\infty$$

For any $T > 0$,

$$\lim_{s \mapsto \infty} \sup_{t \in [s, s+T]} ||\bar{x}(t) - x^s(t)|| = 0$$

$$A = \cap_{t \geq 0} \overline{\{\bar{X}(s) : s \geq t\}}$$

And now, what we are going to do is we are going to show that A has the desired properties. So, you know, we are going to focus on this part in this class.

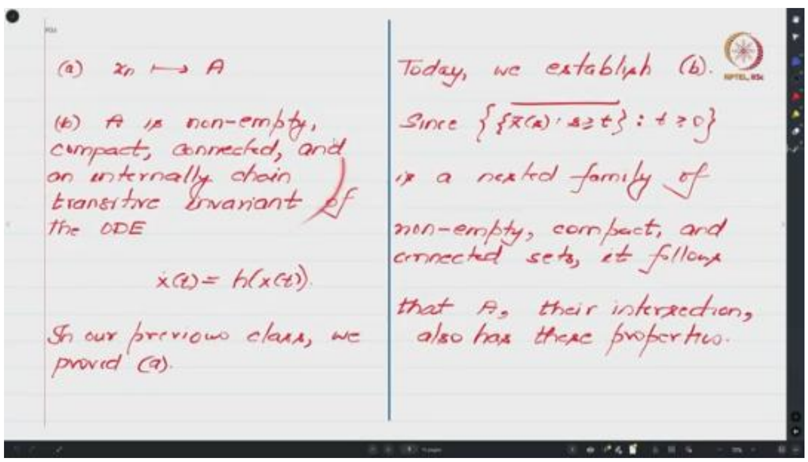

So, let us begin by focusing on these three aspects—that is, non-empty, compact, and connected, right? And that is easy to show in the following sense, right? So, you have this, you know, closure of the tail of your linearly interpolated trajectory for different values of t, and for any t, this is actually a non-empty, compact, and path-connected set.

$$\left\{ \overline{\{\bar{x}(s) : s \geq t\}} : t \geq 0 \right\}$$

Is that okay? I mean, this is a linear interpolation, and hence it is trivially path-connected, right?

And importantly, this family is actually a nested family. Is that okay? And because it is a nested family, right? One can conclude that because each of these sets has these properties of being non-empty, compact, and connected, okay? Their intersection also has these properties, right?

So, that already tells us that A is non-empty, compact, and connected, okay? So, now it remains to show that A is actually invariant and internally chain transitive. So, toward showing these two parts, we are going to prove this claim. So, what does this claim say? It says that if you give me any element in A, then there is a sequence Sn.

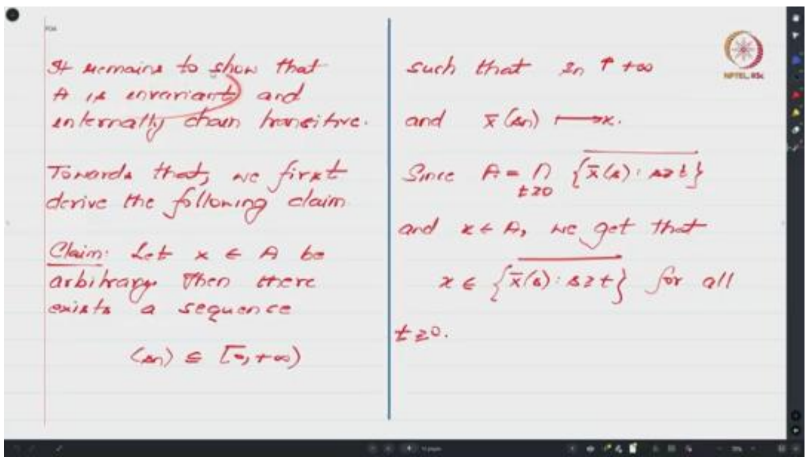

in 0 to infinity such that if you look at the value of your linearly interpolated trajectory at these specific time instances Sn, then they will converge to x. So, let me repeat this claim again. It says that if you give me any x in A, then I can find a subsequence generated from the linearly interpolated trajectory that converges to x. So, this will play a very important role in proving that A is invariant and internally chain transitive. So, how do we go about proving this? So, the first thing we need to observe is that because your X bar SN converges to X, right?

So, in other words, this is the limit of this X bar SN, right? And the limit, you know, is a limit if you only consider the tail of the sequence as well. It helps us conclude that your X, So, I think let me restate this. We have presumed that X is actually in A, and A is the intersection of this closure of this thing. Since X belongs to A, one can conclude that actually X belongs to this closure of this tail for all t greater than or equal to 0.

And now we have to construct this sequence Sn, and the way we will do it is by first setting S0 equals 0, and for any n bigger than or equal to 0, we are going to find an Sn that is bigger than Sn minus 1. That is, we are going to do it in an inductive way. That is, we now have S0, then we will find S1 using some recipe. Once we have found S1, we will find S2 using the same recipe, and so on. So, once you have Sn minus 1, the way you construct Sn is in the following way.

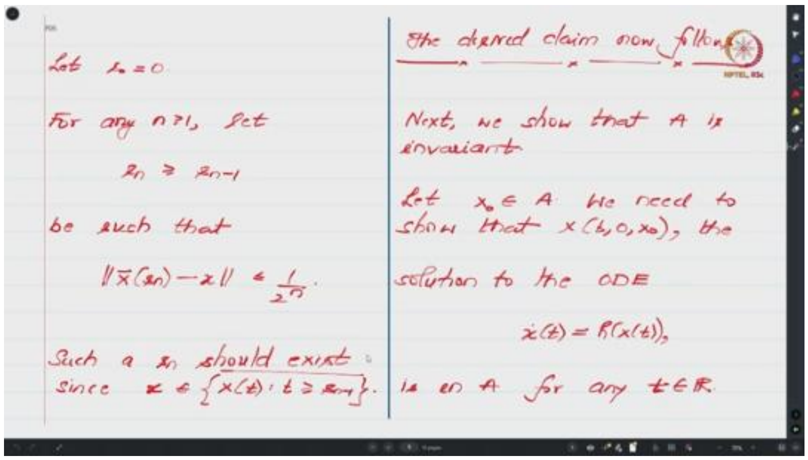

You identify an Sn which is bigger than or equal to Sn minus 1, such that your distance from X bar Sn and X, this distance is less than, let us say, 1 over 2 raised to n. So, why would such an X bar SN exist? Well, recall that X lies in this closed set, right? So, because it lies in this closed set, I can find a point in this set, okay, which is close to X to an arbitrary distance. In this case, we are wanting to find a point which is less than 1 over 2 raised to n.

And such a point should indeed exist because this is a closed set. Now, since for every n, we are able to find a point that is just 1 over 2 raised to n away, and this expression goes to 0 as n tends to infinity, one can conclude that X bar of Sn indeed converges to X, and this is what we needed to show over here. So, this completes the proof of this claim.

$$s_n \geq s_{n-1}$$

$$\left|\left|\bar{x}(s_n) - x\right|\right| \leq \frac{1}{2^n}$$

$$x \in \overline{\{x(t) : t \geq s_{n-1}\}}$$

So, now we are going to use this claim to show that A is actually an invariant set. So, towards that, we will consider an arbitrary point in A, and our goal is to show that if I start a solution trajectory of the limiting ODE from this initial point, then that solution trajectory will lie within A for any T in R. That is what invariant means.

So, if you initialize the solution trajectory within A, And if the solution trajectory remains within A both in forward time and in backward time, that will help us conclude that this set A is actually invariant, and we are going to prove that. Now, from the previous scheme, we know that there exists a sequence of time instances such that this sequence of time instances monotonically increases to plus infinity, right? And X bar of SN converges to X naught, right? Where X naught is this arbitrary initialization that we have chosen, right?

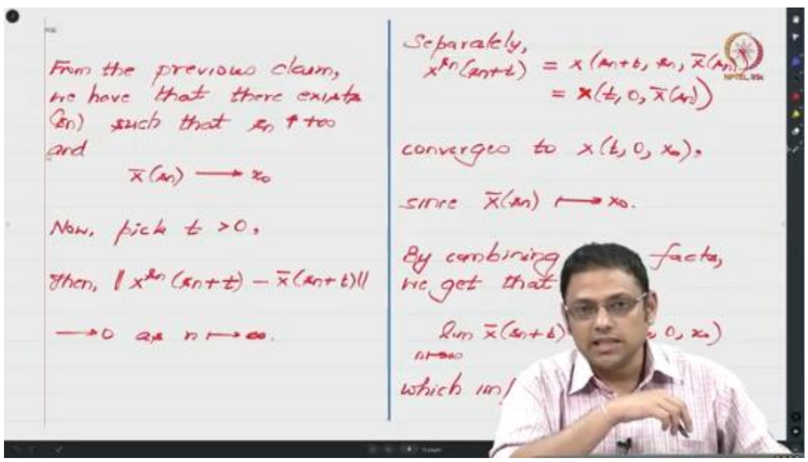

Now, what we are going to do is we are going to pick an arbitrary T, right? And we are going to show that your X... T0X0 belongs to A. So, this is what we are going to show. And toward showing this, what we are going to first do is, you know, since this T is frozen, we are going to appeal to that technical lemma to conclude that this must be true as Sn goes to infinity. So, let us see what this says.

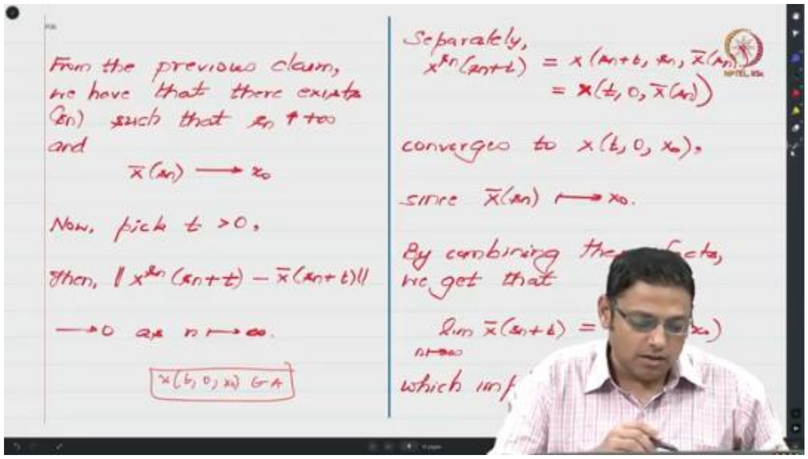

What this is saying is that, You look at X bar SN plus T, right? SN is this sequence that we have already identified. Go T distance ahead, right? And look at the value of the linearly interpolated trajectory and compare this with the value of the solution of your limiting ODE, which passes through X bar of SN at time SN.

Is this okay? So, this is what this is, and you look at the value of this solution trajectory at time Sn plus t, right? And you look at the distance between them. Since this t is fixed, we can appeal to that lemma to conclude that this distance must actually go to 0. I mean, that lemma we had a supremum, right?

Here, we are only focusing on one specific time instance, and hence one can see that this is an easy conclusion from that technical lemma. Now, separately, we can invoke the continuity of solutions with respect to the initial points to conclude the following. In particular, observe that X SN SN plus T. This is the solution of your limiting ODE. This, in our notations that we had introduced previously, can be expressed in this way. So, this solution trajectory is basically the one which, at time Sn, passes through X bar of Sn, and we are interested in looking at its value at Sn plus T. So, this expression, because of the autonomous nature of the ODE,

can be equivalently expressed as follows. That is, at time 0, you start at X bar SN, right? And you are interested in looking at the value of the solution trajectory at time T. Now, we know that X bar SN converges to X naught. This is what we know.

In other words, the initial values of this solution trajectory converge to X naught, and since the solution trajectories are continuous with respect to the initializations, one can conclude that this expression for a fixed T actually converges to XT 0 X naught. By combining this fact, and this fact, one can conclude that the limit of x bar sn plus t—right, this is close to this—and this, as sn becomes larger and larger, becomes close to this. So, those two facts can be combined to conclude that the limit of this actually equals this expression over here, right? Now, because this is the limit of a subsequence generated by your linearly interpolated trajectory, one can conclude that this limit actually lies in A. And since this limit is equal to this, it allows us to conclude that this expression is also in A. And since this t was arbitrary, one can conclude that for any positive time t, this actually lies in A, and hence we have now managed to show that A is positively invariant.

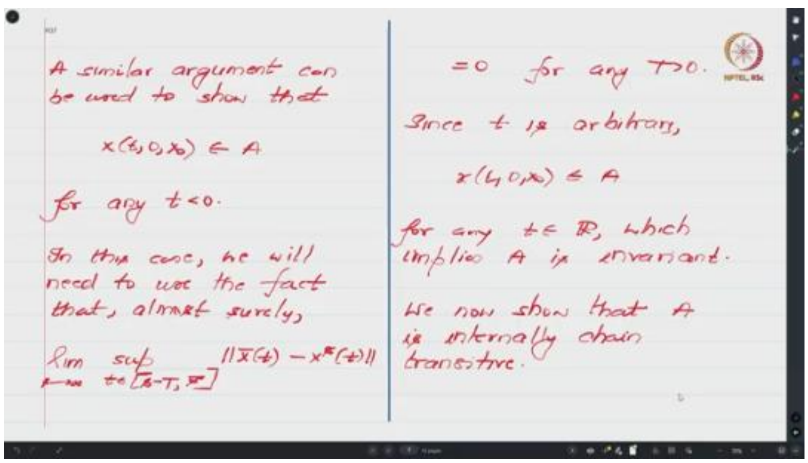

And we can, in a similar fashion, show that A is negatively invariant. For that, we have to repeat the same idea by picking T less than 0 and making use of the second part of that technical lemma. So, recall that we were working with this window from S to S plus T in the previous slide. Now, for showing negative invariance, we will look at the window from S minus T to S and repeat the arguments to conclude that X T 0 X naught lies in A even when T is negative.

And since for any T—both positive and negative Ts—we have managed to show that this belongs to A, we can now conclude that your set A is actually invariant. So, recall that it was remaining to show two things. One is that your set A is invariant, and the second thing that was needed to show was that the set A is internally chain transitive. So, we will now focus on showing that your set A is internally chain transitive. So, you know the formal definition of this internally chain transitive was given two lectures before.

At a very high level, you know, a set being internally chained transitive means that it is path-connected in a very loose way. Path-connected means that, you know, you can find a sequence of solution trajectories which takes you from X to Y through some intermediate points. And when I say path-connected in a loose way, the idea is that you start at X. Run a solution trajectory of your limiting ODE, right? And it will go to the epsilon neighborhood of your next intermediate point, right? After some time t greater than or equal to 0, and then, you know, you start a solution from this intermediate point, and we can show that, you know, after some time capital T.

It will go to the epsilon neighborhood of the subsequent point, and so on and so forth. And in this way, if you start, you know, the solution trajectory of the limiting ODE from the last but one point, you will actually reach the epsilon neighborhood of Y, right, after a time t greater than 0. And, you know, this loose connectedness holds for any epsilon and any t greater than 0. So, we need to now show that indeed this set A that we have has this desired property. So, towards showing that, what we will do is we will pick two arbitrary points X and Y in A and show that this is path-connected in this loose way that I described.

So, formally, what we will do is we will pick some arbitrary epsilon and t greater than 0. And what we will now do is we will pick some delta. Such that delta is strictly less than epsilon over 4 and has the property that if you choose two initializations, x0 and x0 prime, which are delta distance away, then the solution trajectory started at x0 and x0 prime, right, within this window are epsilon over 4 away. So, recall that your solution trajectories are continuous with respect to your initial values. So, if the initial values are sufficiently close, then in any given window, the solution trajectories will also be close.

Also notice that the window over here is 2t. This fact will actually be needed later on. So, keep that thing in mind. So, you know given this epsilon and t we have now chosen a delta such that delta itself is less than epsilon over 4 and it ensures that if the initializations are only delta away then the solution trajectory started at x0 and x0 prime in this window are at most epsilon over 4 away. Now what we are going to do is we are going to pick

N naught bigger than equal to 1 such that for any S bigger than equal to T naught your X bar S lies in the delta neighborhood of A. So, this is your delta neighborhood or the delta expansion of of A and such a Tn0 indeed exists because we have shown that in the previous class that eventually the linearly interpolated trajectory lies in any desired neighbourhood of A. So, here we are focusing on this delta neighbourhood and we are picking a Tn0 such that for all s bigger than equal to Tn0, your X bar s lies in this delta neighbourhood. And on top of that, we also require that this Tn0 satisfy the property that for any S bigger than equal to Tn0, the value of X bar T that is your linearly interpolated trajectory and the solution of your limiting OD starting from X bar S, their distance in this window of size 2T, I think I should again add 2 over here. This distance should be delta away, right? So, we are picking this N0 such that this and this condition simultaneously hold and such a TN0 should exist because we have previously shown that, you know, the tail lies in the delta neighborhood.

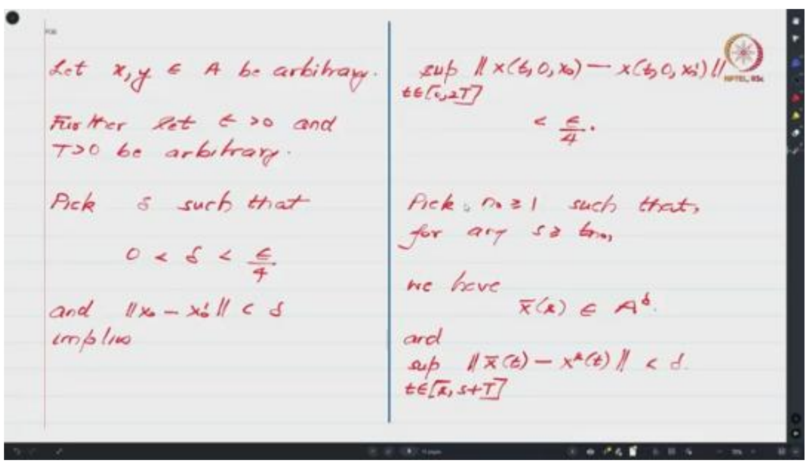

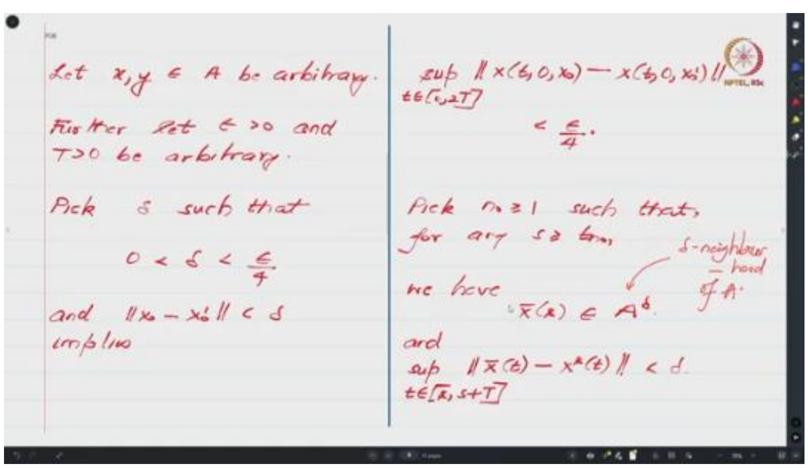

And also that as your S becomes larger and larger this goes to 0 hence eventually this should also be less than delta. So, indeed such a Tn0 exists. Now what we are going to do is we are going to pick N2 and N1 such that N2 is bigger than N1 and N1 is bigger than N0. such that your value of the linearly interpolated trajectory is delta close to X. Recall that X and Y were these arbitrary points that we had chosen. So, X bar TN1 is delta close to X and X bar TN2 is delta close to Y. So, pictorially, if we can presume that this is your set A, right?

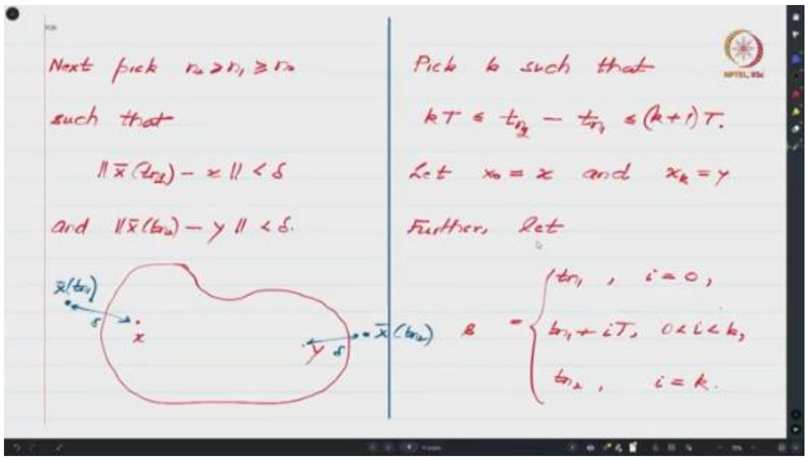

And let us say this is your X and Y that have been arbitrarily chosen. What we are saying is that your X bar TN1 should be delta away from X, and your X bar TN2 again should be delta away from Y. So, notice that X and Y lie in A, but your X bar TN1 and X bar TN2 could lie outside A as well. And, you know, to keep it general, I have actually drawn them outside A, right? All we are saying is that as the value of S becomes larger and larger, okay, the value of your linearly interpolated trajectory will get closer and closer to A.

So, now given that you have this X bar TN1, which is delta away from X, and X bar TN2, which is delta away from Y, what we are going to do is we are going to pick K such that K times T and K plus 1 times T upper bound TN2 minus TN1. So, TN2 minus TN1 is some finite number. And, you know, since T is strictly bigger than 0, if you take or multiply T by K, this number is going to keep increasing. So, identify the last value or the largest value of K such that KT is less than TN2 minus TN1. So, that will automatically ensure that your K plus 1 T is upper bounding this quantity over here.

So, you know, pictorially what it means is that, you know, suppose you manage to show that your K is 3. Now, K could be 3, 4, whatever it is, depending on the value of TN2 minus TN1. Let us say that your, you know, K turns out to be 3, right? So, what it means is that, okay, so if you look at, let us say the value of your linearly interpolated trajectory, right? At TN1 plus T, right, then your linearly interpolated trajectory at TN1 plus 2T, right, and then your linearly interpolated trajectory, which is X bar TN2, sorry, TN1 plus 3T.

Is this okay? And then, if you look at the value, let us say at X bar Tn1 plus 4T. So, somewhere between these two values, your X bar of Tn2 lies. Okay? So, that is the interpretation of this K that is chosen in this fashion.

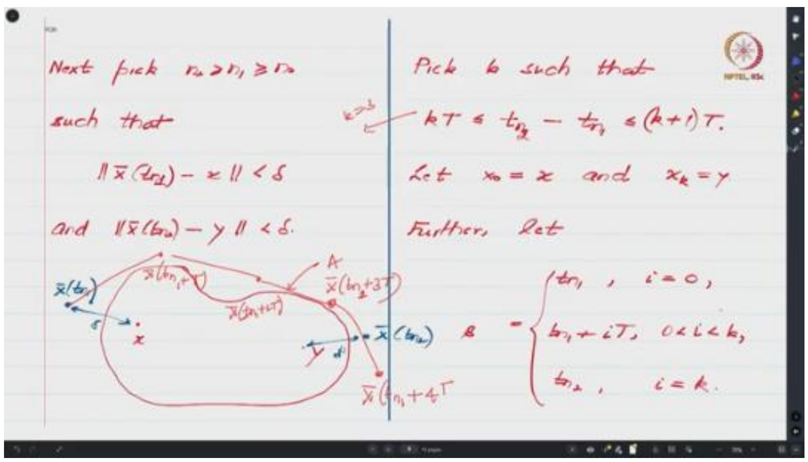

Okay? Right, and now with this in place, we are now ready to pick your intermediate time instances as well as your intermediate points. Okay? So, what we are going to do is we are going to start here and we are going to let this point X be X naught and this point Y be Xk. So, we are going to let this be X naught and let this Y be Xk. And we are now going to choose some intermediate time instances.

So, the first time instance is—so I should actually say SI here. So, for I equals 0, we will let SI be TN1, which is this time over here, and for I between 0 and K, strictly between 0 and K, we will set SI to be TN1 plus IT, right? So, you can see that there is TN1 plus T, there is TN1 plus 2T. So, in this way, we will set SI to be TN1 plus IT.

And in the last case, what we will do is, instead of stopping at TN2 plus 3T, we will actually go to TN2, right? So, for I equals K, we will set you know, the value of SI to be TN2, right?

So, one can then check that your value of SK and SK minus 1, right? could be between 0 and 2T.

On the other hand, the value between, let us say, Si and Si minus 1, where i is strictly less than k, this gap will be exactly T. On the other hand, between Sk and Sk minus 1, the gap could be at most 2T. So, this fact will actually be—I mean, this is the reason why, if you remember, we were working with these windows of length 2T and so on. So now, what we are going to do is, we have identified some intermediate time instances, and now we are going to identify some intermediate points denoted by x1, x2, and all the way up till xk-1. And how we are going to choose these points is, we are going to choose them in a way that xi minus X bar SI is less than delta, and these XIs should be in A. In other words, we are going to pick points in A such that your XI is actually delta away from X bar of SI.

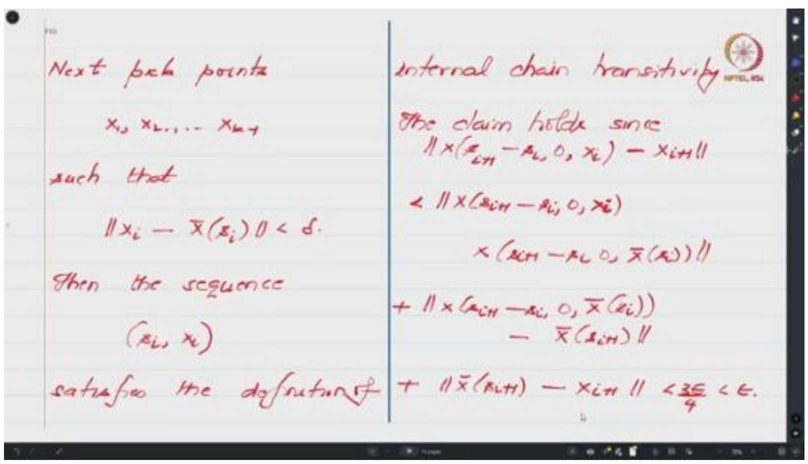

So, if you come back to this picture, we are going to pick a point over here which is delta away from from, you know, this X bar Tn1 plus T. So, we are going to call this as X1. In the same way, we are going to call this as X2, right? And we are not going to pick a delta-close point to this, right? Because this is the kth entry over here, right?

Instead, we are going to directly work with Y, right? So, your X0 will be X. Your, you know, x1 will be some point that is delta away from xtn plus t, your x2 will be delta away from x bar tn plus 2t, and this will be your x3. So, in this way, you have chosen x0, x1, x2, x3. And, you know, using these intermediate points, we will show that if you start from x, okay, we can, in some sense, go to y.

By using solution trajectories of your limiting ODE, right? We are now going to formally show that, okay? So, if you start at XI, Right? And look at the value of the solution trajectory at time instance SI plus 1 minus SI.

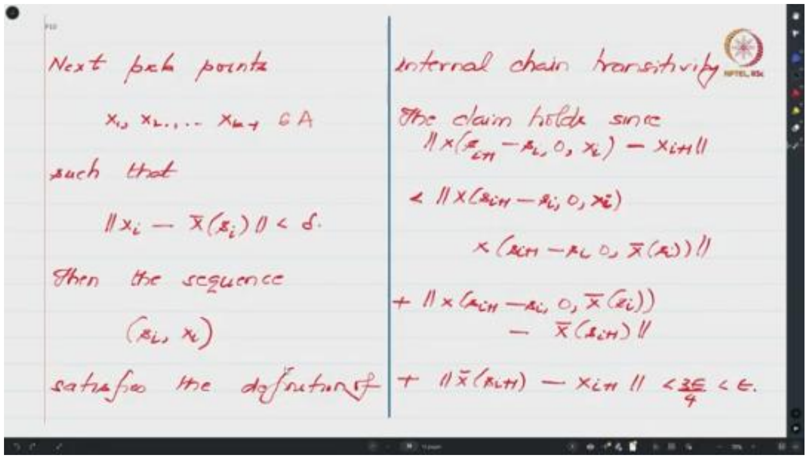

Right? And look at, you know, where your next point is. Right? So, you start the solution trajectory at SI and then ask where this solution trajectory will be compared to SI plus 1. And you can break this distance into three parts.

The first part will be the distance between this solution trajectory and this solution trajectory, where the difference between the two is in the initialization. I should put a negative sign over here. The second term over here is basically this solution trajectory and the value of the linearly interpolated trajectory at SI plus 1, and the last term over here is basically, you know, the distance between X bar of SI plus 1 and XI plus 1. Now, one can see that this term cancels with this, this term cancels with this, and indeed this expression is upper bounded by the sum of these three terms by the triangle inequality and the first term over here. Okay, is less than epsilon over 4 because xi and x bar si are delta away from each other. Why are they delta away from each other? Well, we have chosen xi to be delta away from your x bar si, right? So these two elements or the initializations are delta away, and hence this distance should be upper bounded by epsilon over 4, right? Now, this expression over here

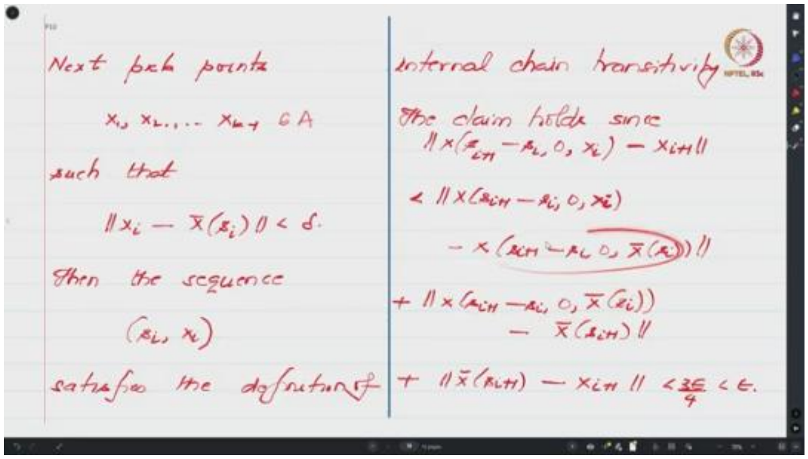

Okay, compares the solution trajectory of your limiting ODE and the linearly interpolated trajectory, both of which start at the same point, that is X bar SI. So, this distance should again be less than epsilon over 4. Right, from the technical lemma, and this distance should be less than epsilon over 4 because, you know, we have chosen our xi plus 1 to be delta away from x bar of si plus 1. Hence, each of these terms should be less than epsilon over 4. So, if you combine all of them, one can conclude that the sum of these three terms will be upper bounded by 3 epsilon over 4, which is strictly less than epsilon.

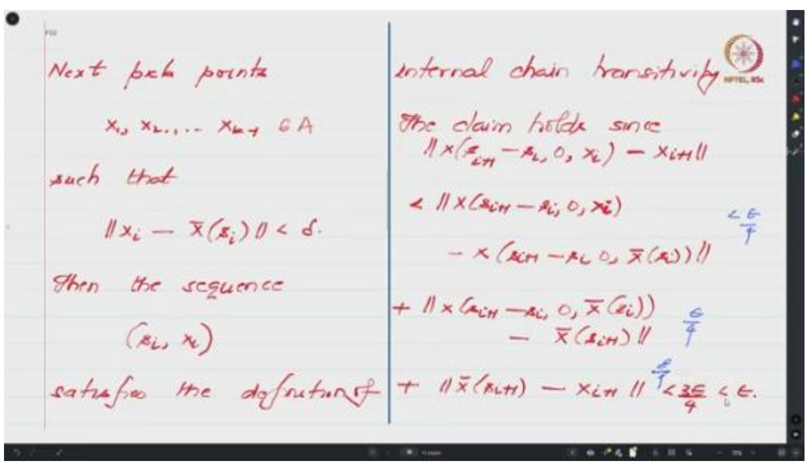

In other words, what we have managed to show is that if you start a solution trajectory from this intermediate point, right? So, if you start a solution trajectory from this intermediate point, then after a time instance t, indeed the distance between x1 and this solution trajectory will be at most epsilon, right? And this is true for, you know, if you start from x0 and go all the way up till x1, or if you start the solution trajectory here, at X1 and look at its value at time instance S1, sorry, S2 and S2 minus S1, and so on and so forth. So, because of this fact, one can conclude that indeed A is internally chain transitive.

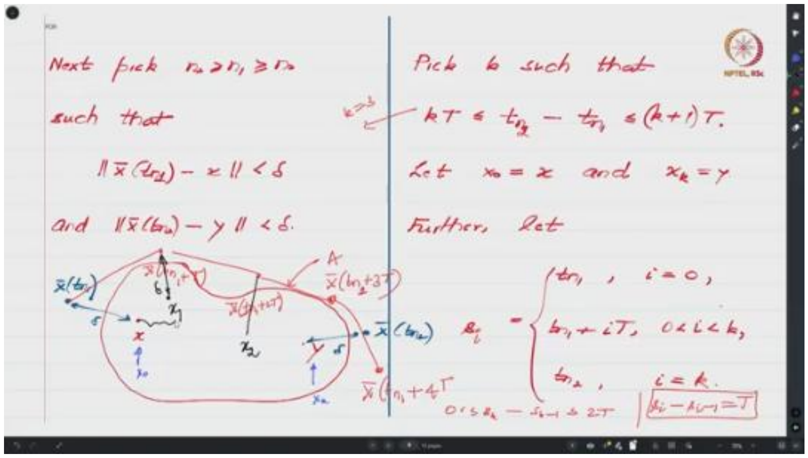

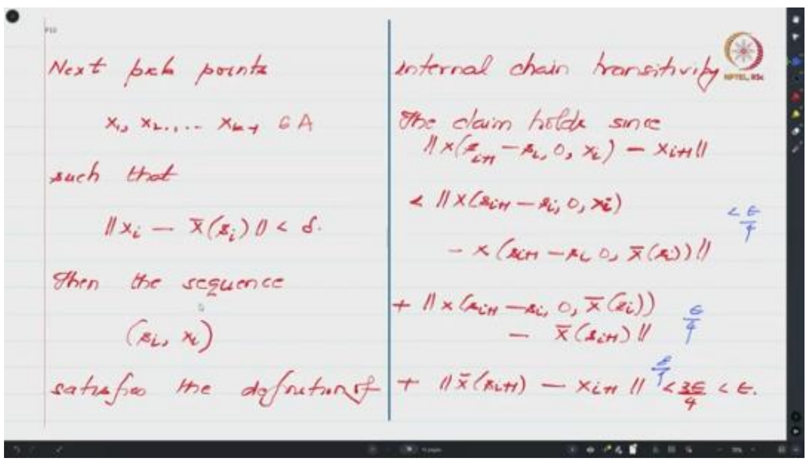

So, this finishes our verification of the proof of the main theorem. Assuming that technical result is true. So, in the next class, we will show how to verify this technical result. Thank you. Hope to see you again next time.