

IntroToLinearDynamicalSystems-Lecture13

Instructor (Stephen Boyd): Let me first answer your – there was a question about one of the homework exercises due today. That’s Homework 5? Right. Okay. And what was the question?

Student: [Inaudible].

Instructor (Stephen Boyd): 10.2? I haven’t memorized them, so you’ll have to tell me what it is.

Student: [Inaudible].

Instructor (Stephen Boyd): What did we ask [inaudible] – I’d have to see the context.

Student: [Inaudible].

Instructor (Stephen Boyd): Oh. It’s your choice, but what’s it for. Is it a simple one?

Student: [Inaudible].

Instructor (Stephen Boyd): Yeah. Okay. Do you know – I believe that to sketch a – not sketch, I guess to get a – there’s a command. This could be completely wrong. I believe it’s “quiver.” Is that right? So some people are shaking their heads affirming that this function in MATLAB will actually plot – I guess quiver is for a group of arrows or something. Don’t ask me. Apparently that’s what it does. That would sketch it, but you’re also absolutely welcome to sketch it by hand, just to draw – take a couple of points, see where they go. That’s the important part. Okay. Any questions about last time? If not, we’ll continue. We’re looking at the idea of left and right eigenvectors. In our next example, very important, you go down [inaudible] pat is a Markov chain. So in a Markov chain, the way we do it, the state is a vector. It’s a column vector, and it actually is a probability distribution. It’s the probability that you’re in each of N states. So PFT is a column vector. Its entries are all positive or not negative, and they add up to one, and they represent a probability distribution. That probability distribution evolves in time by being multiplied by the state transition makers, capital P . I warned you last time that if you see this material in – I guess, by the way – is that coming out. That is straight there. My monitor has this kinda twisted at a five-degree angle or something like that. It’s okay. So I warned you last time that if you take a course in statistics or in other areas that this is the transpose of what you’ll see there. Their matrix P will be the transpose, and they actually don’t propagate column vectors. They propagate row vectors. Yes? Yes, what?

Student: [Inaudible].

Instructor (Stephen Boyd): It is tilted. Ah ha. Well, then. I wonder if you guys can rotate your overhead camera a little bit to make this straight. There we go. Thank you. Okay. So these matrices have the sum over columns equal to one, so that’s what this is. Every

column sum is one. And a matrix like that is called stochastic. Now you can rewrite that this way. If I put a vector in front of all ones – that’s a row vector multiplied by P – I get this row vector here. This is the matrix way of saying that the column sums of P are all one. So this also if you look at it – if you like, I could put a λ in there and say λ is one. This basically says that P – that the vector of all ones is a left eigenvector of P , associated with eigenvalue λ equals one. It tells you in particular P has an eigenvalue of one. But if it has eigenvalue of one, it also has a right eigenvector associated with λ equals one. And that means there some nonzero vector with PV equals V . That’s what it means to have a right eigenvector with eigenvalue one, so PV is V . Some people would say that V is invariant under P . It’s invariant when multiplied by P . Now it turns out that because the entries in P are positive, this is not something we’re gonna look at now. This eigenvector here can actually always be chosen so that its entries are nonnegative, and that means we can normalize it so that the sum of them is one. Now the interpretation then is if you have PV equals V , the interpretation is just beautiful. It’s an invariant distribution. It basically says if you’re in the distribution given by V , that probability distribution, and you update one step in time, you will remain in the same distribution, so it’s an invariant distribution. By the way, that obviously does not mean the state is invariant or anything like that. The state is jumping around. This describes the propagation of the probabilities of where the state is. The state is stochastically invariant. That’s what PV equals V means. Okay. Now we’ll see soon in fact that in most cases, no matter how you start off the Markov chain – and this depends on some properties of P – you will in fact converge this equilibrium distribution. So that’s something we’ll see actually very soon. Okay. So let’s look at this idea of diagonalization. Let’s suppose that you can actually choose a linearly independent set of eigenvectors for N by N matrix A . So we’re gonna call those V_1 through V_N . By the way, we’re soon gonna see this is not always possible. It is not always possible to choose an independent set of eigenvectors of A , but assuming right now it is possible, you write this as AV_i equals $\lambda_i V_i$, and will express this set of matrix vector multiplies as – we’ll just concatenate everything and make it in matrix language. And it says basically AV – A times a matrix formed – that’s a square matrix formed by just concatenating the columns – is equal to V and to multiply each of these V s, that is multiplication on the right by a diagonal matrix, so we have this equation. So you can write this this way. AT equals $T\lambda$, so there you go: five ASCII characters which expresses the fact that the columns of T are eigenvectors associated with the eigenvalues of λI . So a very, very snappy notation to write this. So actually, now I can explain this business. I said mysteriously earlier that you can write an eigenvector this way. And in fact, that’s the way most people write it, but in fact, if you wanna sort of follow this, and make this work for T being a single column or something like that, you can actually write it as this, and that’s – if you see somebody write that, it’s for one of two reasons. Either they’re just kind of like weird, or being perverse by writing the scalar on the right of the vector not the left. Also of course, this requires – to interpret, this requires loose parsing. Or they’re actually quite sophisticated, and they’re simply writing down the scalar version of this eigenvector eigenvector. I should call this the eigenvectors equation. That’s AT equals $T\lambda$. Okay. Now T , because these are independent – T , which we haven’t used yet – this matrix T is invertible because it’s a bunch of – it’s N independent vectors. That’s nonsingular. It can be inverted. I multiply this equation on the left by T inverse, and I get T inverse AT is

λ . Okay? So that's the big – This you've seen before. That is a similarity transformation, and that's a diagonal matrix. It's the eigenvalues, and in fact some people would say the matrix A is diagonalized by a similarity transformation. And in fact, the diagonal matrix is the matrix of the diagonal. Diagonals are the eigenvalues, and the matrix that diagonalizes A is actually the eigenvector matrix, so it looks like that. Okay. Now this is just – it's really in some sense just a change of notation, but it's okay. Suppose you had any matrix T that diagonalized A by similarity. So $T^{-1}AT$ is diagonal. Well, let's see. I'll call these diagonal ones – why don't I just call them λ_1 through λ_N ? And why don't I call the columns of V – of $T^{-1}V_1$ through V_N ? If I do that, and I take this equation, I rewrite this as $AT = T\lambda$. Now I examine these column by column, and column by column that says this. So basically, if you see an equation like $T^{-1}AT = \lambda$, or $AT = T\lambda$ – if you see either of these equations or these like that, it means the same thing. They're just different ways of writing out in different matrix form what this means. Okay? So actually technically, these two make sense. These two make sense even if the V_i are not independent. But for this last one to make sense, obviously you have to have a V_i independent. Okay? So that's the idea. Okay. So we'll say that a matrix is diagonalizable if there is a T for which $T^{-1}AT$ is diagonal, and that's exactly the same as saying A has a set of linearly independent eigenvectors. This is identical. Now sometimes you're gonna hear – I think it's an old word. I hope it's going away. It says that if A is not diagonalizable, it's sometimes called defective. So I don't know where or how that came about, but that's simply what it's called. You still do hear that occasionally. Otherwise, it's called diagonalizable. And it's very important to quickly point out not all matrices are diagonalizable. Here's the simplest example, zero, one, zero, zero, characteristic polynomials S^2 , eigenvalue – only eigenvalue is $\lambda = 0$. There is no other eigenvalue here. Let's try to find now two independent eigenvectors, each with an associated eigenvalue zero. That's the only eigenvalue there is. Well, to say that you're an eigenvector of A with eigenvalue zero, that's just a longwinded way to say you're in the null space. So let's try to find vectors that satisfy zero, one, zero, zero, $V_1, V_2 = 0$ if you like times V_1, V_2 , like that, same thing. Well, the first equation, if you look at the first row, it says $V_2 = 0$. So V_2 has to be zero so V_1 is the only – that's the only thing we have to mess with here. The second row is always satisfied. And there's no way you can pick two vectors of this form that are gonna be independent. It's impossible. Okay? So here's the canonical example of an A which is not diagonalizable. Now we're gonna see actually later today that this issue – it's an interesting one. It's a bit delicate. We'll get to it. It's gonna turn out nondiagonalizability – hey, that came out right. I was about halfway through that and wondering whether it was gonna come out right, but I think it did. I won't attempt it again. Anyway, that property is not – it's something that does – it actually does come up. It actually has real implications, but in many cases matrices that you get are diagonalizable. Not all sadly because that would mean a whole chunk of – a bit of the class we could cut out, which I would not only cut out, but cut out with pleasure. Sadly, can't do that. So there are a couple – we're gonna see the general picture very soon, but for now we can say the following. If a matrix has distinct eigenvalues, so if the matrix – if they're all different, all N eigenvalues are different, A is diagonalizable. Okay? And we'll show this later, but that's just something to know. By the way, the converse is false. You can have repeated eigenvalues but still be diagonalizing. Actually, somebody give

me an example of that. What's the simplest – just give a simple matrix with repeated eigenvalues and yet – identity, there we go. So the identity in \mathbb{R} – what's your favorite number? Seven. Thank you. So that's an \mathbb{R} seven by seven, and very nice choice by the way. I don't know if you chose that, but good taste. \mathbb{R} seven by seven, so it's got – well, it depends if you count multiplicities – well, I can tell you right now the characteristic polynomial is S minus one to the seventh. That is the characteristic polynomial. It has you would say seven eigenvalues at S equals one. So that's about as repeated as you can get. All the eigenvalues are the same. Does I have a set? Can it be – well, I could ask stupid questions like this. Can I be diagonalizable? Can it be diagonalized? And the answer is it already is. So you'd say well what would be T ? You could T equals five. So you could take as eigenvectors E_1 through E_N . By the way, is it true that any set of N eigenvectors of I will diagonalize I or as independent? Let me try – I'll try the logic on you one more time. You can choose a set of seven independent eigenvectors for I . For example, E_1, E_2, E_3 , up to E_7 . Now the other question is this. Suppose you pick seven eigenvectors of I . Are they independent? No. Obviously not, because I can pick $E_1, 2E_1, 3E_1$ and so on, up to $7E_1$. They're all eigenvectors. By no means are they independent, so okay. All right. Now we're gonna connect diagonalization and left eigenvectors, and it's actually a very cool connection. So we're gonna write $T^{-1}AT = \Lambda$.

Before, we had multiplied by T and put it over here, $AT = T\Lambda$. But in this case, we're gonna write it as $T^{-1}A = \Lambda T^{-1}$. Now we're gonna write this out row by row. I'm gonna call the rows of T^{-1} , W_1 transpose through WN transpose. I'm gonna call those the rows. And then in this way, you can right matrix multiply as a batch operation in which you multiply the rows here by all the rows of the first thing by this matrix here. In other words, this is nothing but this. This one is this. And now, if you simply multiply out on the left, the rows of the left hand side are simply this. And the rows of the right hand side are this because I'm multiplying a diagonal matrix by a matrix. If you multiply a diagonal matrix on the left, it means you're actually scaling the rows of the matrix. If you multiply a diagonal matrix on the right, you're actually scaling the columns. You might ask how do I remember that? Well, I remember that when I'm teaching 263 mostly, but then I forget. And then what I usually do is if I have to express row multiplication, I will secretly – because this is not the kind of thing you wanna do in public, let anyone see you doing. I secretly sneak off to the side and I write down the two by two example to see if row multiplication is on the left or the right. And then it's on the left, I'll find out, and I'll say I knew that, and that's how I do it sometimes. Just letting you know that's how I know. But at least for the purpose of – while we are doing 263, I will likely remember. All right, so this says this. If you look at that equation, that's very simple. This says nothing more than the row – than these W_i are left eigenvectors, period. Now in this case, the rows are independent. Why? Because T^{-1} is invertible. Its inverse is T , so the rows are independent, and this means that you have an independent set of left eigenvectors. And they're chosen – in this case, their scaling is very important. Actually, let me just review what this says. It says that if you take a linearly independent set of eigenvectors for a matrix and concatenate them column by column into an N by N matrix, you invert that matrix and you look at the rows of the result. Those are left eigenvectors, so that's what that tells you. Those are left eigenvectors. But it's not as simple as just saying that it's any set of left eigenvectors. They're actually normalized in a very specific way. They're normalized so that W_i

transpose VJ is ΔIJ . Right? Which is – basically this is saying $T^{-1}T$ equals I . Okay? So what you would say is in this case, if you choose the left – if you scale the left eigenvectors this way, they're dual bases. Now I should make a warning here, and that is that if a vector is an eigenvector, so is three times it, and so is seven times it. In fact, any scalar multiple, so when you actually ask for something like left eigenvectors or something like that in some numerical computation, they have to be normalized somehow. Generally speaking, I believe this is the case, most codes – which by the way all go back to something called LAPACK – will return an eigenvector normalized in norm. Notice it'll simply be – have Norm 1. If you just walk up to someone in the street and say “Please normalize my eigenvector,” or “I'm thinking of a direction. Where's the direction?” people will just normalize it by the two norm. There's other cases where they don't. I wanna point out that that will not produce this result. These are not normalized here, so be very careful. These are not normalized. Okay. We'll take a look and see how this comes out. This'll come up soon. Okay. Now we can talk about the idea of modal form. So let's suppose A is diagonalizable by T . From now on, you are to simply recognize this statement as saying this is exactly the same as suppose A has an independent set of eigenvectors V_1 through V_N . If you shove them together into a matrix, I'm gonna call that T . That's what this means. Well, if you take new coordinates to be X equals $T\tilde{X}$, what this means is you are – \tilde{X} are the coordinates of X in the T expansion. Or if the columns of T are the V s, it's in the V s. So you would actually say \tilde{X} gives the coordinates of X in what people would call in this case the modal expansion, or the eigenvector expansion. In other words, instead of writing X , which is in the standard bases or the coefficients, \tilde{X} are the mixture of V_1 through V_N that you need to reconstruct – to construct X . That's what \tilde{X} is. Well, this thing is \dot{X} , and that's $A\tilde{X}$, \dot{X} is $T\dot{\tilde{X}}$, so $\dot{\tilde{X}}$ is $T^{-1}A\tilde{X}$. That's diagonal, and you get this. So in that new coordinate system, you get $\dot{\tilde{X}}$ equals $\Lambda\tilde{X}$. And that means that by this change of coordinate that the autonomous linear system, \dot{X} equals AX is decoupled. So in \dot{X} equals AX , the way you should imagine that is a bank of integrators with a matrix A in a feedback loop. But a matrix A basically takes N inputs and produces N outputs, but it's got all these horrible cross gains. I mean if A has all entries non-zero, it means every output depends on every input, and so if you're mixing the dynamics [inaudible] basically what's happening. When you diagonalize like this, you have completely decoupled all the dynamics and it looks like this. Okay? And that says that at least in this coordinate system, it's very simple. The trajectories give you N independent modes, and the modes are just simply – well, obviously they have nothing to do with each other. They're totally decoupled. And this is called modal form. That's a very common form. Now these can become complex, in which case it's a bit weird, and you have to explain a little bit, and make a story about how the real and the imaginary parts are separately solutions, and all that kind of stuff. Another very common form you'll see is real modal form, and you'll see this for example in mechanical engineering a lot as real modal form for example for a structure. That's how they would describe the dynamics of a structure by giving real modal form. Now in this case, you can – there's actually a way to construct a real matrix S so that $S^{-1}AS$ is not diagonal, but it's diagonal with one by one or two by two blocks like this. Okay? So every time in $T^{-1}AT$, every entry in this which is complex – that means for every complex eigenvalue, you'll actually collect that and its conjugate, and then actually you can take

the real and imaginary apart and so on, and you'll actually get a form like that. I'm not sure – I don't remember if we've actually – hey, we don't tell you how to construct S . That would make an excellent homework exercise. Yeah, okay. So that's the [inaudible], so find S . It's a good exercise to see if any of this makes sense, mess with matrices and things like that. Okay. So you get all these little blocks – by the way, these little blocks like this, you should by now start recognizing. So a little block that looks like σN , σN , ωN , ωN , the characteristic polynomial of this is S minus σN – sorry – squared, plus ωN squared like that. Now assuming here that ω is less than σ – I believe that's the condition here. Oh sorry. Absolute – no, sorry – let me keep this out of here. The roots of this thing are gonna be minus – no, they're gonna be σN plus minus square root of – let's see – I guess it's ωN squared minus σN squared, something like that. Okay? So that's what it's gonna look like. Now these things are the complex eigenvalues, so that is actually negative. Otherwise, this would be two real parts and can be split. And it should kinda make sense because this is the self-exciting component between X – if this were a two-dimensional system with X_1 and X_2 , and these are the cross components which actually give you the rotation. So this condition basically says you get enough rotation. Otherwise, it splits into two. I can tell by facial – just a quick gross look at facial expressions, I've now confused almost everyone, except [inaudible]. Yes?

Student:[Inaudible].

Instructor (Stephen Boyd):How did I look at it and immediately say it? Well, it wasn't totally immediate, but let's – I – two by two [inaudible] two by two matrices, that's where there are basically no excuses, right? So for two by two matrices you should be able to work out one of the eigenvalues with the inverse and things like that. Above that, no one could possibly hold you responsible for it. But two by two, let's just do it. It's $\det(SI - ABCD)$. I mean these are the kind of thing I guess you should know, like that. And so I did this in my head, but I sort of knew what the answer was, so you get this times – minus BC . And now if I like, I could write my quadratic formula which I don't dare do. That would be some horrible expression. This was easier because after all, B was equal to minus C , and A was equal to D , so fewer things to keep track of, but this is what I did in principle. Now one of these two by two blocks, which is a complex mode, they look like this, and they're really quite pretty. They are essentially cross-coupled. It's a pair of integrators cross-coupled. And by the way, if σ is zero, you get the pure oscillation, and this is something we've seen before. In fact, that matrix is zero ω minus ω zero. That's when σ 's zero, so you get this. You get that. And this you should recognize by now as a rotation matrix, I mean maybe – sorry, this is not a rotation matrix. Well, it is, but in this case – this you should recognize as a harmonic oscillator. By the way, you don't have to recognize this as a harmonic oscillator, but we talked about it, and the little bit of playing around with this would convince you that it is. Yes?

Student:[Inaudible].

Instructor (Stephen Boyd):Yeah. Sorry. What's that? Oh, this is not it. Thank you. This is this? Like that? Sorry? With a square.

Student:[Inaudible].

Instructor (Stephen Boyd): Really? Good. Great. Fine. Sorry. Actually, you're right. Of course, this is – fine, so it's $\sigma_N \pm j\omega_N$. There we go. It's that? Thank you for fixing it. How come no one else caught that? Hm. Well, that's one more demerit for this class. That's bad. All right. So diagonalization, it simplifies a lot of matrix expressions. So diagonalization, I'll say a few things about it. It's mostly a conceptual tool. There are a few places where it actually has some actual use. It's widely used for example in mechanical engineering. There in fact is a very famous – there's very famous codes that will take a description of a mechanical thing, and then spit out the modal form, so it'll list eigenvalues, and it'll actually give you the eigenvectors, which in fact are these modal ones, not the complex ones, the two by two blocks. But in fact, mostly it's a conceptual tool, and let's see why. It's gonna simplify a lot of things. So if you have $(sI - A)^{-1}$, that's the resolvent. Generally, if A is two by two, no problem, you can work out what this is. If A is three by three, this is already not a picture – it's nothing you – you don't wanna see the expression for $(sI - A)^{-1}$. Trust me. You can compute it easily for some particular A , but that's another story. However, we can do the following. Wherever you see A , you replace it with $T\lambda T^{-1}$ like that. So we do that here, and now I'll show you a trick. This identity matrix you write as $T T^{-1}$. Okay? Now I pull a T out on the left, and a T^{-1} out on the right in this inner matrix. I invert a triple product that's the same as the inverse the other way around, so I get $T (sI - \lambda)^{-1} T^{-1}$. By the way, inverting a diagonal matrix, that's fine. That's easy to do. You invert the diagonals, and you get this. That's the resolvent. Okay? By the way, this is sometimes called the spectral resolution of the identity or some – there's some name for it. There's a name for this way to represent the resolvent. Actually, let me say a little bit about that. Some of you might know about the idea of residues in complex analysis. Then again, maybe none of you know about residues or partial fraction expansions. Partial fraction expansions? Come on. Somebody's heard of that. Did you learn about that in some stupid signals and systems course? Is that – yes. Is that where you heard about it? Okay, great. So this says the partial fraction expansion of the resolvent is this. It's really quite cool. Let me try to get it right. Oh, I think I can get it right. It's this, and I'm using informal syntax here. So that's the partial fraction expansion of this. Partial fraction expansion of a rational function is to write it out as a sum of terms each of which is one over one minus – s minus a pole, so that's the partial fraction expansion. Okay? The other way to do is say that these Rank 1 matrices, $V_i W_i^T$ are the residues of this function at – that's the residue of this function at the pole λ_i . Okay. So the diagonalization simplifies tremendously this, the resolvent. It's also true for the powers. For example, if you raise a matrix to power, if you know it's eigenvectors and eigenvalues, it's very straightforward because you simply write $T\lambda T^{-1}$ to the K . Now when you do this, you just line these K times, and the T^{-1} here annihilates the T to the left of it. And this happens in all of these, and you get $T\lambda^K T^{-1}$. Okay? So that means it's very straightforward in fact to calculate powers of a matrix. And in fact, this already is a method perhaps not – maybe not a good one, but at least conceptually this gives you a very good way for example of simulating a dynamical system. Or if someone walks up to you on the street, and asks what's A to the one million, probably the worst thing you could do would be to

write a loop that keeps multiplying an N by N matrix by A and let it run a million times. That would be one way to do it. And the cost of that would be ten to the six times N cubed if you did that loop. Okay? Now you can also do the following. You could also calculate the eigenvectors and eigenvalues, and although I'm not gonna get into the details, that could also be done in order N cubed, like five N cubed or something like it. This doesn't matter, but I just wanna make a point here. Once you calculate this and this, which costs N cubed, let's say five, this is nothing but a bunch of calls. It costs basically zero – it costs N , which is basically dominated by the N cubed. So this can be done in five N cubed. Okay? And the point is that is a whole lot smaller than that. So diagonalization in this case actually gives you a very serious advantage in how to compute something. Okay. Let's look at exponential. You wanna look at E to the A in general. Calculating E to the A is a pain. It's not fun, unless A is two by two or has some other special structure like diagonal or something like that. It's actually quite difficult. Let's write out the power series, and if you write out the power series here, not surprisingly when you make a – we've already seen when you have a power, it's the same as simply putting the T and the T inverse here, and having these things – these are all diagonal matrices, and that's nothing but this. Exponential of the diagonal matrix is e^{λ} and diag commute for a matrix exponential. You get this like that, and that gives you a very simple way to compute the exponential. That's not quite how it's done, but it's one of the methods that can be used is this. Okay. Now in fact, this idea extends to analytic function, so if you have a function which is given by a power series, that's an analytic function – you don't need to know this, but it's just kind of cool, so it doesn't hurt. So if you have an analytic function that's a function given by a power series – it could be a rational function. It could be exponential, anything else – then it's absolutely standard to overload an analytic function to be called on N by N matrices. So F of A is given by $\beta_0 I$ where the power series expansion of F is this. Okay? You've seen one specific example. You've seen so far the exponential. This allows you to work out this thing. So here for example, we would have the following. We would have F of A is T times – well, let me write it this way: diag of F of λ I times T inverse. There's your formula like that. Okay? So it gives you a very quick way. This is actually something good to know about because there are a lot of times when you do see things that are polynomials of matrices, rational functions of matrices, and things like that. Those do come up a lot, and it's good to know that if you diagonalize, they will simplify. They'll simplify actually just to this single analytic function applied to the eigenvalues. Okay? Actually, I'm not quite sure why this part didn't get into the notes. That's very strange. Okay. Let's look at the solution of X dot equals AX . Well, we already know what it is. It's simply this. It's X of T is the X of TA times X of zero. That's what it is. That's the solution. And we already have a rough idea of what this looks like and things like that, or we have some idea of it anyway. Although, this is not the kind of thing a person can look at right off the bat – just take a look at a four by four matrix or for that matter forty by forty matrix – look at it and say, “Oh, yeah. That's gonna be a rough ride in that vehicle,” or something like that, or “Yeah, that's gonna be some business cycles. I can see them.” I mean there's just no way anybody can do that, unless A is diagonal or something like that. Okay. Well, let's start with X dot equals AX . T inverse AT is λ . X of T is E to the TA X of zero. That's the solution, but this thing is this, and then this has really the most beautiful interpretation because T inverse X of zero, I write it this way: W I transpose X of zero. That's a number.

Then it's multiplied by this thing, which actually tells you for that eigenvalue eigenvector whether it's – this tells you whether it grows, shrinks, if it's complex, oscillates, and all that kind of stuff. And then this gives you the V_i , reconstructs it. So the interpretation of this formula is really quite beautiful, and every single can be interpreted. It's this. What happens is you take the initial state $X(0)$, you multiply by W^T , and you get something very specific. That is the component of $X(0)$ in the V_i expansion. So for example, if $W_3^T X(0) = 0$, that has a meaning. It says if you expand $X(0)$ in a V expansion, you will not have any V_3 . That's what this says. Okay? Because that's what this does. It decomposes it this way. In fact, let me write that down right now. $X(0) = \sum W_i^T X(0) V_i$. That's the expansion. And that's true for any $X(0)$. Okay? This looks fancy. This comes from the fact that – I mean this is actually quite straightforward. This basically is a restatement of this. There's nothing [inaudible]. Okay? So W_i^T , which are the left eigenvectors, they decompose a state into the modal components, if you wanna call the V_1 and V_N the modal components. That's what this does. All right, that's fine, so that you decompose it. Then this thing, time propagates that mode. That's what $E = e^{\lambda T}$ does. It propagates the i th mode in time, very simple formula. Why? That's the point of a mode. A mode propagates in a very simple way. It grows. It shrinks. It oscillates. That's about it so far, by the way. And then it comes out along a direction V_i , so all the parts of this are supposed to make sense, so that's the interpretation of this. Okay, but now we can actually ask some really interesting questions and answer them. You might ask this. You have $\dot{X} = AX$. Now remember, we say by definition A – the system is stable. By the way, you would normally say the system $\dot{X} = AX$ is stable. Sometimes, as a matter of slang, you'll hear people talk about A being stable, but that's – it should be understood that's slang. So this system is stable provided all solutions of $\dot{X} = AX$ converge to zero, they all decay. But now we're asking this. For a general A , for what initial conditions do you have $X(T) \rightarrow 0$ as $T \rightarrow \infty$? By the way, this one answer, you can always give, no matter what A is. What's that answer? Zero. If $X(0) = 0$, then it stays zero, period. And therefore, it goes to zero. So the initial state zero, no matter what A is, gives you at least one trajectory that converges to zero. I mean converge is a little bit technical there. It is always zero, but that means it converges to zero. Okay. Now the way to answer this is you divided the eigenvalues into those with negative real part, so let's say that's the first S of them, and the others, so these have nonnegative real part. Now we can answer the question lots of different ways. One is this: you say now that's just a formula for $X(T)$. This thing will go to zero, provided the following holds. The first S terms in here shrink. The remaining $S+1$ through N do not shrink. Therefore, this will go to zero provided these numbers are zero for $S+1$, $S+2$, and so on. That's one way to say. So one way to say it is that these numbers should be zero for $S+1$, $S+2$, up to $S=N$. By the way, that is identical to saying that you're in the span of V_1 through V_S . Why? Because $X(0) = \sum W_i^T X(0) V_i$ like that. You have this. Therefore, to say that these are zero from $S+1$ to N means that $X(0)$ is a linear combination of V_1 through V_S . Okay? So these are two different ways to say this. And there'd be all sorts of names people would call this. They would say that – they would refer by the way to this span. They would call that the stable eigenspace or something like that. That would be one – or some people would just call it the stable subspace, and the idea would be this. If you start

in this subspace, the trajectory will go to zero. It might oscillate. It'll go to zero. If you're not in this subspace, you will not. So that's how that works. Okay, so that's the kind of question you can answer now. And finally we'll handle this issue of stability of discrete time systems. So suppose the matrix is diagonalizable, and you have the linear dynamic [inaudible] XT plus one is AX of T , then the solution is trivial. It's just powers of A . But if you write A as $T \lambda T^{-1}$, then A to the K is this. Now I understand – I know how powers of complex – I know what powers of complex numbers do. That I can actually handle, and so you get this. Powers of complex numbers go to zero only if their absolute value is less than one. Their imaginary part tells you about how much of a rotation in degrees you get at each step, but their magnitude tells you how the magnitude scales, and you realize that XT plus one is AX of T is stable if and only if the eigenvalues are less than one in absolute value. Okay? And it turns out this is gonna be true even when A is not diagonalizable, so I don't mind stating it as a fact right now. XT plus one is AX of T is stable if and only if all the eigenvalues have magnitudes less than one, so that's the condition. Actually, as in the continuous time case, there's a much more subtle statement. The spectral radius of a matrix is defined to be the maximum of the absolute values of the eigenvalues. Okay? And so one way – this is called a spectral radius. It's denoted ρ . This is relatively standard notation. What this says is X – the discrete time autonomous system XT plus one is AX of T is stable if and only if the spectral radius of the dynamics matrix A , or update matrix, or whatever you wanna call it – the spectral radius is less than one. That's the condition here. Now more generally, ρ of A gives you the growth or decay magnitude, asymptotic. So for example, if ρ is 1.05 – in other words, there is at least one eigenvalue with a magnitude of 1.05, it says that X of T will grow asymptotically. It depends on the initial condition, but it can grow as fast as 1.05 to the T . If the spectral radius is 0.7, it says that after ten steps roughly, the state has decayed roughly by 0.7 to the ten. That's a small number. Okay? So this is the continuous time analog of the maximum of the real part of the eigenvalues of a matrix. That's what this gives you. Okay, so enough on all that. We'll now do a bit on – maybe I'll even cover it, the Jordan canonical form. So here I actually have to ask, how many people have seen – or perhaps right verb is been subjected to the Jordan canonical form? So a handful of people. Did it make any sense at the time? How should we interpret this?

Student:[Inaudible].

Instructor (Stephen Boyd):It did? Okay. Good. Let's look at it. So the Jordan canonical form is essentially – it's as close as you can get to a diagonalization when you can't diagonalize the matrix. So let me explain that. It's this. Any N by N matrix – any, no exceptions – can be put in something called Jordan canonical form by a similarity transformation.

So there's a matrix T , obviously invertible because I'm about to refer to T^{-1} , for which $T^{-1}AT$ is J . J – that's the Jordan form – is a block matrix. Each of these blocks, which is called a Jordan block, looks like this. It's a bunch of λ s with ones on the superdiagonal. So that's a so-called Jordan block. By the way, a one by one Jordan block is a little matrix that looks like this. Okay? So a diagonal matrix is a special case of a Jordan form. It's the special case when there are N Jordan blocks and each block is one

by one. So basically, you have these little ones in the superdiagonal. We'll see what that means soon, what the ones mean. So that's a Jordan block. Now a couple of comments about Jordan blocks. First of all, the Jordan form is upper bidiagonal. That's a name meaning it's got a diagonal, and it's got – one diagonal above it is nonzero. It's upper – it's much more than that because in fact there's ones in the upper – it's in the zero one, the upper diagonal, and not only that, it can only be one if the lambdas are repeated there. So diagonal is a special case of N Jordan blocks. And it's gonna turn out the Jordan form is unique. Now you have to interpret that very carefully. It is not of course on the details of the mathematics of linear algebra, so it's not like we're gonna get crazy with all this, but it's important to understand what it means to say it's unique. It says that basically if two people calculate a Jordan form for a matrix, they actually can be different. One difference is simply this. They might order the blocks in a different way. However, the following is true. If two people work out a Jordan form, they have different T s here possibly, then there's a permutation – a block permutation which will change one Jordan form into the other. So the way you would say this is you would say the Jordan form is unique up to permutations of the blocks. So the things people can – the types of things people cannot agree on is what is J_1 . No one can agree on that because it depends on what you chose to put as the first Jordan block. However, you can't – no one can possibly disagree about the numbers of Jordan blocks for example, and the sizes of them, and the sizes associated with a certain eigenvalue. So for example, if you say this has three eigenvalues, this eigenvalue is associated with one Jordan block of size two by two, everyone computing a Jordan decomposition will actually – will agree with that. Okay. Now I should also mention Jordan canonical form is – it's an interesting thing. It is almost strictly a conceptual tool. So it's used to show things, to illuminate ideas, and things like that. It is actually not used in almost any numerical computations. Okay? So if you go to the web or something like that – if you go to Google and type let's say – if you type something like “source code Jordan canonical form,” you will get – actually, what you'll mostly get is you'll get a bunch of harangues about how terrible it is, and no one should ever compute the Jordan canonical form by – numerically, and so on and so forth. That's probably what you'll get, but you'll find some strange things there, but it's basically not done. Even when you do find algorithms for it, every paper will start this way. It will say, “It's well-known that you basically – it doesn't make any sense numerically to compute the Jordan form.” It goes on, and it says, “But let's suppose you did. You really had to. Then this paper's about how you might do it, if you were to do it, but we don't recommend it.” So that would be the kind of abstract you'd find. Not that this matters. I'm just mentioning it. Okay. Maybe it's not never, but it's awfully close, and boy do you have to justify yourself if you actually do anything like this in any numerical application. All right. Now the characteristic polynomial of A is – of course, if J is block diagonal – so the characteristic polynomial of – actually, under a similarity transformation is the same. Wasn't that a homework problem? It wasn't? That's terrible. Well, similarity – wait a minute. Oh well. That's – maybe it shouldn't have to be. Was it one? Well, I'm glad to see though that everyone thought about that problem a long time and really – in fact, that's great because it's actually below even your consciousness now. It's so ingrained –

Student: I think it's the current homework.

Instructor (Stephen Boyd): It's what? Oh, the current homework. Oh well, that would explain it because that would've been terrible. I mean it's a quick calculation, but the characteristic polynomial under a similarity transformation, it doesn't change the eigenvalues. So the eigenvalues of A are the eigenvalues of this thing. That's a block matrix. Eigenvalues of a block matrix are the eigenvalues of all the blocks stuck together. Eigenvalues of that matrix, that's upper triangular. Eigenvalues of this matrix are λI with a multiplicity N . The characteristic polynomial in fact of this is S minus λI to the N here. That's just $S - \lambda I$. Okay. So basically, this tells you the following. If you wanna get the characteristic polynomial of the matrix, you take $\det(S - \lambda I)$ – it's the eigenvalues associated with the blocks raised to the block size. And now we immediately see the following. Once you believe in the Jordan canonical form, which I will not show how – I will not go through the week long proof that any matrix has a Jordan canonical form, especially because the computational algorithmic payoff is – to say dubious is putting it very nicely, so I won't go through that, but assuming you believe it, and you should – that is after all what we have mathematicians for, and they assure us that it's true. Then it says immediately that if a matrix is diagonalizable, its Jordan form must be – you can only have block sizes one by one. To say that a Jordan form has block sizes one by one says it's diagonalizable. That's basically what it says. Okay. Now this – when you see repeated eigenvalues now – so in fact, let me explain how this works. If you see repeated eigenvalues, it means maybe you have a nontrivial Jordan form. Oh, I should mention something here. If the Jordan blocks are all one by one, this is diagonal. People would call that a – if any block is two by two or bigger, people call that a nontrivial Jordan form, meaning diagonal is just diagonalizable. So if you see – what this says is the following. If the eigenvalues are distinct, your matrix is diagonalizable. And if someone says, "Yeah? What's the Jordan form?" You'd say, "I just said it's diagonalizable." Okay. Jordan form is N Jordan blocks, each one by one. That's the trivial Jordan form. If you see repeated eigenvalues, it does not guarantee that the Jordan – you're gonna have nontrivial Jordan form. In fact, somebody quickly give me an example of a matrix that has repeated eigenvalues and yet has a trivial Jordan form.

Student: [Inaudible].

Instructor (Stephen Boyd): I. What size? This is very important. We talked about this earlier. Seven. Thank you. So seven by seven – the seven by seven identity matrix has seven absolutely equal eigenvalues. Its Jordan form is trivial, which is a pedantic way of saying it's diagonalizable, and so – on the other hand, you can have more – in fact, let's talk about seven by seven. That's kinda big. Let's talk about like four by four. So here, that's the identity. Eigenvalues all are one. It's got four Jordan blocks of size one by one. Okay? How about this one? Eigenvalues are all the same. The eigenvalues of that matrix are one, one, one, one, period. What's the Jordan block structure? Well, there's one block here, and it looks like that. So you can have a block of two, and then two separate blocks of one. Are there others for this one? Well, I could do this. I could have a block of – let's see if I can get it right. No. Yes. If I'd given myself enough room, it would've been right. How about that? That matrix – what's the block – what am I doing? My god. Okay, let me just get rid of that. I didn't do that. There we go. Okay, here. Since I can't even get it straight, I'll show you the blocks. Okay? What about this one? What's the Jordan – I

mean the block size here you would describe as it's got two Jordan blocks, one three by three, one one by one. By the way, eigenvalues, this matrix, this matrix identity, all the same. Any others? One more one. So we could have a single Jordan block – I don't know what I'm doing. Here we go. One, one, one – there we go. It's a single block of size four by four. And that would be the – and any others? Let's list all possible Jordan forms of a four by four matrix with four eigenvalues one. There's one we missed. What is it?

Student: Two two by two.

Instructor (Stephen Boyd): Two two by twos, so that's – exactly. So you can have this like that, and that's it. These along with I are the five possible Jordan forms for a matrix with four eigenvalues of one. Okay? Of course, a natural question in your mind would be – well, let me list some. The first might be who cares. So we'll get to that. And the second might be – and it's related – is what's the implications. What would it mean? How would you know? How would this show up for example in dynamics or something like that of a system? How would you know you had one of these, and what would be any consequences of it? Okay. And we'll get to that. I promise. So the connection between the eigenvalues and the Jordan blocks and sizes is a bit complicated, but it all comes from this. It says that basically if you have – the characteristic polynomial is a product of S minus λI to the block size I . And the null space for example of λI minus A is the number of Jordan blocks with eigenvalue λ . And we can check that because what happens is if you look at the null space, if you look at λI minus A – I will multiply by T inverse in T like this, and T inverse in T goes in there and annihilates itself. It's basically – that's λI minus J , and that is equal to a block matrix that looks like this. It's λ minus λ one. And then there's some minus ones on the superdiagonal like that. And I won't draw the other blocks. Okay? Now if you wanna know what's the null space of this matrix, you have columns – at the leading edge of each Jordan block, you have a column whose only nonzero entry – possibly nonzero entry is λ minus λI . So if λ is equal to λI , you get a zero column, and that means that matrix is gonna drop rank. It is not gonna be invertible. So that's what – this happens. So in fact, this will happen. Every match at the beginning of a Jordan block, you will get a zero column, and that says in fact the dimension of the null space of λI minus A is exactly equal to the number of Jordan blocks associated with λI . So over here, let's look at that. What is the – let's look at the null space of λ , which is one, minus the matrix A . And let's look in different – this is one times I minus A . Let's look at the different cases. If you take I , and I ask you what's the null space of one times I minus A , that's the null space of the four by four matrix zero. What's the null space of the four by four matrix zero?

Student: [Inaudible].

Instructor (Stephen Boyd): It's what? It's R^4 . It's all four vectors. So it's four-dimensional in this case. What is the null space of I minus A for this matrix? What is it?

Student: [Inaudible].

Instructor (Stephen Boyd): Well, I wouldn't say \mathbb{R}^1 because that means – that's just the set of real numbers. It's the set – it's one-dimensional, which is I think what you meant. It's one-dimensional, and it's all vectors of the form something zero zero zero – yes, [inaudible] something zero zero zero. It's one-dimensional in this case. Why? There's one Jordan block. It makes perfect sense.

This case, if you take I minus this matrix, these become minus ones and these become zeroes, and then you ask what is the dimension of the null space of that matrix, and the answer is two. That's two Jordan blocks. Same here, and in here it's actually – the dimension of the null space is three. Okay? So basically, the amount by which – the amount of rank that λI drops – λI minus A drops when λ is an eigenvalue tells you something – well, actually it tells you exactly the number of Jordan blocks. That's not enough by the way to give you the full block structure. That comes out of λI minus A raised to various powers. And I'm not gonna go into this, but we can – in fact, I'm not gonna go into that, except that we can actually just – let me just ask a couple of questions. Suppose a matrix has eigenvalues minus one – with multiplicities, minus one, three, and five – five by five matrix. Let's enumerate all possible Jordan forms for that matrix. Let's start. What are the possible Jordan forms? What's the simplest one?

Student: Trivial.

Instructor (Stephen Boyd): The trivial one which is just diagonal, minus one, minus one, minus one, three, five – and if someone asked you how many Jordan blocks, what are the sizes, what would you say here? How would you describe the trivial – it's diagonalizable here.

Student: Five one by one.

Instructor (Stephen Boyd): Yeah, so you'd say it's five one by one. But you'd also say by the way, it's pedantic to talk about Jordan blocks when a matrix is diagonalizable. That should also – that should be the second part of your reply when someone asks you about this. Okay. Are there any others? Any other possible Jordan forms? For example, could I have – could the eigenvalue three correspond to a Jordan block of size two? No. Out of the question because its multiplicity is one. Same for five. So these two – no matter what happens, this matrix has two Jordan blocks of size one by one, one associated with eigenvalue three, one with five, period. And the only one where there's any ambiguity would be this little block of repeated ones, and what are all possible Jordan forms for three repeated eigenvalues of minus one? We've got one that's diagonal. What else? Two in one and what?

Student: [Inaudible]

Instructor (Stephen Boyd): And three. Okay. In this case, if I asked you – if I told you what the dimension of the null space of minus one minus I minus A is – if I told you that number, could you then uniquely determine the Jordan form of A ? I'm getting this out

and ready. What do you think? You could. And the reason is the dimension of the null space of $\text{minus } I \text{ minus } A$ can be either one, two, or three. If it is three, it means A is diagonalizable, end of story. If it is two, it means this – there is one block of size two by two and one of size one by one. If it is one, it says there's a single Jordan block, period. And therefore, you have determined it. Warning! That is not always the case. If I told you that a matrix is four by four, has four eigenvalues of one, and the dimension of the null space of $\text{one } I \text{ minus } A$ is two, that does not tell you the Jordan form of A . Why? Because you don't know if it is this one or this one. Each of these has two Jordan blocks, and you don't know which it is. Okay? So that's the idea. Okay. Let's look at – well naturally, the columns of T in $T^{-1}AT = J$ are called generalized eigenvectors. These – you group these according to the blocks, and these are the columns that you associate to the Jordan blocks. If you call these – if you split these out as columns, then you'll get something like this. The first one comes out just the way you think it does. Or sorry, not the – yeah, the first one comes out just the way you think it does. It's $AV_1 = \lambda V_1$ here. That's the first one. But the next ones because of that upper triangular – sorry, that upper diagonal, you inherit this one. Each AV_i is actually the previous one plus λ times this, and so these are called generalized eigenvectors. You actually won't need to know this. They don't come up that often, but you will see this every now and then. You'll see people refer to generalized eigenvectors. Now for a – if you have $\dot{X} = AX$, if you put a change of coordinates, if you put it in Jordan block form, basically that splits the dynamics into – it splits the dynamics into I guess K separate – independent blocks. Each one is a Jordan block. Now you should have an overwhelming urge to draw a block diagram of a Jordan block system, and this is it. It's a chain of integrators. The chain of integrators by the way corresponds – that's what that upper block of – that upper diagonal of ones is this – gives you a chain. So you should start thinking of an upper block of ones as giving you things like shift. It's a shift or it's a chain in this case like that, and so on. And then the lambdas are simply wrapped around this way. So interestingly, people who do engineering and mathematicians both refer to sometimes Jordan blocks as Jordan chains for totally different reasons. People in engineering refer to it as a chain because it's got this chain of – it's dynamics built around a chain of integrators. And in math, it's a Jordan chain because it's a chain of subspaces. So this only shows why if you're in engineering, so that's the dynamics you see. By the way, when you see this, if you remember things like – so actually, let me actually explain a little bit now because the main thing is to get a rough idea of what on earth does it mean if you have $\dot{X} = AX$ and A has a Jordan block. This says that some of the dynamics is sort of connected in – would you call that series, or cascade, or something like that? That's what it means. It means that some of the dynamics feeds into the other. Remember what the diagonal system looked like. It's N boxes that look like this. So the Jordan blocks are actually gonna – it's gonna turn out it's gonna have to do with dynamics blocks that cannot be decoupled. That's what it's gonna be. It's gonna be dynamics blocks that cannot be decoupled because they're connected in cascade, not in parallel. Okay. And we can look at things like the resolvent and the exponential of a Jordan block. If you look at the resolvent, you see you have this upper thing here, but if you take the inverse of that, the inverse of an upper triangular matrix is not that bad to work out, and it looks like this. Actually, it's quite interesting because now you can see something else. You can see that when you take the resolvent of a Jordan block, you're

gonna get powers – you’re gonna get S minus lambdas to negative higher powers. Didn’t have that before in the resolvent. So it turns out it’s gonna correspond to – repeated pols in the resolvent are gonna correspond to Jordan blocks. Could work out the Laplace transform, and this will actually at least give you part of the idea of what the meaning of these things is. When you work out the exponential of a Jordan block, it turns out sure enough you get this E to the T lambda part. We’re hardly surprised to see that, but now you can see what a Jordan block does. It gives you polynomials. So I think what I’ll do is – let me say a little bit here. This tells you what you needed to know. When you see X dot equals AX , and let’s make it simple – let’s say all the eigenvalues are lambda, period. Okay? Now I can tell you what it means for this – what the Jordan blocks in A – if A is diagonalizable, the only thing you will see in the solution will be things that look like that, period. If there’s a Jordan block of size two by two, you will not only see exponentials, but you will see terms of this form like that. That’s only if there’s a Jordan block of size two by two or larger. If there’s a Jordan block of size three by three, you will see not only TE to the lambda T , but T squared E to the lambda T . Another way – you can turn it around and say that if you see a solution that looks like that here, that says that there is a Jordan block of size K plus one there. Did I say that right? Yes, K plus one. That’s what it says. So Jordan blocks are basically gonna be the matrix attribute which you’re gonna associate with T to the K – these terms which are a polynomial times an exponential. Okay? And let’s actually just look at one example just for fun, and then we’re gonna quit. Let’s look at X dot – I believe this might have come up yesterday in your section, so there. I allocated on the page enough for a four by four. There you go. Thank you. It’s fixed. Let’s look at that. What are the eigenvalues? What on earth have I done? My god. That was a terrible crime. There we go. Okay. But you didn’t violate the eight-second rule or something like that. When you write something that stupid down, something should say something within some amount of time. Okay, fine. What are the eigenvalues?

Student:[Inaudible].

Instructor (Stephen Boyd):All zero. Okay. So just based on that, what do you expect to see on the solution when you look at X ? Constants, right? And someone says, “Anything else?” Now in this case, what do the solutions look like? The solutions here – that’s a Jordan block of size – a single one. You are gonna see solutions that look like this. E to the zero T , that’s one. You’re also gonna see T , T squared, and T cubed. The solutions of this X dot equals AX for this thing are gonna be polynomials. Everybody cool on that? And they’re polynomials of up to degree three. Now let’s do one more thing. Let’s change that so that it looks like this. Here’s the block structure. What do you expect to see? Not expect. What would you see?

Student:[Inaudible].

Instructor (Stephen Boyd):You don’t have this, and you don’t have that, but you do have this. And finally, if it was just this, if it’s all zero, you don’t even expect – you just see constants. And of course, that’s correct because X dot equals zero – the solution is that everything is just constant. Okay? So the neural – I mean you really wanna kinda

understand all of this, but then the real neural connection you needed to make is dynamically Jordan blocks are these annoying terms of the form – they correspond to these. That’s what they – that tells you there’s a Jordan block somewhere in the neighborhood. Okay. So we’ll quit here.

[End of Audio]

Duration: 74 minutes