

Instructor (Stephen Boyd): Finish up Jordan Canonical Form. I'll probably say a little bit at the end of the section and then we'll start, in fact, the next section of the course.

So if you go down to the pad here, we can talk about Jordan Forms. You're going to want to move, I think, this way. I don't know, though. I think you got it. Okay. So last time we looked at the Jordan Canonical Form, which has many ways to describe it. One is to say that it's the closest you can get to diagonalization, if the matrix is not diagonalized. So you find a matrix T . So the T inverse AT is block diagonal. It's more than block diagonal.

Each block is, itself, upper by diagonal. It's got a bunch of repeated eigenvalues on the diagonal. On the super diagonal, it's got one, so that's the Jordan Form. It's a generalization of a diagonal matrix, is a Jordan Block. That corresponds to Jordan Blocks, which are all one by one.

Okay. Now, we looked at various things involving Jordan Form, and the real question was, what does it mean? I think we saw some of that from a dynamic point of view. We can get some of that by looking at the exponential of a Jordan Block. So the exponential of T times a Jordan Block looks like this. You get the familiar, E to the T lambda. So that you'd expect. That's the eigenvalue.

But now you can see that you're getting these powers of T . So Jordan Blocks should be associated, in your mind, with – when I say Jordan Blocks, I mean nontrivial Jordan Blocks. Jordan Blocks bigger than one by one – should be associating [inaudible] with these polynomials $X E$ to the [inaudible].

So that's the same thing in an undergraduate course. [Inaudible] forms is associated with repeated poles. So in an undergraduate class, you see S plus 1 squared in the denominator. You're going to see $T - E$ to the minus T in some solution or in the inverse of the plus transform.

So here, Jordan Block is going to take the role of repeated poles in an undergraduate case. And by the way, that's exactly correct. You can see here that if you look at the resolvent of a Jordan Block, you actually get these higher powers like this. And if it's a one by one block, you just get that. There's an interesting way to think of this. If you see that a matrix, for example, has four eigenvalues at minus one. Then you look at the resolvent. That's minus one, I , minus A , inverse. And you look at the pole and poles of the entries of the resolvent at the eigenvalue minus one.

In fact, if all of those have degree one, it means that it's four one-by-one Jordan Blocks. If you get degrees as high as three or something like that, that tells you all of them are in one block together. So degrees in the denominators in resolvents, powers of T multiplied by exponentials, these are all the symptoms of nontrivial Jordan Blocks. Okay.

Now, the same way we talk about a solution of the form E to the λT times V where V is an eigenvector as a mode of a system, you can talk about generalized modes. Now, unfortunately, the nomenclature for generalized modes is horrible. You look at 15 books, you'll get 15 slightly different definitions of what a generalized mode is, so it's not something that you need to worry about. But you should also know that the nomenclature is not consistent.

So here, if you take anything that's in the span – so T_I is the sub matrix of T . It's the vectors associated – it's the columns of T associated with the I Jordan Block. If you take any linear combination of those, you'll find that that's an invariant subspace, that you stay in that subspace. If you start in that subspace, you stay there. So you stay in the subspace, and what you get is this. All the solutions have an exponential, and then they have polynomials multiplied by them, so that's what they look like.

These are sometimes called generalized modes of the system, something like that. But again, this nomenclature is not uniform. So things of the form, you know, T , E to the λT , T squared, those are generalized modes.

Now the inverse of T , if we call it S , then if we take the inverse and then we break up the inverse, S , row by row, conformably with the T s, we get matrixes S_I . That will extract out from an initial condition, the coefficients of X of zero, for example, in a generalized eigenvector expansion. So that's what this is. This pulls out the coefficients of the block in a generalized eigenvector modal expansion of X zero.

Then you propagate them forward in time. That's this, and then you reconstitute them like that. So these are the analogues of the formulas that we looked at before when A was diagonalized. So this is the picture. That's how that works.

I mean, all of it really happens a lot of times because these things just get more complicated. They look quite similar, but the indexing gets horrible because these things have different sizes and all that sort of stuff. Okay.

We're going to do one application of Jordan Canonical Forms. It's a generic application. This is really what it's used for. It's used for showing something about matrixes. So if you want to show something about square matrixes, many things are simplified by using Jordan Canonical Form. The way that works is this. You should always warm up by assuming the matrix is diagonalizable.

So if someone says, "Show that – anything. Just make up something about square matrixes." Then you want to start by showing it for the diagonalizable case, okay? Because if you can't do the diagonalizable case, there's no point even worrying about the horrible headache of the Jordan Canonical Form.

So you first warm up by showing it for the diagonal matrixes. Then you step back and you say – now you've done that, you say, "Now, let's do the Jordan Canonical Form." Now let's do the real case, and you'd do Jordan Canonical Form. So that would be the

way to do it. We'll just see an example. It's a classic one. It's also important in its own rights. It's called the Cayley-Hamilton Theorem.

Actually, before we start this, we're just going to have a discussion about matrixes, about N by N matrixes. So if you have an N by N matrix, like this, I can have the set of N by N matrixes. That is, itself, a vector space. Its dimension is N squared. Okay?

If you have a basis with lots of bases, here's one. E_{IJ} , which is $E_{I, I}$, $E_{J, J}$ transpose, that is the matrix that has a one in the IJ entry and nowhere else. Okay? Then you have IJ equals one to N . So these are independent matrixes. I mean, this is sort of the analogue of just $E_{I, I}$, in the general case.

So the set of, say, ten by ten matrixes is a 100 dimensional vector space. This is one basis for it. Now let me ask you this. If I take any matrixes. Let's take any matrixes. So if I write down A_1 up to A_{101} , and there in \mathbb{R} ten by ten. What can you tell me about 101 vectors in a 100 dimensional space?

Student:[Inaudible].

Instructor (Stephen Boyd):They are not linearly independent. Right, which is to say, "They are linearly dependent." So many ways to write it, but one, you can say, one of these matrixes is a linear combination of the other 100. Okay? So in particular, if I write down the matrixes A_1 up to A_{101} , one of those matrixes is a linear combination of the previous ones. You agree? That just comes from the fact it's 100-dimensional.

Now to say that the powers are dependent basically says there's a polynomial because a linear combination of powers is a polynomial. It says there's a polynomial of that matrix that is zero. One polynomial that's always zero is the zero polynomial, but I mean a nonzero polynomial.

So there exists a polynomial, P , for which P of A is zero, and P has a degree – well, in this case, it could have a degree of 101. I could argue in a fancy way that it has to have a degree of 100. Okay? So there's a polynomial – a nonzero polynomial of a matrix makes it zero. Now comes the interesting part. It turns out that in this case, there's actually one of degree ten. In other words, you don't have to go up to 100.

When you take the powers of a matrix, you go A , A -squared, and you think of these as vectors. They're independent. But by the time you hit N , they're dependent, period, and they're dependent, of course, from then on. So it says that the dimension of the subspace spanned – that's the Cayley-Hamilton Theorem. Cayley-Hamilton Theorem says this. If you have any end-by-end matrix and you plug in its characteristic polynomial, if you evaluate the matrix – evaluate the characteristic polynomial with the matrix, that's an overloading here, to talk about a polynomial evaluated at a square matrix.

You get zero. This says something very interesting. It says lots of things. Basically, it says the following. If you take I , A , A -squared, A -cubed, and the powers of a matrix, it

says by the time you get to N , that A to the N is a linear combination of the previous A s. So you're in an N -squared dimensional space, but you actually only sweep out by powers, at most, a subspace of dimension N . So that's what this is saying.

It's going to have important implication. You can actually already think of some of the implications. You've already seen a lot of cases where powers of A come up. They come up in discreet time-dynamical systems, for example. We've seen that.

Actually, let's look at a wrong proof of this. I'll show you a proof of this that's wrong, but it sounds really good. If I did it quickly, you'd probably go for it. Actually, I could probably get away with it. I don't know, but that's only because the class is this early. Let me just try it.

So the wrong proof goes like this. It says, "What's the big deal?" Tie of A is debt – well, how do you evaluate a polynomial? You just plug – wherever you see S , you put in A . So you go A times I minus A , and I know what that is. That's zero, and that's zero. That would seem to be the end of the story. Like it? Oh, for sure, if I'd done that quickly, you absolutely would've bought that.

In fact, especially if I waved my hands wildly and said, "Obviously," and "Clearly," and things like that, and then quickly moved on, you would've gone for it for sure. I can just tell. No? Okay.

So it's a completely wrong proof here. You can't just plug a matrix and have A – actually, the correct thing to do here would be to say, "This original matrix over here, this S , actually looked like this." So if you really propose to plug A in there, you should've made a matrix that looks like this. Like that. You can see now, things are going downhill. It's just not working. Okay? So this is a non proof.

Let's poke through and see a real one. Let's just do an example first. It's a stupid one. Here's a matrix, two by two. One, two, three, four. You work out the characteristic polynomial. That's S -squared minus five S minus two. The Cayley-Hamilton Theorem says that if I plug in – if I evaluate this polynomial with A , then I'll get the zero matrix. So I work out A -squared and minus five A and minus two I . These may or may not be correct. I guess this term is correct. That's probably A -squared.

The point is, if you actually worked out what all this was, you'd actually get zero. We can audit it. We'll do an audit. We'll check the one-one entry. So seven minus five minus two, it worked. Okay? So that's the Cayley-Hamilton Theorem.

Very interesting, it says that this matrix, it's square, is actually a linear combination of the identity and itself. That's – in fact, you can read that right off of here. If you have this as zero, it basically says A -squared is five A plus two I . So it says that power of A is a linear combination of I and A .

Okay. Now, if you think carefully about what that implies, it's actually really quite interesting. It says that if you take a matrix to any power, we're talking about positive powers, now. If you take a matrix to any power at all, it is always a linear combination of I , A , A -squared up to A and minus one. Period. Okay? So that makes perfect sense.

How do you know that? Well, we can show that very quickly. What we do is if you take a power of A , what I'll do is I'll do polynomial division, and I'll divide S to the P by the characteristic polynomial, and I can write it this way. It's a quotient polynomial times this plus – and then a remainder. This thing has to have degree less than N . Okay? Which is the degree of this, my divisor.

So I get a remainder polynomial. In fact, the remainder polynomial is exactly what the coefficients are in this expansion. Let me actually say a little bit about this because it's actually kind of interesting. It says if you have any analytic function of a matrix – oh, we just put a new problem on this. But you've already seen it for exponential. So if I have any function which is defined by a power series, so this is α_0 plus $\alpha_1 U$ plus $\alpha_2 U$ -squared and so on. Then we overload this analytic function to apply to end-by-end matrixes, and it's just this way. α_0 plus $\alpha_1 A$ plus $\alpha_2 A$ -squared and so on. Like that.

You've already seen this once. You've seen the exponential. In fact, you can do this for, actually, any analytic function, even ones that – we'll just – just any analytic function like this. This is an analytic, add zero. So any analytic function, you can define this.

Okay. Now, let's see what Cayley-Hamilton Theorem has to see about this because it says something very, very interesting. It says the following. It says that when I go up to A to the N , there's a term, $\alpha_N A$ to the N , and then I get α_N plus one, A to the N plus one. Like that. These keep going, but every term from N and on is actually a linear combination of I up to A^{N-1} .

You know what that means? That means any analytic function of a matrix actually must have the following form. It must look like, you know, γ_0 plus $\gamma_1 A$ – and you can stop at $\gamma_{N-1} A^{N-1}$. Period. So the exponential of a matrix is a linear combination of I , A , A -squared, to A^{N-1} . By the way, so is the inverse. We'll get to that in a minute, and that's the basis of a huge family of methods for solving large problems. Okay. This is all a consequence of the Cayley-Hamilton Theorem.

All right. Let's look at the inverse. That's even more interesting. It has huge practical implications. Absolutely huge. Here it is. Let's take the inverse of a matrix. Now, of course, not all matrixes are invertible, so somewhere in here, we're going to have to encounter something that tells us that the matrix is invertible. So I'm sort of on – I'm saying, "Please suspend –" assume it's invertible. Somewhere in here, we better encounter something about that.

What we do is we write this. This is a characteristic polynomial. It's A to the N plus A and minus one, A to the N minus one, plus all the way up to A zero I equals zero, here. Okay. And what I do is simply this. I put this one term on the other side, and I divide by A zero, and I get this.

IN fact, this says it's more than that. I then factor out A . This is the inverse there. That's A inverse. It says that the A inverse is an explicit linear combination of I A up to A N minus one. Okay? In fact, the coefficients, we can say exactly what they are. They have to do with the characteristic polynomial. They're A one over minus A one over A zero, minus A two over A zero, minus one over A zero. The catch is this. A zero, the constant term in the characteristic polynomial, is something like – does anyone – do we have a homework problem on that?

We may have just assigned it just now. It's either debt A , or maybe there's a plus/minus one in front. I think there might be a plus/minus one in front of it, or is it just debt A . Anybody remember this? What is it? Debt minus A ? That's fine. So I'll just write it this way. It's plus/minus one debt A . It depends on – what is it?

Student:[Inaudible].

Instructor (Stephen Boyd):Here? I am. You mean like right there? Wow, that's cool. And now, through the miracles of SSH, it'll be fixed in a few minutes. I lost Jacob though. Okay. Thank you.

Of course, I could argue that's – when you get certain levels of development, you get to a level where you're so cool that you can write a scalar plus a matrix. It's understood. The I is understood. I actually haven't achieved that level of development, but had a reached it, I could write stuff like that. Right. Okay. All right.

So this is actually just an interesting fact. I'm just going to say a little bit about it because it's fun, and it's just a little two minute aside, and it's just for – this is just for cultural background. This is cool stuff everyone should know about. Just make a short story about this.

So let's talk about, for fun – this is just a little aside, and it's just for fun. Let's talk about solving $A X$ equals Y where A is square. Okay? Rule of it's two by two. You can do that by hand. The three by three, I don't know. Somebody might force you to do it by hand. Maybe you should do it once. Who knows? If you do it more than once, it's a complete waste of time, especially because the method is not used for anything. However, I suspect that many of you have been subjected to that more than once, which I think is the right number. Although, one could make a very good argument for zero, too.

All right. Then you grow a bit older, and you realize that the way you solve this – and we're just talking about calculating this. It's not a big deal. The way you do this, actually, is you use a computer. You might use MATLAB, but you must remember that, in fact, MATLAB is doing absolutely nothing but parsing this. It's passing it to a high quality,

open-source software called LAPACK. I'm just telling you this because these are things you should know, okay?

The computational complexity is on the order of N^3 . So it grows rapidly with N , and let's see. On a current machine, 1,000 variables, 1,000 equations can be solved in, easily, under a second. That makes some predictions. How about 100 variables, 100 equations. If I can do under a second and it's in cubed, how about 100 variables, 100 equations? You should know these numbers, actually, because they have serious significance. How fast can you solve 100 equations and 100 unknowns?

First of all, if you use your stupid grade-school method, otherwise known as the method taught in most linear algebra classes, by hand, that would take you a long time. You can laugh at that, but, you know, people actually did solve linear equations by hand. So for example, as part of the Manhattan project and things like that.

Okay. So how fast is this for 100 by 100? It's ten times smaller. Scale's like N^3 . What's the factor? It's 1000 to one. So how fast can you solve 100 equations and 100 unknowns? Under a millisecond. Now, I know you're not impressed by anything, but you should sit and just think about that for a second. You want to be impressed by that, you can imagine how long it would take you to do – well, that's silly. That's stupid. That would be, like, have you ever done JPEG encoding by hand?

It's really hard. So it's kind of stupid, but there's very few people who have really thought through and realized the implications that you can solve 100 linear equations and 100 unknowns with total reliability in under a millisecond. So these are just amazing numbers, right?

I'll leave it there. No problem. So you go up to 1,000. Everything's fine. 2,000, 3,000. It's growing like N^3 , so these things become macroscopic times. So for example, 10,000 by 10,000, assuming the memory is there, is going to be 1,000 seconds, so we're talking ten minutes or something like that. Right? Something – you know, 15, 20 minutes, okay? Something like that.

So these methods kind of – unless you're going to do this on some exotic cluster of machines and all that kind of stuff, these methods, for just a normal person, they kind of lose – they become inconvenient at around 2,000, 3,000. Okay? Then you move into another regime where you take into account the sparsity pattern of A . So the fact that A 's got a lot of zeros in it. There's methods for that, where you avoid all the zeros and things like that.

These methods will get you to 50,000, 100,000 variables, things like that. I mean, quite reliably. Then you move into the really big problems, and that's problems where X has a dimension of a million or 10 million. The people who did this – these come up all the time, for example, if you do real medical imaging, as an example. Then X represents some density, or something like that, in voxels. So the numbers are just huge. They can

be very, very big. Even just a little slice that's – if you have, whatever, 256 by 256, these are big numbers.

If you solve PDEs, you're basically solving equations like that. But for X , maybe 10 million or something like that. These base methods, where you basically factor A and all that, they're not going to work. In fact, you can work out what 10 million cubed is, multiply that in seconds, and you'll find out that this is not going to work. It doesn't matter because you don't have the storage anyway.

So how do you solve huge equations like this? It turns out for many of those huge equations, the one thing you do have is you have a fast method, given Z , to multiply A by Z . That might be through some kind of Fourier transform, Radon transform, something special in your problem, specialized code. It can be all sort of things, but you have a fast multiply. In other words, you don't – first of all, you can't even store a 10 million by 10 million matrix. So in fact, what you really have is you have a method or a function that evaluates matrix vector multiply.

Now, if I have a method, that's all I have is a method for matrix vector multiply. By the way, this is often associated with an inversion problem. This is a simulator. For example, let's do medical imaging. X is 10 million variables. That's some density and a whole bunch of – 10 million voxels. If I ask you, "What is $A Z$?" I'm actually saying, "Suppose the density were Z . What would my MRI or whatever it is, my PET or whatever it is, what would I measure?" That's what multiplying by A does. That's the forward simulator.

So basically what I'm saying is this. Imagine a situation where you have access to a forward simulator that's fast. It's not just stupid. It doesn't – you can't even store a 10 million by 10 million matrix. So it does it fast. If you do this, then given Z , I can actually calculate – if I can calculate $A Z$ fast, I can call it again, and I can get A -squared Z . I can call my simulator three times, and I can get A -cubed Z .

That means, actually, I can make N calls to my simulator, N minus one, and I can calculate Z , $A Z$, up to A and minus one Z . By calling N minus one times my fast simulator and multiplying each time. Everybody agree? Okay. We're doing this for 10 million variables.

Now, you just calculated these. Let me show you something super cool. Look at that. If you knew or estimated the spectrum, you've actually just solved the equation. If you plugged in these coefficients here, you would actually have gotten – you would've ended up with A inverse Z . Everybody see what I'm saying? The point is, it's not a minor fact that the inverse of a matrix is a linear combination of $I A$ up to $A N$ minus one. It's not just – I mean, this has serious implications, one of which is kind of at the root of the ability to solve absolutely huge equations.

Of course, I haven't told you how you can calculate these. It turns out, you can actually calculate these using something called conjugant gradients, but that's beside the point. The big picture is this. Inverse of a matrix is a linear combination $I A$ up to $A N$ minus

one. Those, you can actually – although you can store those matrixes if their 10 million by 10 million, you can actually get A to the eight times Z by doing eight forward simulations. There was a question.

Student:[Inaudible]

Instructor (Stephen Boyd):What's that? I'm evaluating what?

Student:So multiplying A by Z doesn't [inaudible]?

Instructor (Stephen Boyd):Doesn't save anything? What's that? Oh, how do I get A-squared Z?

Student:[Inaudible].

Instructor (Stephen Boyd):Oh, yeah, yeah. Exactly. So if I did this the stupid way, by storing A as a matrix and then actually doing a matrix vector multiply, right, which would be beautifully blocked out for me and optimized for my cache sizes and things like that, it would be fine. But the point is that you work out all the arithmetic and find out that the whole thing's over N-cubed. So you have to have a much faster method of multiplying A by the – in fact, the typical methods, instead of N-squared, might be $N \log N$.

Less than N is unusual because you just have to write the answer. If someone – I have to give you Z, you have to read all the entries in Z, and you have to give me the answer back. So it's not going to be faster than N. If it's N-squared, you might as well just do it the old matrix vector multiply method. So all the action is in between, and it's usually $N \log N$.

Student:I was going to say, you can store all of those vectors [inaudible].

Instructor (Stephen Boyd):No it wouldn't. Oh, sorry. You have a good point. Yeah, that's a good point. Yeah, so you'd add them up one by one. Right. That's a very good point. You wouldn't calculate all these vectors separately, in which case, you've just stored a million by a million. Of course, you couldn't do it. You're absolutely right.

So what you'd do is you'd actually evaluate the polynomial the right way where you'd accumulate. You'd multiply by one thing. What's that called, when you evaluate a polynomial that way? It's got some famous method. Luetin method. I don't know. Who knows? But thank you, that's a very good point.

Okay. All right. Well, that was just as aside, just for fun. Let's look at the proof of the Cayley-Hamilton Theorem, and you always start something like this by assuming – you warm up by assuming A is diagonalizable. So we say, we assume A is diagonalizable, we have T inverse eight. T is lambda. Characteristic polynomials in terms of the eigenvalues

is just this thing worked out here. Then we work out π of A , and that's going to be πT $\lambda^{-1} T^{-1}$.

Now, when you take a polynomial of a similarity transform, you can pull the T out. You've seen that a couple of times. This is π of λ . That's a diagonal matrix. So we just have to show the π of λ , zero. But this thing is a diagonal matrix. Each entry looks like this. It's λI , minus $\lambda^{-1} I$ is the first one.

So each of these – not one of these matrixes is zero, but each of them, the I one of these matrixes has a zero in the $I I$ position. So when you multiply out all these diagonal matrixes, that's the same as multiplying – for example, the I th entry is you multiply all the I th entries. One of the them is zero. In fact, the I th one. That means the whole thing is zero. So it was that simple.

Now for the general case, what you have to show, we've worked out – that's the Jordan Form. What you have to show now is that the characteristic polynomial evaluated at the I th Jordan Block is zero. So that is this. Here, you have to multiply $J I$ minus $\lambda^{-1} I$ to the $N I$.

This is the first term. Actually, I'm not going to care about any of these matrixes, because this is the one that's going to do the trick, the I th block. The I th block here, when I subtract $\lambda^{-1} I$, the zeros go away. I get these upper triangular matrix, which you should recognize as an – that's a downshift. I should remember these someday, but anyway, let me just do a quick experiment here.

Or is it an upshift? It's an upshift. Thank you. It's an upshift. I knew that. I wasn't confused. All right. So that's an upshift matrix, and you raise it to the $N I$ power, which is the size. If you call upshift N times on an N by N matrix, there's nothing left – sorry, on a vector, there's nothing left. So this is actually the zero matrix here. So it works. It's just a little more complicated, but it works.

Okay. So that finishes up the Jordan Canonical Form. I have to say, I wish it were an issue – I wish it were something that was only math and didn't have any real consequences. Unfortunately, I can't tell you that's the case. So it does have consequences. In fact, systems with nontrivial Jordan Block occur all the time.

Maybe the most famous one – we'll get to it, but the most ubiquitous one, it's probably each of you right now. Probably, in fact, almost certainly, each of you has a piece of electronics with you that involves a nontrivial Jordan Block. So if anything you have with you, which would be an iPod, computer, anything, cell phone, there's going to be an FIR filter in there somewhere for something. Certainly for the audio and probably for other things as well.

If you have an FIR filter, I got news for you. That's a Jordan Block. Actually, for the same reason that it's an upshift. If you look at what an FIR filter is, it's nothing – well, I'm getting ahead of myself. It's something that looks like this, right? It's a bunch of delayers,

like this, with an input, and then a bunch of coefficients multiplied like this. So a bunch of coefficients, and then these things are all added up. That's the output. There you go.

If you work out what the A matrix is for this, and $X_{t+1} = A X_t + U$. A is the upshift matrix. So that means you're carrying, on your person, a nontrivial Jordan Block. Not that it matters, but, you know. Actually, it's good to know about these things because it means that there are things that hold for diagonalizable matrixes that aren't true for systems with Jordan Blocks. If you imagine that those things are true, you could easily get into serious and actual trouble.

I don't mean actual trouble in the sense that you'd have a mathematical misconception. Many mathematical misconceptions, actually, are harmless, except when you're in a math class. These would not be harmless because they would actually have real implications. You'd make assumptions about things, make predictions about how things would work or not work, and you would just be wrong.

Okay. Well, let's move on to the next section of the class which is, we're going to look at linear dynamical systems with inputs and outputs. So we'll look at this. It's not too much different. Some things are kind of interesting here.

So here, we're going to bring in this term, and that's the output – we didn't even have that before. I mean, we could have, but no reason to worry about it. So you have $A X + B U$. So here, you've distinguished these two terms. People would call this, actually, the – it depends on the field. You would call that the drift term, would be $A X$, and then $B U$ would be called the input term. The drift term, of course, that makes perfect sense. It's sort of what would happen if you were zero. So that's what the drift term is. There's other names for that. I can't remember them right now.

Okay. The picture is something like this. Up until now, the picture was this. You'd have a face plane, and you'd draw the state here. If there is no input, you simply calculate A times X . You get \dot{X} , which is actually – it's a vector, telling you where X is going instantaneously. You would draw that, rooted in X , here. Now you can see the direction X is going in and the magnitude that vector tells you how fast it's going. That's just \dot{X} . The point is that what makes it interesting is that X then moves a little bit, and where it goes, the drift term, changes a little bit. So it's now undergoing some curve. Actually, you know exactly what it does now. It's sines and cosines, invariant planes and all that stuff. So you know how that works.

Now what happens is this. We have inputs, and inputs allow you to change the velocity vector in a very simple way. Here's a real simple example. Here's \dot{X} , if you have U equals zero, here. That's $A X$ would be this velocity vector here. If U is one, so this is B , you actually go in this direction. If U is minus 1.5, you go here. Okay? So for example, if I – now you can actually visualize, at least on a vector field, the affect of messing with U . So imagine U as a joystick.

Suppose U , for example, were 10,000. What would be the velocity – what would be the direction of X in this case? It's pretty much aligned with B at that point. So in fact, you would say it this way. When U is 10,000, the BU term has completely overwhelmed the natural – you know, the natural dynamics or the drift term. Basically, X is now going in the direction B , and with a very big velocity. Okay? So that's the picture.

In fact, you can imagine – you know, you imagine your choice of U basically says that your velocity vector can lie – if it's rooted here, it can lie along this line here, and that's your choice in U . Okay. So let's get some interpretation of these things. If you write this out column-wise, you get \dot{X} is AX plus B one, U one down to – I guess it's M inputs and P outputs. I think these are converging on standards – I mean, they're just, of course, conventions.

These are one of the [inaudible] use of M , and what you can think of this is this. It says that the state derivative, that's \dot{X} , it's an autonomous or drift term, and you get one term per input. So you get this. So each input gives you, essentially, another degree of freedom for \dot{X} , assuming the columns of B are independent, which is often the case. Not always, but often the case.

So assuming the columns are independent, it gives you another degree of freedom for \dot{X} . So that's the picture. You can also write it row-wise. So you can say that $X\dot{I}$ is – that's the drift term. It basically says it's an inner product of your input vector with – now that's a row of the B matrix, and that tells you that. So, for example, if you see that a – what would it mean if the third column of B is huge for that system? It has a meaning. I don't want details, just gross meaning. If the third column of a matrix B is huge, what does it mean about it?

Student:[Inaudible].

Instructor (Stephen Boyd):It says what?

Student:[Inaudible].

Instructor (Stephen Boyd):Exactly. So the third input, that says that system is extremely sensitive to the third input. The gain from the third input or whatever is high or something like that. It says that a small value of U three – that's the third input – is going to cause a huge deviation of \dot{X} from its drift direction toward the actual direction.

What if the third column of B is huge. What does that mean? It's got a meaning, too. Suppose it's way bigger than all the others. That says something. What does it say?

Student:[Inaudible].

Instructor (Stephen Boyd):It says what?

Student:[Inaudible].

Instructor (Stephen Boyd):No, I'm talking about the third row of B. The third row, sorry. The third row of B is huge, it basically says this. It says that all of the inputs have a tremendous affect on X_3 on the third component. Now, X_3 , by the way, can couple back in and have an affect on X_1 , X_2 and X_4 all those. No problem, but the immediate affect on where the thing is going mostly affects X_3 .

All right. A block diagram of the system is this. Looks like that. Your input comes in. This is called the feed-through term because I guess it just feeds right around here. That's the feed-through term. B converts the input to basically \dot{X} terms. \dot{X} has two components. That's \dot{X} , is the input to a bank of integrators. It's $A X$. That's the autonomous system there, plus $B U$. Then exit ϕ comes out, gets multiplied by C to form the component of the output due to the state. That's added into $D U$. That's the feed-through term.

So everything here has the obvious interpretation. So for example – a lot of these we've seen. For example, if I told you that C two five is huge, something like that. What does it mean? It means that the second output, Y_2 , is mostly dependent on X_5 . Okay? You can go through it and make all the interpretations, but I think they're kind of obvious.

This block diagram is interesting when there's structure. So here would be an example. Suppose A is block upper triangular, and B has this – is also, in a conformal way, the bottom half is zero. You get something like that. If you draw the block diagram of this out, you get this. You get U coming in. It multiplies B one. I should say, that's the autonomous system. We've already seen that picture, actually, and it's kind of interesting. We interpreted this as saying that X_2 affects X_1 . But X_1 does not affect X_2 because there's no arrow going down.

Now when you – note that U also does not affect X_2 . So the real system looks like this, and you can see a lot of things here. You can see, for example, X_2 is not affected by U. This is actually fairly important because it says – if someone says, "Could you please take the state to zero," or, "Take it to this desired state," then you'd say, "I can't." They'd say, "Why?" You'd say, "Because with my input, I can't affect X_2 ." They'd say, "I can't believe it. We paid so much for those actuators, and you still can't do it. We're going to find somebody." Anyway. You get the idea.

Okay. Now we'll look at the analogue of transfer function. I'm actually just curious how many people – I guess you've read, EE, you can't possibly avoid this, right? I don't know about other people in other departments. How many people have actually seen transfer function? Is that everyone? Come on. You're just holding back. You haven't heard of transfer function before? That's cool. What departments are you from? What are you in?

Student:[Inaudible].

Instructor (Stephen Boyd):CS. Okay, that's consistent with things I know. And your department?

Student:CS.

Student:CS. Um hm. A pattern is emerging. Okay. That's fine. No problem. It's not that big of deal. Actually, it was a very big deal last century. It's going to be much less of a deal this century. Of the things – if you're going to not know about something, it should be something whose D importance D T is negative. You made a good choice. Yeah. So all right. Good.

All right. So let's take the Laplace Transform. By the way, that statement holds for the Laplace Transform, too. Let's take the Laplace Transform, X [inaudible] equals $A X$ plus $B U$. You get this. We did this before, but with this. We have the Laplace transform now of the input here. That's $B U$ of S . We just solve for this. We get X of S . There's our friend, the resolvent. It's X of zero. That's the term we saw before, and if you take the inverse Laplace transform that, you get the matrix exponential.

Then here, I have something interesting. I have a product. Actually, I have a product of two things. I know what the inverse Laplace transform of this is. I know what the inverse Laplace transform of that is. That's U . Little U of T . Now, this is the product, and the inverse Laplace transform of a product is the convolution of the two.

So you get X of T is E to the $T A X$ of zero, plus this interval, zero to T , E to the T minus τ , $B U$ of τ , $D \tau$. So that's the solution.

By the way, if you don't know about Laplace transform, it's not that big a deal. You could also get this formula directly, I think. They did, I believe, in the 18th century or maybe the very early 19th century using something called integrating factors or who knows. That's the solution. So I think it's that simple. You can also just check it by differentiating.

All right. We are going to interpret this. That's important, even if you're in CS, actually. It's very important. Erecting this stupid barrier, learning about Laplace transforms and transfer functions and stuff is a pity because this one is actually really important for everyone to understand, so let's look at it.

So it says that the state has two terms. One is this term. We're very familiar with that term. That's basically what happens if there's no input. This term is really interesting. It's a convolution, and it's a convolution of the input with something over here. The function, E to the $T A B$, that's called the input to state impulse matrix. We'll see why in a minute.

$S I$ minus A inverse B , the resolvent times B , that's called the input to state transfer matrix because – we'll see why in a minute. Well, I can tell you why right now. If the initial state were zero, you could see that the Laplace transform of the state is the Laplace transform from the input times this thing. So this is called the transfer matrix, from input to state. Okay.

Now we'll plug in the readout equation. You get Y of S is C, SI minus A inverse, X of zero, and then you can rewrite the whole thing. This is quite familiar to us, in the time domain. That's this thing. Basically, that's the component of the output due to the part of X that's autonomous or something. The zero input X . I mean, every field has a different name for that. Yeah.

Student:[Inaudible].

Instructor (Stephen Boyd):This thing? Oh, no. I call that – that's the convolution. I was going to say I call, but not only do I call it, but everybody would call that the convolution. What made you suspect the legitimacy of this convolution? You said it was almost – because what?

Student:[Inaudible].

Instructor (Stephen Boyd):The [inaudible] the integral. Ah huh. Okay. Well, so my response – I mean, I agree with you. Often, you see a convolution where the integral goes from minus infinity, plus infinity, something like that indeed. However, here, we'd be in some – we're actually – I guess we wouldn't actually. Yeah, you could write it that way. The zero to T is only because you start from here, from X of zero. So I'll get to that later.

But this is the convolution. This is what – if you walk up to someone and say, "What's the convolution?" Actually, they'd give you two things. The minus infinity, plus infinity and this one. They coincide if you do things like agree that U is zero for negative T , for example. Question?

Student:[Inaudible].

Instructor (Stephen Boyd):Can you –

Student:[Inaudible]

Instructor (Stephen Boyd):Yes.

Student:[Inaudible].

Instructor (Stephen Boyd):Right. That is correct. That's right. Right. Now convolution in the scalar case is commutative. Right. Absolutely. In the matrix case, it's not, but it inherited that just because these are matrixes, so that's a very good point. But we'll get to these. Okay.

All right. So the output looks like this. It's got this term, and then it's got this term here, which is a function, which comes up a lot. $C S I$ minus A inverse B plus D times U of S . In a time domain, it looks like that, here. You get all sorts of things. This function up here, this matrix, is called the transfer matrix or transfer function of the system.

This thing, here, well, if I evaluate it with T, so C E to the T A B plus D times the delta function at zero is called the – that's called the impulse matrix or the impulse response of the system. We're going to find out why soon.

Okay. So the simplest thing is this. If you have no initial condition, than this term goes away, and I have Y is H U. And it looks like that. By the way, this is – I don't want to make fun of it. In a time domain, you'd write Y equals H star U. This is, by the way – if anyone said, why on earth do you Laplace transforms and all that kind of stuff, this is why you do it. You do it because you have some huge structure or something like that with 40 variables. You have six inputs and things like that. Fantastically complicated. Each input affects this thing dynamically in some different way.

They all couple together. It's complicated enough, even if there are no inputs. You throw in the inputs, it's complicated. Very, very complicated. That doesn't mean we won't understand it, you can't write three-line programs that actually work out – write small snippets of code that actually do serious things, work out good inputs and all that kind of stuff. Bottom line, though, is that it's really complicated. However, when you take Laplace transforms and the input is zero, you get down to this beautiful formula. It actually goes back, oh, basically way into – because basically it said, "It's just multiplication." But you're multiplying complicated objects, Laplace transforms, so that's what it comes down to.

So this is kind of why you do this. Actually, it's interesting, this religious arguments we have, not real ones, but they are, actually. It turns out, in ISL, let's say, there were several faculty members who lines up and said that the Fourier transform and the Laplace transform are absolutely fundamental. They're fundamental. That's the real thing. Then others of us – I'm on the other side – said, "No, no, no. Convolution is the real thing." If we were smarter, if we had the ability to look at this, and actually, after a while you can get good at this. To look at it, then you can look at it and say – look at a convolution and go, "Yeah, that's gonna give a good ride," or, "I'm not getting in," or something like that.

Then you wouldn't need this, but they're just different representations of the same thing. Anyway. All right. Now here, HIJ, if H is this transform matrix, that C S I minus A inverse B plus D. If you take IJ, that's the transfer function from input UJ to output YI in the absolute undergraduate EE. But also ME and some other areas who've learned about transfer functions. That's the transfer function. From the Jth input to Ith output. Okay.

Let's look at this impulse matrix. That's this thing. Here's what it is. It's very interesting. It says that when the initial state is zero, the output – they have a beautiful number of – I mean, how many Askey characters – five Askey characters represents the input/output behavior of any linear dynamical system. It's just Y equals H [inaudible] with U. It's an integral, but that's it.

By the way, this has got everything in it. You've got the coupling from different inputs to different outputs. You've got coupling across time. That's what convolution is. So it's all there, five Askey characters, and it's basically this. It's a beautiful equation.

By the way, this is an equation everybody should understand, even if – whatever you think of Laplace transforms and transfer functions and everything, this one is unbelievably important. Actually, we should go over it and understand absolutely everything about it.

All right. So let me see if I can get this right. All right. Here it says – let's actually do scaler because we do have some people who maybe haven't seen convolution before. So we're going to do scaler first. This is review for everybody else, and I'm going to write is zero to T, H – I'm just trying to make it look like down here.

T minus tau, U of tau, D tau. So you want your scalers, H – everybody's a scaler. So what this says, I'll give my interpretation. For example, when I used to teach 102 – EE 102. So the interpretation is this. It says that the current output is actually an integral, but let's call that something like a linear combination. It's a mixture of the input in the past because I'm only going to refer to U of tau here, between zero and the current time. So it's a mixture.

The question is, "How much do you throw into the mixture to form the current output?" The answer is you multiply it by weight, which is H of T minus tau. So in fact, if T is seconds, tau here should be seconds ago. That's what they are. That's the actually formal unit of tau here, is seconds ago. To see if this makes any sense, I could ask you the following.

Suppose that H of T is – let's say that H of seven is really big and positive. What does it mean? What if H looks like this? It goes like this. What does that mean when you have a convolution? What does a convolution do?

Student:[Inaudible].

Instructor (Stephen Boyd):It says what?

Student:[Inaudible].

Instructor (Stephen Boyd):It basically says the following. It says that when you form the current output, this tells you how much the current output depends on what the input was. If this is tau, these should be seconds ago. So it means the following. It means that H – it does depend on what it was in the recent – in the last five seconds a little bit, but basically it depends a lot on what the input was seven seconds ago. Everybody got that?

Then it depends a little bit on what it was eight, nine, ten and less and less, but basically – so it says that basically in this case, the simplest thing is – this convolution integral is actually something like a seven-second delay. Everybody got this? Now, if H of T, for example, decays, it means that this system has fading memory. Why? Because if H of T – as T gets bigger, H gets smaller, the interpretation of H of 100 is how much the current output depends on what the input was 100 seconds ago. That's what it means.

If I make 100, 200 and 300, and it's getting smaller and smaller and smaller, it says the amount by which the – what happens, 100, 200, 300 seconds ago, affects the current output. It's getting smaller and smaller. It's got fading memory. Okay? Something like that.

Actually, once you see this, you should just sort of think of this as everywhere you go. For example, let's let Y be – let's make Y river flow, and let's let U be rainfall in a region. Let's talk. What does H look like? What does it look like? I mean, just grossly, obviously. I just want a gross idea of what H might look like. What do you think? Go ahead. What?

Student:[Inaudible].

Instructor (Stephen Boyd):I don't know. I want – I don't want – just want no exponential, no math. I don't want to hear any math. You can use words like big, small, positive, negative, no more. Okay, what? It should decay?

Student:[Inaudible].

Instructor (Stephen Boyd):Okay. Well, it might do something like this. It might be small at first, and that would represent immediate run off. That stuff that got into the river first, but I think, actually, it would take a while because the water comes down, it goes down through tributaries into creeks and things. Then it builds up. So it might actually – actually, the flow might depend quite a bit on what it was several days ago.

If you can believe this story or not, for the record, I have absolutely no idea how this works. I'm making it all entirely up, just in case you – but I have no idea. I'm just making it up, but it could be right. It might look something like this.

First of all, could it be negative? It's probably not negative because you're basically – if it's negative, it says, "Hey, that's kind of cool. The height of the river today actually is suppressed by additional rainfall 2.2 days ago." So I don't know. It might look something like this. It might be small, and it might kind of go up like this, and then eventually it would die down again when all the water's come out. Something like that. Does that look cool? So that's H .

If this has a – boy, does this have a meaning. So I'm just saying, this is important. You should just think everything is this way, actually. You don't look convinced. How about pharmacokinetics? You inject – U is the amount of a drug you inject or ingest or something like that. Then Y would be the – what its concentration is in the bloodstream. What do you think it's going to look like? Probably something like this, but we'll just leave it that way. It depends how it's administered and all that kind of stuff. Anyway, you get the same story.

Okay. Back to this. There we go. Now this is a full multi-input, multi-output system. Here's the whole thing. By the way, this would be kind of interesting. This would be, if we went back to rivers, the height at the river, so Y of T is the height of the river at some

place, I , at time, T . $U_{J \text{ of } T}$ is the rainfall in some sector, J , at time T . My question is, how does it affect it? This gives you the whole answer. This is assuming a linear model, which might or might not be the case for that.

Here's a linear model. It says that the total contribution, you sum across J , and by the way, these are not just random indices. When you're summing across J , you're actually summing – if you write code for this, it shouldn't be called J . You're summing across inputs. You're summing across the input labels, right? So that's really what's happening. That's what this sum is. This is the [inaudible] across the input labels, and then here, this integral, you can think of it as a sum, if you like. Some kind of a generalized sum, and it's actually integrating H, I, J of T tells you how much output I depended on what input J was τ seconds ago. I think that kind of made sense.

We'll look at some example, and I think this will make more sense. Okay. We have the step matrix or step response. This is not fundamental, but this is traditional, so it's good to cover it. It builds on the last centuries' stuff. So the step matrix or step-response matrix is this. It's simply the running integral of the impulse response. It actually has a simple interpretation. It's quite interesting, actually. It's when – if you want to know what S_{IJ} of T is, it basically says you apply the following input.

U_{J} is one, and all the other U s are zero. So it says, "Apply an input where the J th input is one," and this – if you multiply that, that will give you – and you look at S_{IJ} of T . You'll get the output I versus time, okay? So S_{IJ} of T gets Y_I when U is E_J .

This, you can actually – if A is invertible, you can actually work out a formula for that in the time domain, which is this. This is easy enough to verify or do from Laplace transforms. Okay. So let's look at examples because that will sort of make it all clear. So here's an example. It's three masses stuck between two walls. There are springs in between each of the masses and the walls, and there's dash pots or some kind of dampening mechanism.

We're going to apply two inputs. One is a tension between mass one and two, and the other is a tension between mass two and three, okay? So these might be implemented by a solenoid or a Piezoelectric actuator. Something. It doesn't matter. So we're able to – U one positive means we're pulling the two masses together, these two. U one negative says we're pushing them apart. Those are forces.

We'll take all constants one. So these are kilogram, one newton per meter, one newton per meter per second. Okay? So that's the picture. We're going to take – X is going to be six dimensions. The state, it's going to be some positions from some offset of these three masses and their velocities. Then you can work out what \dot{X} is. It's just this. \dot{X} is this top row. It's easy. That's zero I . It says that the block top – that's Y_1, Y_2, Y_3 differentiated – is equal to the block bottom. That's what this top – this is zero I .

Then these are various things I threw in at the right places, and they could well be – they might – I think they're correct. That's what they are. Actually, I'm sure they're right. Then

here, for example, let's just audit one column of B. So the first column says – when you see that, it says something really interesting. This column, of course, is how U_1 affects \dot{X} .

What it says is really interesting. It says it has no affect immediately on the position, and that's correct. If you apply a newton to something, you will not see any immediate affect on the position. What will happen is when you apply a newton to something, it acquires a velocity. Then because it acquires a velocity, it acquires a displacement. So this kind of makes sense.

Down here, it says its affect on the first mass is one, and minus one on the second. That's exactly right. Look at this. In the first mass, if I pull with a newton here, I'm pulling this to the right, which is positive displacement, and I'm pulling this one to the left. This tells you how much velocity is acquired per newton. The fact that it doesn't touch – there's a zero here. This zero is because this force is only between these two masses and not that one. Okay? So that's just – that's my audit of just a random column.

By the way, you should do that. Just, always, when you look at equations, just audit them. Make sure you understand what it means and all that stuff.

Eigenvalues of A turned out to be minus 1.71 plus minus $j0.71$. Minus one, and then minus 0.29. So we know already, qualitatively, what the dynamics are going to look like. That is something that decays in about two seconds. Even though, technically, it's oscillatory, it's gone. So that's a bump. That's a thud is what that is. That's more like a bump. That's more like a thud. Okay?

There's also one that – this decays in four seconds or whatever. It's not even done one cycle. So even though it's oscillatory, it would be kind of like a bump is what it's going to be. This one – that's interesting. This one decays, oh, let's say in about five over 0.29 – 15. Is that about right? I think this takes 15 seconds. I can cheat. Yes. Yes. About 15 seconds for this to decay in this mode. Meanwhile, its period would be – we can calculate that from this. The period is going to be – if it's 0.71, whatever – that's 1.4 to ten. Does that sound about right?

Remember if I'm wrong, it reflects on you because there are a lot more of you than there are of me. What do you think? Ten? 15 second decay, ten second period. That's also, by the way, not an impressive oscillation, right? It means it's going to get one and a half oscillations in before it's gone. I think I did that right. We'll find out.

All right. So the impulse matrix from U_1 is this. You can organize impulse matrixes, by the way, lots of different ways. You can plot them different ways, and they're kind of interesting. You are, after all – in an impulse matrix, you are plotting something that depends on three things. It depends on I, that's the output index. It depends on J, that's the input index. And it depends on T, which is the time-lag variable, or it's tau in seconds ago or T in seconds ago.

So here we've plotted that, and this shows you the impulse response from U one. So let's even think about what the impulse response would be. The impulse response would be this. I would take – I would grab hold of mass one and mass two. I'd grab them, and I would apply an extremely large tug, tugging them together, but very briefly. That's what would happen.

Now, actually, you can integrate the dynamics of this in your head. You should be able to, right? It's very simple. If you pull these two masses and let go real quick, here's what would happen. Of course, they would acquire a velocity, and they'd start moving toward each other. But when they move toward each other, first of all, this spring is pushing them apart. There's damping from here. This spring is pulling this guy back to the wall. This one is pulling – is extending that spring. It's pulling this guy.

Then what happens is this guy starts moving to the left, and then these things reach some zero velocity part. Anyway, then the whole thing oscillates, and what do you think? I'm just saying – look, it's not a pretty picture, but this is how you integrate things in your head. So this is the right way. Anyway. So the thing oscillates for a while, and then the damping mechanisms remove a bunch of energy and stuff like that. Let's just see if this is sort of consistent with what happens.

This says that when you tug on these, the first mass rapidly moves to the right. This moves to the right. The second mass, that's this guy, rapidly moves to the left and, in fact, it looks like it moves a little bit more. So I'm just making that up, but it looks to me like a little tiny bit more. Then you have to go back here and explain why, when you tug these things, does this mass actually only move out – this one moves a little bit farther to the left.

This is just for fun. This is just explanation, but if someone said, "Can you explain it," how would you explain it? What do you think?

Student:[Inaudible].

Instructor (Stephen Boyd):There you go. So this one is tied to a hard wall, so it's being restrained by a hard wall. This one, in fact, is being restrained by a mass connected through a spring to a wall, and this mass actually starts moving to the left, too. So the whole system is looser here, there's less pulling it back and so on. So that's why it goes a bit far.

By the way, I think I mentioned this before, when you do things like this, you have to be able to give an explanation like this. They don't have to be correct. If your story-telling skills are good enough, people will go for them. It's fine. You have to – I mean, first you just learn to do this, but then you have to be able to do the following. To get a story for this, and then have someone say, "Oh, by the way, it's wrong, and actually, these are switched." So the thing that you just gave a two-minute intuitive explanation for why one thing goes farther than the other, turns out it's wrong.

Then you have to actually not be embarrassed and go, "Right. I can explain that, too." That's what you have to do. So this is very important. It's not a formal part of the class, although it is possible, maybe we should have a story-telling portion of the final, just for fun. I mean, just a brief one. Maybe the right way – I mean, to really know if you're a master, you'd have to give a good story both ways. So we'd give you some phenomenon. We won't tell you which is which, and you have to give a convincing story.

I think that should probably be submitted in, in video to us because it's not something you write on a piece of paper. Then we would – I'd get a panel, and we would evaluate each one. Would we go for it or not. Probably, both of the things we'd ask you to explain would be false, actually, because that's the real – anybody can explain something that's true. That's super easy, right? But really, I think the way you know if you got this – you should start.

By the way, start by explaining things that are true. When you've mastered that, you can move on to the more advanced topics. We'll get into this later on various occasions. Okay. This shows the response from you, too. There's a symmetry here which is kind of obvious.

Now we'll look at our second example. It's an interconnect circuit. So it's quite real, quite important. It wouldn't be this simple, but it would look something like this. You have a tree of resistors and a capacitor at each node to ground, and I have a voltage source here. That's going to be our input, so I have a scaler input system.

I'm interested in the voltage at every, single node, like this. So I'm interested in voltage. This is quite useful if you wanted a picture for where this comes up. This could be a driving gate. You could include this resistor in the gate if you like and maybe some of that. This is a driving gate. The gate goes high, shifts from zero to one. Actually, it's just about right. You're lucky because VDD is one, right about now, on some processors.

So this goes – this flips from zero to one volt. T equals zero, and what happens here, of course, you can figure out. Again, if you have background in EE, you can integrate the dynamics of this in your head. That's easy to do. So for example, we'll make everything a one here. For EE, if this thing flips from zero to one volt, what is the voltage there right at T equals zero? What's the voltage here? These are all zero. It's zero, right?

What is it – that's one, one, one, one, one. What is it one millisecond later? My time scales are all off, but that's fine. What?

Student:[Inaudible]

Instructor (Stephen Boyd):It's still zero. That is actually, technically, a correct answer, but I said, no, I want the next level of accuracy from zero. All right, let me ask you this. Right at T equals zero, this thing flips up to one volt. That's zero. There's a volt across this resistor. It's one ohm. How much current is flowing?

One amp. So one amp is flowing into a one [inaudible] capacitor. What's the DVDT? One volt per second. Okay. So in one millisecond, the answer is the voltage here is one millisecond, about. It's not quite right because, actually, as the voltage here rises, the voltage – the charging current decreases slightly, you know. So that's fine.

The voltage here, though, it does nothing. It's not even – there's no current flowing here. Once the voltage here builds up, then you get current flowing across here, and that – by the way, that slows down the buildup here. This thing builds up, and then once the voltage here builds up, current starts spilling into this last one. So again, very roughly – and by the way, this description I'm telling you is true. I will tell you when I'm lying maybe. We'll see, but for this one, this is actually – so we should expect this one to pop up, linearly at first, but go up. It should slow down for a little bit because more current is being sloughed off to charge these guys.

This should be slower, and then this should be the last one. We'll see, in fact, that's going to be the case. Let's look at that. That's the step response, and sure enough, here it is. I claim that you could even, accurately go down here and describe things like changes and slope here this way. So you could even say things like – in fact, if you zoomed way in here, if you zoomed way in right at the beginning, what would S one look like? If I zoomed in the plot right here – we'll just do this, then we'll quit.

If I zoomed in, I wanna know, what do the voltages of these nodes look like between zero and one millisecond. That's for ones everywhere. We need somebody in EE. We already discussed it. This one looks like a line going up at one volt per millisecond. It's linear. What do you think the voltage in here looks like?

Student:[Inaudible].

Instructor (Stephen Boyd):Zero. That's actually a very good – that's a valid answer, but I want the next level of accuracy. What's the voltage here and here look like?

Student:[Inaudible].

Instructor (Stephen Boyd):What does it look like? It's quadratic. It's quadratic here. It's quadratic because it's an RC circuit being charged by a ramp. This goes up quadratically. Now, of course, when T is small, T-squared is really small. That's why I accepted the answer when someone said zero. That was perfectly cool.

What do you think it looks like here? It's cubic, and that means it's very, very small. I think I said that correctly. You know what I can tell from this? You need – we need to do exercises on intuitive integration. You'll leave this class, you can type in $\int \int \int$ times T times A, make plots, all sorts of stuff like that. But you need to do this.

The only real reason is so that you can make up stories. You can't just say, "No, that's what it does," because people don't like that. They want the human element. They want you to look at this and say, "No, no. Sorry. That's one. Switching time is that." But they

want you to say, "Yeah, but my –" And you say, well, see, at first it was charging quickly, but then work here was being sloughed off, and of course these were growing quadratically. But then, you see, they caught up. That's what they want.

We'll work on that. I think maybe for the final, I think we're going to – we'll do something. Okay. We'll quit here.

[End of Audio]

Duration: 78 minutes