[MUSIC PLAYING]

PROFESSOR: During the course, we've developed a number of very powerful and useful tools. And we've seen how these can be used in designing and analyzing systems. For example, for filtering, for modulation, et cetera. I'd like to conclude this series of lectures with an introduction to one more important topic. Namely, the analysis of feedback systems.

And one of the principle reasons that we have left this discussion to the last part of the course is so that we can exploit some of the ideas that we've just developed in some of the previous lectures. Namely, the tools afforded us by Laplace and z-transforms.

Now, as I had indicated in one of the very first lectures, a common example of a feedback system is the problem of, let's say balancing a broom, or in the case of that lecture balancing my son's horse in the palm of your hand. And kind of the idea there is that what that relies on, in order to make that a stable system, is feedback. In that particular case, visual feedback.

That specific problem, the one of balancing something, let's say in the palm of your hand, is an example of a problem, which is commonly referred to as the inverted pendulum. And it's one that we will actually be analyzing in a fair amount of detail not in this lecture, but in the next lecture.

But let me just kind of indicate what some of the issues are. Let me describe this in the context not of balancing a broom on your hand, but let's say that we have a

mechanical system which consists of a cart. And the cart can move, let's say in one dimension, and it has mounted on it a bar, a rod with a weight on the top, and it pivots around the base. So that essentially represents the inverted pendulum. So that system can be, more or less, depicted as I've indicated here. And this is the cart that can move along the x-axis. And here we have a pivot point, a rod, a weight at the top. And then, of course, there are several forces acting on this.

There is an acceleration that can be applied to the cart, and that will be thought of as the external input. And then on the pendulum itself, on the weight, there is the force of gravity. And then typically, a set of external disturbances that might represent, for example, air currents, or wind, or whatever, that will attempt to destabilize the system. Specifically, to have the pendulum fall down.

Now, if we look at this system in, more or less, a straightforward way, what we have then are the system dynamics and several inputs, one of which is the external disturbances and a second is the acceleration, which is the external acceleration that's applied. And the output of the system can be thought of as the angular displacement of the pendulum. Which if we want it balanced, we would like that angular displacement to be equal to 0.

Now, if we know exactly what the system dynamics are and if we knew exactly what the external disturbances are, then in principle, we could design an acceleration. Namely, an input. That would exactly generate 0 output. In other words, the angle would be equal to 0. But as you can imagine, just as it's basically impossible to balance a broom in the palm of your hand with you eyes closed, what is very hard to ascertain in advance are what the various dynamics and disturbances are. And so more typically what you would think of doing is measuring the output angle, and then using that measurement to somehow influence the applied acceleration or force. And that, then, is an example of a feedback system.

So we would measure the output angle and generate an input acceleration, which is some function of what that output angle is. And if we choose the feedback dynamics correctly, then in fact, we can drive this output to 0. This is one example of a system

which is inherently unstable because if we left it to its own devices, the pendulum would simply fall down. And essentially, by applying feedback, what we're trying to do is stabilize this inherently unstable system. And we'll talk a little bit more about that specific application of feedback shortly.

Another common example of feedback is in positioning or tracking systems, and I indicate one here which corresponds to the problem of positioning a telescope, which is mounted on a rotating platform. So in a system of that type, for example, I indicate here the rotating platform and the telescope. It's driven by a motor. And again, we could imagine, in principle, the possibility of driving this to the desired angle by choosing an appropriate applied input voltage.

And as long as we know such things as what the disturbances are that influence the telescope mount and what the characteristics of the motor are, in principle we could in fact carry this out in a form which is referred to as open loop. Namely, we can choose an appropriate input voltage to drive the motor to set the platform angle at the desired angular position.

However, again, there are enough unknowns in a problem like that, so that one is motivated to employ feedback. Namely, to make a measurement of the output angle and use that in a feedback loop to influence the drive for the motor, so that the telescope platform is positioned appropriately.

So if we look at this in a feedback context, we would then take the measured output angle and the measured output angle would be fed back and compared with the desired angle. And the difference between those, which essentially is the error between the platform positioning and the desired position would be put perhaps through an appropriate gain or attenuation and used as the excitation to the motor. So in the mechanical or physical system, that would correspond to measuring the angle, let's say with a potentiometer.

So here we're measuring the angle and we have an output, which is proportional to that measured angle. And then we would use feedback, comparing the measured angle to some proportionality factor multiplying the desired angle. So here we have

the desired angle, again through some type of potentiometer. The two are compared. Out of the comparator, we basically have an indication of what the difference is, and that represents an error between the desired and the true angle. And then that is used through perhaps an amplifier to control the motor.

And in that case, of course, when the error goes to 0, that means that the actual angle and the desired angle are equal. And in fact, in that case also with this system, the input to the motor is, likewise, equal to 0.

Now, as I've illustrated it here, it tends to be in the context of a continuous time or analog system. And in fact, another very common way of doing positioning or tracking is to instead implement the feedback using a discrete-time or digital system. And so in that case, we would basically take the position output as it's measured, sample it, essentially convert that to a digital discrete-time signal. And then that is used in conjunction with the desired angle, which both form inputs to this processor. And the output of that is converted, let's say, back to an analog or continuous-time voltage and used to drive the motor.

Now, you could ask, why would you go to a digital or discrete-time measurement rather than doing it the way I showed on the previous overlay which seemed relatively straightforward? And the reason, principally, is that in the context of a digital implementation of the feedback process, often you can implement a better controlled and often also, more sophisticated algorithm for the feedback dynamics. So that you can take a count, perhaps not only of the angle itself, but also of the rate of change of angle. And in fact, the rate of change of the rate of change of angle.

So the system, as it's shown there then, basically has a discrete-time or digital feedback loop around a continuous time system.

Now, this is an example of, in fact, a more general way in which discrete-time feedback is used with continuous systems. And let me indicate, in general, what the character or block diagram of such a system might be.

Typically, if we abstract away from the telescope positioning system, we might have a more general continuous-time system. And around which we want to apply some feedback, which we could do with a continuous-time system or with a discrete-time system by first converting these signals to discrete-time signals. Then, processing that with a discrete-time system. And then, through an appropriate interpolation algorithm, we would then convert that back to a continuous-time signal. And the difference between the input signal and this continuous-time signal which is fed back, then forms the excitation to the system that essentially we're trying to control.

And in many systems of this type, the advantage is that this system can be implemented in a very reproducible way, either with a digital computer or with a microprocessor. And although we're not going to go into this in any detail in this lecture, there is some discussion of this in the text.

Essentially, if we make certain assumptions about this particular feedback system, we can move the continuous to discrete-time converter up to this point and to this point, and we can move the interpolating system outside the summer. And what happens in that case is that we end up with what looks like an inherently discrete-time feedback system.

So, in fact, if we take those steps, then what we'll end up with for a feedback system is a system that essentially can be analyzed as a discrete-time system. Here we have what is, in the forward path, is basically the continuous-time system with the interpolator at one end and the continuous to discrete-time converter at the other end. And then we have whatever system it was in the feedback loop-- discrete-time-- that shows up in this feedback loop.

Well, I show this mainly to emphasize the fact, although there are some steps there that we obviously left out. I show that mainly to emphasize the fact that feedback arises not just in the context of continuous-time systems, but also the analysis of discrete-time feedback systems becomes important. Perhaps because we have used discrete-time feedback around a continuous-time system. But also perhaps because the feedback system is inherently discrete-time. And let me just illustrate

one, or indicate one example in which that might arise.

This is an example which is also discussed in somewhat more detail in the text. But basically, population studies, for example, represent examples of discrete-time feedback systems. Where let's say that we have some type of model for population growth. And since people come in integer amounts that represents essentially the output of any population model, essentially or inherently represents a sequence. Namely, it's indexed on an integer variable. And typically, models for population growth are unstable systems.

You can kind of imagine that because if you take these simple models of population, what happens is that in any generation, the number of people, or animals, or whatever it is that this is modeling, grows essentially exponentially with the size of the previous generation.

Now, where does the feedback come in? Well, the feedback typically comes in, in incorporating in the overall model various retarding factors. For example, as the population increases, the food supply becomes more limited. And that essentially is a feedback process that acts to retard the population growth. And so an overall model-- somewhat simplified-- for a population system is the open loop model in the absence of retarding factors. And then, very often the retarding factors can be described as being related to the size of the population. And those essentially act to reduce the overall input to the population model.

And so population studies are one very common example of discrete-time feedback systems.

Well, what we want to look at and understand are the basic properties of feedback systems. And to do that, let's look at the basic block diagram and equations for feedback systems, either continuous-time or discrete-time.

Let's begin with the continuous-time case. And now what we've done is simply abstract out any of the applications to a fairly general system, in which we have a system H of s in what's referred to as the forward path, and a system G of s in the

feedback path. The input to the system H of s is the difference between the input to the overall system and the output of the feedback loop.

And I draw your attention to the fact that what we illustrate here and what we're analyzing is negative feedback. Namely, this output is subtracted from the input. And that's done more for reasons of convention then for any other reasons. It's typical to do that and appropriate certainly in some feedback systems, but not all. And the output of the adder is commonly referred to as the error signal, indicating that it's the difference between the signal fed back and the input to the overall system.

Now, if we want to analyze the feedback system, we would do that essentially by writing the appropriate equations. In generating the equivalent system function for the overall system, it's best done in the frequency or Laplace transform domain rather than in the time domain. And let me just indicate what the steps are that are involved. And there are a few steps of algebra that I'll leave in your hands.

But basically, if we look at this feedback system, we can label, of course-- since the output is y of t, we can label the Laplace transform of the output as Y of s. And we also have Y of s as the input here. Because this is the system function, the Laplace transform of r of t is simply the Laplace transform of this input, which is Y of s times G of s. So here we have Y of s times G of s.

At the adder, the input here is x of s. And so the Laplace transform of the error signal is simply x of s minus r of s, which is Y of s G of s. So this is minus Y of s G of s. That's the Laplace transform of the error signal. The Laplace transform of the output of this system is simply this expression times H of s. So that's what we have here.

But what we have here we already called Y of s. So in fact, we can simply say that these two expressions have to be equal. And so we've essentially done the analysis, saying that those two expressions are equal. Let's solve for Y of s over x of s, which is the overall system function.

And if we do that, what we end up with for the overall system function is the algebraic expression that I indicate here. It's H of s divided by 1 plus G of s H of s.

Said another way, it's the system function in the open loop forward path divided by 1 plus, what's referred to as the loop gain, G of s times H of s. Let's just look back up at the block diagram. G of s times H of s is simply the gain around the entire loop from this point around to this point. So the overall system function is the gain in the forward path divided by 1 plus the loop gain, which is H of s times G of s.

Now, none of the equations that we wrote had relied specifically on this being continuous-time. We just did some algebra and we used the system function property of the systems. And so, pretty obviously, the same kind of algebraic procedure would work in discrete-time. And so, in fact, if we carried out a discrete-time analysis rather than a continuous-time analysis, we would simply end up with exactly the same system and exactly the same equation for the overall system function. The only difference being that here things are a function of z, whereas if I just flip back the other overlay, we simply have-- previously everything is function of t and in the frequency domain s. In the discrete-time case, we've simply replaced in the time domain, the independent variable by n. And in the frequency domain, the independent variable by z.

So what we see is that we have a basic feedback equation, and that feedback equation is exactly the same for continuous-time and discrete-time. Although we have to be careful about what implications we draw, depending on whether we're talking about continuous-time or discrete-time.

Now, to illustrate the importance of feedback, let's look at a number of common applications. And also, as we talk about these applications, what will see is that while these applications and context in which feedback is used are extremely useful and powerful, they fall out in an almost straightforward way from this very simple feedback equation that we've just derrived.

Well, the examples that I want to just talk about are, first of all, the use of feedback in amplifier design. And we're not going to design amplifiers in detail, but what I'd

like to illustrate is the basic principle behind why feedback is useful in designing amplifiers. In particular, how it plays a role in compensating for a non-constant frequency response. So that's one context that we'll talk about.

A second that I'll indicate is the use of feedback for implementing inverse systems. And the third, which we indicated in the case of the inverted pendulum, is an important context in which feedback is used is in stabilizing unstable systems. And what we want to see is why or how a feedback system or the basic feedback equation, in fact, let's us do each of these various things.

Well, let's begin with amplifier design. And let's suppose that we've built somehow without feedback, an amplifier that is terrific in terms of its gain, but has the problem that whereas we might like the amplifier to have a very flat frequency response, in fact the frequency response of this amplifier is not constant. And what we'd like to do is compensate for that.

Well, it turns out, interestingly, that if we embed the amplifier in a feedback loop where in the feedback path we incorporate an attenuator, then in fact, we can compensate for that non-constant frequency response. Well, let's see how that works out from the feedback equation.

We have the basic feedback equation that we derived. And we want to look at frequency response, so we'll look specifically at the Fourier transform. And of course, the frequency response of the overall system is the frequency response of the Fourier transform of the output divided by the input. Using the feedback equation that we had just arrived, that has, in the numerator, the frequency response in the forward path divided by 1 plus the loop gain, which is H of j omega times k.

And this is the key. Because here, if we choose k times H of j omega to be very large, much larger than 1, then what happens is that these two cancel out. H of j omega here and in the denominator will cancel out as long as this term dominates. And in that case, under that assumption, the overall system function is approximately 1/k.

Well, if k is constant as a function of frequency, then we somehow magically have ended up with an amplifier that has a flat frequency response.

Well, it seems like we're getting something for nothing. And actually, we're not. There's a price that we pay for that. Because notice the fact that in order to get gain out of the overall system, k must be less than 1. So this has to correspond to attenuator. And we also require that k, which is less than 1, times the gain of the original amplifier, that that product be greater than 1. And the implication of this, without tracking it in detail right now, the implication in this is that whereas we flatten the frequency response, we have in fact paid a price for that. The price that we've paid is that the gain is somewhat reduced from the gain that we had before the feedback. Because k times h must be much larger than 1, but the gain is proportional to 1/k.

Now, one last point to make related to that. One could ask, well, why is it any easier to make k flat with frequency than to build an amplifier with a flat frequency response? The reason is that the gain in the feedback path is an attenuator, not an amplifier. And generally, attenuation with a flat frequency response is much easier to get than gain is. For example, a resistor, which attenuates, would generally have a flatter frequency response than a very high-gain amplifier.

So that's one common example of feedback. And feedback, in fact, is very often used in high-quality amplifier systems. Another very common example in which feedback is used is in implementing inverse systems.

Now, what I mean by that is, suppose that we have a system, which I indicate here, P of s-- input and output. And what we would like to do is implement a system which is the inverse of this system. Namely, has a Laplace transform or system function which is 1 over P of s.

For example, we may have measured a particular system and what we would like to design is a compensator for it. And the question is, by putting this in a feedback loop, can we, in fact, implement the inverse of this system? The answer to that is

yes. And the feedback system, in that case, is as I indicate here.

So here what we choose to do is to put the system whose inverse we're trying to generate in the feedback loop. And in this case, a high-gain in the forward path. Now for this situation, k is, again, a constant. But in fact, it's a high-gain constant.

And now if we look at the feedback equation, then what we see is an equation of this form. And notice that if k times P of s is large compared with one, then this term dominates. The gain in the forward path cancels out. And what we're left with is a system function, which is just 1 over P of s.

And a system of this type is used in a whole variety of contexts. One very common one is in building what are called logarithmic devices or logarithmic amplifiers. Ones in which the input-output characteristic is logarithmic. It's common to do that with a diode that has an exponential characteristic. And using that with feedback-- as feedback around a high-gain operational amplifier.

And by the way, the logarithmic amplifier is nonlinear. What I've said here is linear, or the analysis here was linear. But that example, in fact, suggests something which is true, which is that same basic idea, in fact, can be used often in the context of nonlinear feedback and nonlinear feedback systems.

Well, as a final example, what I'd like to analyze is the context in which we would consider stabilizing unstable systems. And I had indicated that one context in which that arises and which we will be analyzing in the next lecture is the inverted pendulum. And in that situation, or in a situation where we're attempting to stabilize an unstable system, we have now in the forward path a system which is unstable. And in the feedback path, we've put an appropriate system so that the overall system, in fact, is stable.

Now, how can stability arise out of having an initially unstable system? Well, again, if we look at the basic feedback equation, the overall system function is the system function for the forward path divided by 1 plus the loop gain, G of s times H of s. And for stability what we want to examine are the roots of 1 plus G of s times H of s.

And in particular, the poles are the zeroes of that factor. And as long as we choose G of s, so that the poles of this term-- I'm sorry, so that the zeroes of this term are in the left half of the s-plane, then what we'll end up with is stability.

So stability is dependent not just on h of s for the closed-loop system, but on 1 plus G of s times H of s. And this kind of notion is used in lots of situations. I indicated the inverted pendulum. Another very common example is in some very high-performance aircraft where the basic aircraft system is an unstable system. But in fact, it's stabilized by putting the right kind of feedback dynamics around it. And those feedback dynamics might, in fact, involve the pilot as well.

Now, for the system that we just talked about, the stability was described in terms of a continuous-time system. And the stability condition that we end up with, of course, relates to the zeroes of this denominator term. And we require for stability that the real parts of the associated roots be in the left half of the s-plane.

Exactly the same kind of analysis, in terms of stability, applies in discrete-time. That is, in discrete-time, as we saw previously, the basic discrete-time feedback system is exactly the same, except that the independent variable is now an integer variable rather than a continuous variable. The feedback equation is exactly the same.

So to analyze stability of the feedback system, we would want to look at the zeroes of 1 plus G of z times H of z. So again, it's those zeroes that affect stability. And the principal difference between the continuous-time and discrete-time cases is the fact that the stability condition in discrete-time is different than it is in continuous-time. Namely, in continuous-time, we care for stability about poles of the overall system being in the left half of the s-plane or the right half of the s-plane.

In discrete-time, what we care about is whether the poles are inside or outside the unit circle. So in the discrete-time case, what we would impose for stability is that the zeroes have a magnitude which is less than 1. So the basic analysis is the same, but the details of the stability condition, of course, are different.

Now, what I've just indicated is that feedback can be used to stabilize an unstable

system. And as you can imagine there's the other side of the coin. Namely, if you start with a stable system and put feedback around it, if you're not careful what can happen, in fact, is that you can destabilize the system. So there's always the potential hazard, unless it's something you want to have happen, that feedback around what used to be a stable system now generates a system which is unstable.

And there are lots of examples of that. One very common example is in audio systems. And this is probably an example that you're somewhat familiar with. Basically, an audio system, if you have the speaker and the microphone in any kind of proximity to each other is, in fact, a feedback system.

Well, first of all, the audio input to the microphone consists of the external audio inputs, and the external audio inputs might, for example, be my voice. It might be the room noise. And in fact, as we'll illustrate shortly if I'm not careful, might in fact be the output from a speaker, which represents feedback.

That audio, of course, after appropriate amplification drives a speaker. And if, in fact, the speaker is, let's say has any proximity to the microphone, then there can be a certain amount of the output of the speaker that feeds back around and is fed back into the microphone.

Now, the system function associated with the feedback I indicate here as a constant times e to the minus s times capital T. The e to the minus s times capital T represents the fact that there is, in general, some delay between the time delay between the speaker output and the input that it generates to the microphone. The reason for that delay of course, being that there may be some distance between the speaker and the microphone. And then the constant K2 that I have in the feedback path represents the fact that between the speaker and the microphone, there may be some attenuation.

So if I have, for example, a speaker as I happen to have here, and I were to have that speaker putting out what in fact I'm putting into the microphone, or the output of the microphone, then what we have is a feedback path. And the feedback path is from the microphone, through the speaker, out of the speaker, back into the

microphone. And the feedback path is from here to the microphone. And the characteristics or frequency response or system function is associated with the characteristics of propagation or transmission.

If I were to move closer to the speaker and I, by the way, don't have the speaker on right now. And I'm sure you all understand why. If I move closer, then the constant K2 gets what? Gets larger. And if I move further away the constant K2 gets smaller.

Well, let's look at an analysis of this and see what it is, or why it is, that in fact we get an instability in terms-- or that an instability is predicted by the basic feedback equation.

Now, notice first of all, that we're talking about positive feedback here. And just simply substituting the appropriate system functions into our basic feedback equation, we have an equation that says that the overall system function is given by the forward gain, which is the gain of the amplifier between the microphone and the speaker, divided by 1 minus-- and the minus because we have positive feedback-- the overall loop gain, which is K1, K2, e to the minus s capital T. And these two gains, K1 and K2 are assumed to be positive, and generally are positive.

So in order for us to-- well, if we want to look at the poles of the system, then we want to look at the zeroes of this denominator. And the zeroes of this denominator occur at values of s such that e to the minus s capital T is equal to 1 over K1 times K2. And equivalently that says that the poles of the closed loop system occur at 1 over capital T, and capital T is related to the time delay. 1 over capital T times the log to the base e of K1 times K2.

Well, for stability we want these poles to all be in the left half of the s-plane. And what that means then is that for stability what we require is that K1 times K2 be less than 1. In other words, we require that the overall loop gain be less-- the magnitude of the loop gain be less than 1. If it's not, then what we generate is an instability.

And just to illustrate that, let's turn the speaker on. And what we'll demonstrate is feedback. Right now the system is stable. And I'm being careful to keep my distance

from the speaker. As I get closer, K2 will increase. And as K2 increases, eventually the poles will move into the right half of the s-plane, or they'll try to. What will happen is that the system will start to oscillate and go into nonlinear distortion.

So as I get closer, you can hear that we get feedback, we get oscillation. And I guess neither you nor I can take too much of that. But you can see that what's happening-- if we can just turn the speaker off now. You can see that what's happening is that as K2 increases, the poles are moving on to the j omega axis, the system starts to oscillate. They won't actually move into the right half plane because there are nonlinearities that inherently control the system.

OK, so what we've seen in today's lecture is the basic analysis equation and a few of the applications. And one application, or one both application and hazard that we've talked about, is the application in which we may stabilize unstable systems. Or if we're not careful, destabilize stable systems. As I've indicated at several times during the lecture, one common example of an unstable system which feedback can be used to stabilize is the inverted pendulum, which I've referred to several times.

And in the next lecture, what I'd like to do is focus in on a more detailed analysis of this. And what we'll see, in fact, is that the feedback dynamics, the form of the feedback dynamics are important with regard to whether you can and can't stabilize the system.

Interestingly enough, for this particular system, as we'll see in the next lecture, if you simply try to measure the angle and feed that back that, in fact, you can't stabilize the system. What it requires is not only the angle, but some information about the rate of change of angle. But we'll see that in much more detail in the next lecture. Thank you.