DENNIS FREEMAN: Hello and welcome. As before, as you might be expecting, the big new news is, of course-- well it's not new news. The big news is exam, tomorrow night. This should be very familiar stuff to you.

I apologize, a couple people asked me right before lecture about Fourier because there was Fourier on one of the previous exams. There is no Fourier. I should have probably axed that problem. So in my head I was trying to be complete disclosure. This is what the exam looked like. I didn't mean to send the false message that there is Fourier. There is no Fourier. Next time we'll have lots of opportunities to ask Fourier questions, so no Fourier this time. On the other hand, yes, Bode, Root Locus, that kind of stuff, yes. So, that kind of stuff you should get on top of. Questions or comments about the exam? I guess the one thing to remember is it's in Walker. It's not on the third floor of 26. Other than that, you should be kind of in a good position to know roughly what to expect. Yes?

AUDIENCE: When you say two pages of notes, do you mean two separate 8 and a half by 11 cards?

DENNIS FREEMAN: Yes. So the idea is if you made one 8 and a half by 11 front and back for the last exam, please just bring that. Don't bother making a new one. But you might want to make a new one that contains stuff that we've done since then. That's the idea. But we of course don't care what you do. Micro, you know, if you want to bring-- so, no electronics. But optical magnification is fine. If you want to bring a little microscope, it's fine. We don't mind. If you can have an all optics, no electronic, microfiche reader, fine. So the idea is just to send the message that you shouldn't be spending a lot of time leafing through 300 pages. That's what we don't want you to do.

AUDIENCE: Question.

DENNIS FREEMAN: Yes

AUDIENCE: Can we bring our abacus, or is that not allowed?

DENNIS FREEMAN: Can you bring your?

AUDIENCE: Abacus.

DENNIS FREEMAN: Abacuses are welcome. In fact, I would love to race somebody with an abacus. I will take you on with a slide rule.

[LAUGHTER]

Okay, so I learned with a slide rule and the question was-- Well, no, you don't want to know that. I don't want to tell you that, actually. Okay, other questions, comments? Abacuses, fine. Non electronic devices.

So today is the final lecture on feedback. What we've done over the course of the last week and a half or so is looked at a variety of applications. Last time we looked at how to perform-- how to improve performance with feedback. We looked at how you could increase the bandwidth of an op amp, increase the speed of an op amp, change a motor so that its natural output, which is a voltage controlled velocity, could be perturbed into a voltage controlled position.

Today I want to just do a few more examples. Because what I want you to do is leave this place thinking feedback's really useful. It's not just torture. I mean it is torture. We do torture you and you're supposed to figure out how it works quantitatively and all that. But it's not just torture. It's actually very useful. And that's what I want you to take away. The first application today that I want to think about is reducing unwanted parameter variation. So the idea that I want you to think about is what do you have in a power amplifier of the type that you would drive speakers-- of the type that you would use to drive speakers. That's a kind of a complicated electronic device. It's kind of hard to make high power. You end up having to generate high power in order to make a loud speaker that fills this hall for example. And if you were to try to do something bigger like a concert, it's even harder.

And the problem with those kinds of devices is that they are active. They generate signals that are bigger that generate more power at their output than there was available at the input. And we're very good at such things. Unfortunately, such things are intrinsically non-linear. So we actually take advantage of non-linearities in materials in order to do that. The result though is that at a high level the thing's still non-linear. And we can see some of that. And not only non-linear, but other things. They can be temperature dependent.

So if you were to make a high power amplifier it's very likely that the performance would depend on temperature. Who cares, it's always room temperature. No, not inside the amplifier. It may start out at room temperature every time you turn it on, but within a few minutes it's going to be much higher than room temperature. And if you didn't do something about it, the gain would change. If you didn't do something about it in your car, it would track the weather, which may be good. Maybe it would provide a distraction and provide utility to winter. But short of that normally we don't want the sound level that comes out of a stereo system to depend on temperature.

So how do you take a device that intrinsically depends on temperature? Here I am imagining that I have a power amplifier whose gain can go between 8 and 12, say, depending on what the temperature is. How could you fix that so that its behavior was more consistent from day to day, from hour to hour. And of course, not too surprisingly, the answer is use feedback. So how do you do that? So you want to connect the amplifier to the speaker. After all, it's inside the amplifier that you have all the transistors that are able to drive the speaker. So you don't want to perturb that connection. But what you can do with feedback is monitor what was the voltage that actually got produced and try to rig the feedback system to coerce it to the voltage you wish it were.

So the way you can do that is to generate a feedback loop. Precede the power amplifier with a voltage amplifier that does not necessarily significantly increase the power. That makes that kind of an amplifier somewhat easier to make. So imagine that this power amplifier has big chunky transistors because it has to generate amps in order to make your speakers work. This, by comparison, could be a much less power hungry kind of amplifier. And it's easier to make, and in particular would be easy to make it with a gain of 10 or 100 or 1,000 or 10,000. We're very good at doing that kind of stuff.

Then the idea of feedback is to compare the signal that you generated to the signal that you might have desired. So here's the signal you generated. Here is the signal you desired. I'm comparing it through a monitor that's feeding back beta times the output voltage. The reason I do that can be seen if you just work through the math. So if I just write the closed loop transfer function for x goes to y, the easy way to think about that's Black's formula, of course, right? So think about the forward path, k times f naught divided by 1 plus the loop gain. So 1 plus the gain around the loop is beta k f naught. And now just think about what would happen if I coerced it so that I made k the thing that's under my control big. If I make k big enough, then I can arrange so that beta k f naught, independent of beta and k, I can always choose k so that the product is big compared to the one. If I do that, I can ignore the one. If I ignore the one, the k's cancel, the f naughts cancel, and I'm left with whatever beta.

Beta didn't depend on f naught. That's what I wanted. It didn't even depend on k, other than k had to be big enough so that the division worked. So the result is that you can make a stable gain by incorporating an amplifier with an unstable gain. Unstable meaning that it varies in time, not that it's unstable in the sense that we often use the word unstable. I should avoid that. The parameter of the gain changes. It changes between 8 and 12. And I can illustrate that with this kind of a picture. If I plot the variation at f naught.

I've showed a range here from 0 to 20. I'm imagining that in a practical situation it could vary between 8 and 12. And then in red, I'm showing the gain that would result if there were no feedback. And as you can see that over the range of 8 to 12, the output gain, the total gain of the system, varies between 8 and 12, not surprisingly. That's what it would mean to have no feedback.

But if I crank in feedback, if I set the k to 100 and feedback ratio to 1/10, the 100 is sufficient to make the one negligible which means that my overall gain is 1 over 10 determined only by the feedback ratio, beta. And so the result is this line shows the solution to that equation. And you can see that it's ruler straight through the interesting region. Does that makes sense? So what I've been able to do is use feedback in order to control the value of a parameter that was otherwise variable, or sensitive to something, often temperature. Is that all clear? That kind of a thing--

Question? I got rid of all the variations, except beta. What about variations in beta? Right, the premise was I couldn't control this, it was going to vary between 8 and 12. But what if beta changes. What if beta went from 1/8th to 1/12th, would that affect things?

AUDIENCE: So yeah, so we could put a feedback loop [INAUDIBLE].

DENNIS FREEMAN: You could put in another feedback loop. That doesn't sound like it's going to converge quickly. Yes?

AUDIENCE: Typically we want to be able to control beta anyway.

DENNIS FREEMAN: Right

AUDIENCE: There's nothing users control.

DENNIS FREEMAN: That's right, we'd want to control it. Is there any reason we-- but that's what we wanted to do with f naught too. We wanted to control that. We'd like that to be 10. Yes.

AUDIENCE: It will depend on how that beta would change. Changes, say, from 1/8th to 1/12th, probably nothing to worry about. If it changes from 1/8th to 1 over 8,000, then you may want to think about what k and f are.

DENNIS FREEMAN: So what do you think will happen? Do you think it will vary a lot or a little compared to 8 to 12?

AUDIENCE: Probably not that much.

DENNIS FREEMAN: Why?

AUDIENCE: Because the k is already-- 1/8th and 1/12th is not really much of a difference in terms of--

DENNIS FREEMAN: So why do I think about beta differently from the way I think about f naught? Is there a good reason for that or is that just wishful thinking, right? Engineering by wishful thinking. I wish I could get a beta that was exactly right. Yes?

AUDIENCE: For amplification, beta needs to be less than 1 [INAUDIBLE] so [INAUDIBLE] the resistor, it's not going to have a very big temperature [INAUDIBLE]

DENNIS FREEMAN: Precisely. So the problem is getting gain in a stable and reliable fashion, not in getting attenuation. We can control the value of a resistor with great precision. And so it's easy, by comparison, it's relatively easy to control the gain of a resistor network, and much harder to control the gain of a transistor network. Transistors have gain. Gain is hard to make in a non temperature dependent fashion. Division by a resistor, it's not hard to make a resistor. I mean it's not trivial, but it's not hard. Even in the dark ages-- even in the dark ages, resistors were made by winding wires, and they were very stable with temperature or could be made to be very stable with temperatures.

So one of the tricks in feedback is that we usually-- so we move the-- we move the parameter dependence out of the forward path and into the feedback path with the idea that we're moving it out of the part that has active parts like transistors and into the part that we can build with passive parts. And we almost always do feedback that way. Not always, but almost always.

Another important point is that this could work at a very high speed. As an illustration of that I want to think about what would be inside of that f naught box. Here's a very simple idea of what you could put in there that's based on two power transistors. This is not a course in power transistors. I am not going to give you a very good model for a transistor. There's other courses that do much better justice to that topic. But what I do want to say is that the idea would be that turning those transistors on and off represents a significant source of distortion, or potentially a source of distortion.

So for example, the simplest possible model I could think about would say that in order to generate, in order to push this voltage to a higher number, what I want to do is turn this transistor on because this will pull this voltage up toward the positive rail. So if the input is such that I'd like this to go up what I would like to have happen is that this transistor would come on and conduct current this way, which would suck this voltage up to there. To do that, the voltage here has to exceed the voltage there.

For a BJT, that's not the only kind of transistor I'm just illustrating one particular kind of example. For a bipolar transistor of this type, the way you would turn that transistor on would be to have this voltage exceed this voltage by at least some threshold voltage, nominally 0.6 volts, which means that the kind of relationship between v out and v in that you would get is illustrated here. If the input voltage differed from the output voltage by less than plus or minus the threshold voltage, then nothing would happen because neither transistor would be on. If the input exceeded what was desired for the output by more than 0.6, the transistor on the top would start to come on, and you would get roughly linear growth in the output voltage in relation to the input voltage.

Similarly with negative. If you make the voltage at the input sufficiently negative, you can cause the bottom transistor to come on and suck the voltage down toward minus 50. But again, you have to exceed the difference. The difference has to exceed something like the threshold voltage, normally about 0.6 volts, which means then that you get this piecewise linear approximation to the relation between the input and the output voltage.

The problem with that is that it distorts the waveform. So if I imagine a purely sinusoidal input, what I get at the output then is not a purely sinusoidal output. And, in fact, that's quite objectionable for many purposes. In particular, in the example where we're listening to a sound, if we were to listen to this sound it would sound very different from the way a sine wave would sound. So the way I can do that is, not surprisingly, fix it with feedback. Again imagine that this is the complicated part. v f naught, followed by the speaker, put some gain in front of it, gain's cheap. We can do that run relatively well. And put in feedback so that I can directly compare how is the speaker voltage compare to the voltage that I'd like to listen to.

The result then is that, depending on how big is the gain, I can change how big is this dead zone. The zone of input voltages that produce no change in output. So if I turn up the gain, here I am illustrating k equals one. If I turn it up to k equals two the dead zone gets smaller-- k equals four, eight, 10, whatever. So the idea is that by adjusting the gain, I control how much feedback there is. And I can control how distorted the wave form looks. And I can illustrate how bad it is to not do that by comparing an original clip of music to different values of the feedback. So what I'm going to do is play this many clips, some original-- the original that came off the mp3. Put it through the system with no feedback, that is to say there isn't a k here. There's no feedback loop. You just put v in a v out. Then k of two, four, eight, 16, and then back to the original. And if you listen, you should be able to hear the effects of the distortion. OK. Ready?

[MUSIC PLAYING]

You hear?

[MUSIC PLAYING]

OK. Everybody get to hear it? So the point-- although the acoustics in this room are terrible-- but even somebody with as old ears as mine, I can hear a very big difference even between k equals 16 and the original if I listen to it over headphones, or if I listen to it in a better environment. So the point is that we're very sensitive to it. In fact, 16 isn't even close. I mean, in a real kind of an amplifier you would use gains of at least 1,000 in order to make it acceptable. So the point is that we can use feedback. So the overall point is to try to illustrate the kinds of things we can do with feedback. Here what we're doing is making a system that's intrinsically non-linear more linear. Questions? Comments? Left turn approaching.

So what I want to close with then is two more examples that are much more fun. They're examples of the use of feedback to stabilize unstable systems. This is kind of connecting back to the very first lecture in this series that Russ gave where he did some very complicated control problems by doing feedback. I'm going to do some slightly simpler ones. But they're hard enough to be clear that they were a hard problem, and easy enough that every one of you should be able to write the equations and figure out what I'm about to do. If you are interested to do the more complicated things like having owls come in and swoop, that's a research project. So this is supposed to be in the middle ground between research and trivial, right? So I'm going to do some examples now of some unstable systems to illustrate how you would use feedback for stability.

First example, is magnetic levitation. You've probably heard of magnetic levitation. Magnetic levitation is hard. We've talked about magnetic levitation since the 70s at least, perhaps the 60s. It's hard because magnetic systems tend to be intrinsically unstable. Not always, but the easy ways you can imagine using a magnet to levitate things. We'd all like to make anti-grav right? Anti-grav would be a really good thing. And it seems like you ought to be able to do that with a magnet, right? So the idea is what I would like to do is levitate a mass with a magnet. Anti-grav, right? We've had that since Star Trek, right? It's time somebody built it. But it's hard even though we can make magnets. Even though we can make things that seemed like they would do it. It's hard because it's intrinsically unstable.

This particular system is intrinsically unstable in the sense that if the mass, let's say that it was at equilibrium by which I mean, in the absence of somebody touching it, it would stay exactly the way it is. That's what I mean by equilibrium. Let's say it was at equilibrium so that the magnetic force I found was just sufficient to hold up the weight, Mg. If I came over and just very gently bumped the mass up a little, what would happen to the force? Force will go up. So you're supposed to remember this from some subject that you took before, right? Yes, no, yes? It started with an 8, yes?

So the idea is something to do with if you have a magnet of that type, you generate magnetic fields. You remember they're not supposed to cross, right? One of the rules is they don't cross.

[LAUGHTER]

But the idea is that the density of flux lines decreases as you leave the magnet, and the force that's generated has something to do with the change in the density of flux lines. If you move an object through a region where the number of flux lines doesn't change, it takes no force. If you move an object into a region that has fewer or more flux lines, it requires force to make that move. That's the idea. So, since the flux lines are thinning as you move away, if I push up it will generate-- I'm moving the mass into a place where there is a greater density of flux lines, that will generate force up on the mass. That's bad, right?

I had the force exactly where I needed it to be in order to levitate the mass. If the force goes up, what happens? It accelerates up. Well, OK, so it accelerates up, then what? Well, it gets closer. Then what? It accelerates faster. So, the problem is that the position is unstable, or we might say quasi stable. So what we can do is make a model of this to try to understand how we could fix it with feedback. Imagine that I represent the magnet by a function that maps for every value of the position of the mass why it generates a force, f m. Then, because of the shape of the field lines, I'm expecting that the force increases on the right of the curve and decreases on the left of the curve. If I'm sitting for y zero at equilibrium what matters is the difference between y and the equilibrium place.

And the important force is the difference between the magnetic force and Mg, the weight. So I get a simple relationship that says that this force, this vertical line, f, should be m a. F equals m a which is m y double dot. So having generated that kind of a model I can now imagine putting that into a feedback system. So now I put the model into the feedback system and this junk over here, this is f equals m a. f gets multiplied by one over m, that a. a is y double dot. Integrate it to get y dot. Integrate it again to get y. But y was the thing that controlled the magnet in the first place. So I'm trying to cast magnetic levitation into a frame for understanding feedback.

So now if I want to further simplify things I can replace the magnet model, at least for small deviations in position, by a linear gain. So I'm making a straight line approximation to this curvy line. And then we have something that we analyze all the time, simple feedback system that has integrators, gains, that kind of stuff. It's surprisingly similar to the mass spring dashpot. Springs, by contrast to magnets, springs levitate stably. What happens if I push-- so I'm suspending a mass from a spring, if I push up on the mass how does the force in the spring change? Gets less. That's what you want. It's no longer sufficient to hold the mass up, so the mass falls. Push up, the mass tries to fall. Pull down on the mass, it tries to pull up. Spring levitates stably. Magnet levitates unstably. OK?

Almost the same system though. Here's a way of thinking about the magnetic levitation like we just saw. Here's an equivalent diagram for the spring levitation. Almost the same thing. The critical thing is that minus sign and that plus sign. So what's the relationship between the poles of the spring and mass system and the poles of the levitation system? Look at your neighbor. Say the word poles. We say that a lot in this class. How do the poles relate in those two systems?

[AUDIENCE CHATTER]

OK, so draw the poles in the air for the two systems. How do they relate? Excellent. Excellent. I saw one right answer. Somebody else draw some-- I'm not sure if I caught that one. So how do the poles relate? Come on, somebody drawn it in the air. I've got one right answer. OK, your air drawings are a little cryptic. Here's a not so air drawing. So here's a picture of what they look like for a spring and mass system. And all you do is think about the differential equation, turn it into system function, there's two poles, factor the denominator, you find that the two poles are on the j omega axis. j omega axis, that sounds bad. Yeah, it is bad, right? what happens if they're on the j omega axis? Oscillate. If they're right on the j omega axis, it oscillates forever. That's exactly what a mass spring dashpot does if you have no damping. OK, it's bad. But it's not nearly as bad as magnetic levitation.

Magnetic levitation, you get the same distance from the origin, but it's turned to plus and minus something. That one. I like. That one I don't like. That one is unstable. So the question is, how can you make a magnetic levitation system? How do you make a system that we use a magnet to levitate a mass when the physics are intrinsically stable? And the answer is-- feedback, of course, yes.

So the idea is that here is our system that was intrinsically unstable. What we'd like to do is somehow change the physics. One of the ways we can change the physics is by modifying the current that runs through the magnet. If we changed the current from i naught to 1.1 times i naught, then the forces everywhere would be bigger. We could use that idea to introduce a second feedback path. There is the intrinsic one that comes just from the physical layout. But I'm going to add a path that senses the position of the mass and feeds back a change in current. And if I do this carefully, I can arrange this path, the intrinsic one has positive feedback that was bad, but the one I add has negative feedback, which is good.

Then the overall system looks more like this where I've got k, which I cannot affect, that's physics. But I've got k two that I can affect. And what's showed here is a root locus. As I change k two, what happens to the poles? And what you can see is since they were equally disposed on the two sides of the j omega axis, as I turned up the gain-- you can do this just by doing-- factor the polynomial. Right? So I factored this to generate this plot. As you increase k two the poles go toward each other and then they split, and they become complex. Complex is good. But this still has the problem that the mass spring dashpot had if it didn't have any damping. So it's better because at least it's stable-- Oh well no it's not stable. It's marginally stable. It's now oscillatory. It no longer has eigenfunctions that grow with time. The problem is they don't shrink either. They state constant in amplitude.

But you can fix that, and what I'll do now is show a demo that was done by Professor James Roberge where the idea was you don't have to just put the extra gain in. What Jim did was put a zero too. So these are the intrinsic poles from the magnetic levitation. This is a zero introduced by electronics. The zero has the effect of bending the poles, the closed loop poles around, so that you can then make it stable. So [INAUDIBLE].

So here's Jim's device. And the goal now is to smash fingers because it will try to levitate. So this is a hunky magnet. This is a hunky power supply. And that is-- well this is the power supply. This generates DC voltage that drives the amplifier that's inside this box. So the idea then is we have this metal ball. So we have this metal ball, and it's kind of heavy, right? You know, it's not like some nuclear-- it's not uranium or anything, but it's heavy. And the idea is going to be to use the hunky magnet. Let's see, can we see? Oh yeah, OK, so you can see it up there. So you can see maybe that here is the pole of the magnet, which is right there. So there is the pole, a different kind of pole.

[LAUGHTER]

Not to be confused.

[LAUGHTER]

So that's the iron core of the magnet. And you can see the turns up here. And here you can see some light. So what's going on here is a source of light is being aimed toward a photo sensor so that when the mass, the ball, is levitating at exactly the right place, there's a signal that gets fed back into the control box to do the thing that I have illustrated here. OK, everybody knows what's going on? I'll smash my fingers now and then ask Lorenzo to do it correctly. So what I'm going to do now is try to move the mass into the air so it levitates. so successful would be you're watching up there and there's going to be a ball up here, and it will levitate. Unsuccessful is I'll have a black and white-- a black and blue fingernail.

So, let's see-- Lorenzo is laughing at me.

[LAUGHTER]

LORENZO: See your hand can move. OK.

DENNIS FREEMAN: Nope, that's not going to work.

LORENZO: You don't have a way to move?

DENNIS FREEMAN: Yes.

LORENZO: Yes.

DENNIS FREEMAN: OK, this will work better. If I get coached enough, I can do things.

LORENZO: Yeah! Oh!

DENNIS FREEMAN: Almost!

[LAUGHTER]

OK, it was close. No, no, no. Failing once is fine. Failing twice, no that's not acceptable. Wonderful. OK, so that's magnetic levitation made possible by feedback. That's the point, OK? Everybody sort of knows what's going on?

OK, so to redeem myself we'll leave the demo set up and after lecture you can come down and you can do it and smash your fingers. Questions, comments, everybody knows the point? So we'll leave it there just for entertainment value.

LORENZO: It will get hot.

DENNIS FREEMAN: Oh, OK. OK. You don't want to burn it up, huh? Good. Thank you. Thank you.

[APPLAUSE]

So the next demo is the famous inverted pendulum. The idea is that if you have a pendulum-- Yes?

AUDIENCE: Was k oscillating also? With the--

DENNIS FREEMAN: Correct. Correct. So, that's a good question. OK, so did I screw up? Not a chance, right? Or for those of you who know me-- what would you conclude if it was oscillating and I told you I thought I built this? I screw up the circuit. I screwed up the root locus diagram. Everything's hunky dory, and I did everything right. None of the above. Yes?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Positions. OK, so flaky hardware. Universal solution right? Could be. Could be. Could be. There's a different explanation, which is possible, but your explanation may be right. What's another explanation besides flaky hardware? Yes?

AUDIENCE: [INAUDIBLE] complex poles.

DENNIS FREEMAN: Complex poles, but didn't I stabilize them? Leading question. Didn't I stabilize them? Leading question. The answer is, didn't I stabilize them? Yes? No? Completely cryptic he is.

AUDIENCE: I mean they can still have complex values.

DENNIS FREEMAN: They could have complex-- how?

AUDIENCE: [INAUDIBLE] don't know where they are.

DENNIS FREEMAN: And so what depends-- what determines where they are?

AUDIENCE: Sorry?

DENNIS FREEMAN: What determines what they are?

AUDIENCE: Whatever your-- whatever you set [INAUDIBLE] k or gain.

DENNIS FREEMAN: k

AUDIENCE: Yeah.

DENNIS FREEMAN: Exactly. So I get to control this k two thing. And as I change k two, the poles start out here, they go toward each other. So if k equals zero they're here. As k gets bigger they do this. So it's possible that they're just simply here. Right? The poles could still be complex. If that were true, you would expect the oscillations to die with time. Right? If the oscillations didn't die with time, it's still possible-- well, if they don't die with time, then I guess I screwed something up. Right? Because even if they just come here, they're oscillating but they're on the left part of the axis so they should decay eventually.

So I guess to the extent that they didn't die eventually, something screwed up with the analysis, the hardware is screwed up, something. Yes?

AUDIENCE: So is root locus due to k or k two.

DENNIS FREEMAN: k two. You can draw a root locus for anything. You can change any parameter and sketch the locus of points where the roots move to, hence the name, root locus. This is for k two. So z zero was fixed here, held constant, and k two was modified to trace out that loop.

The next example is pendulum. So just like spring mass system, spring levitated mass in a stable fashion, although perhaps oscillatory. So this pendulum is also stable in that same sense. Might be oscillatory, but it will eventually stop. Of course pendula are not very interesting. What's interesting is inverted pendula where the rule is what I have to do is stabilize this by moving my hand to compensate. So let's say that it starts to move this way. I've got positive feedback that tends to make it continue. So if I started at equilibrium, by which I mean a state that would persist if nothing changed. So if I started at equilibrium, but then perturbed it by breathing, or whatever, or by just not being quite at equilibrium, it'll move and then I have to do something to compensate, otherwise it'll fall down. And there are people who can do this better than me.

But we don't need to have-- I don't need to be coordinated. I can build a system that's coordinated instead. That's the value of feedback. So I can be uncoordinated. Before I do that-- no, let's do this first. So, here is a mechanical system that uses feedback to stabilize-- I'm not going to read anything right? So here is a pendulum, an inverted pendulum. There's some stuff down here. There's no active parts. This stuff is a potentiometer, just like the potentiometer on the back of the motor that I demonstrated. And what it's doing is providing a signal to electronics that will let the electronics say, oh, it's moving this way.

Now, the cart is on wheels, and the wheels are connected to a hunky motor. If the weight starts to move this way what should the hunky motor do? Push the cart that way. So that's the idea. So I've got a one dimensional inverted pendulum and I'm explicitly modeling the idea that I'm-- that the base is moving. So over here is an outline of the way you think about that-- the way you can think about that. I got the car and it's in a not invert-- the car as supporting a pendulum. The pendulum is in a non inertial frame because the mass-- let's say that I'm approximating the pendulum as having all of its mass at the end. That's an approximation, but it's not a critical approximation.

So I can imagine now that I've got an input variable, which is x, the position of the car. And I've got a feedback signal that's trying to feedback to me how much is theta, the angle that the pendulum is making. And so what I want you to do now-- so if you just write 801 kinds of stuff you can imagine that the if-- So, first off, you've got-- instead of f equals m a, you have torque equals i omega.

Everybody's happy about torque equals i omega, right? And, if I want to be in a non inertial frame, one way to account for the fact that the car is moving is to compute torques about the point of insertion. Then I can represent the effect of the moving car as a pseudo force on the mass. That pseudo force is in proportion to the acceleration of the car. And I've got the weight of the inverted pendulum trying to pull it in the other direction. And if you think about all this stuff you get this kind of an equation. So this is i omega dot, i omega dot equals, and then I have to worry about torques. There's the torque that's due to the weight, and there's the torque that's due to the cart movement.

so the question then is, what are the poles of this system? We want to think about all of these kinds of systems using 003, because it's such a powerful thing. So what are the poles of the inverted pendulum? Look at your neighbor. Say the word pole.

[AUDIENCE CHATTER]

Where do the poles--

[AUDIENCE CHATTER]

Bump the car.

[AUDIENCE CHATTER]

OK, so where's the poles? Sketch them in the air. No that didn't work at all. So to think about where are the poles you can think of-- the hard part is just the sinusoidal trigonometric non-linearity, so what do I do with that?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Throw it away, absolutely. So I'll assume that the sine of theta is--

AUDIENCE: Theta.

DENNIS FREEMAN: Theta. And the cos of theta is?

AUDIENCE: One.

DENNIS FREEMAN: One. For small angular deflections, just make a small angle approximation. When you do that it comes out surprisingly to look just like the levitated mass problem. You get real poles on each side of the j omega axis. So in fact, you have to do a similar kind of stability thing to make this work. And so that's what we'll demonstrate now. So the idea is going to be to turn on electronics to stabilize the inverted pendulum so that I don't need to do it with my hand.

LORENZO: I'm not sure it's going to work, because I [INAUDIBLE] the cord. So if [INAUDIBLE]

DENNIS FREEMAN: It always works. He does it really well.

LORENZO: That works.

DENNIS FREEMAN: So it's a little more stable than me, but not much. I mean, I'm pretty good, right?

[LAUGHTER]

DENNIS FREEMAN: I can keep it going too, right? Well, sort of. So is it any better than I am? So can you share that it's better or worse than me.

LORENZO: It's better than you.

DENNIS FREEMAN: It's better than me? Prove it. Go ahead.

LORENZO: The way [INAUDIBLE].

DENNIS FREEMAN: So you ought to do this. So this guy. So can somebody think of something that's different about this from the magnetic levitation problem?

AUDIENCE: Like right now it looks like it's going to converge.

DENNIS FREEMAN: But it's doing a really good thing. It's converging around here. What would happen if I was doing my thing and it started to fall this way, so I started to follow it. That's actually very complicated. So there's actually another feedback loop inside here that's trying to make it stay in the center. So if I tried to pull that one out. Wap!

OK, so there's many feedback loops going on here. There's the one that you can see. But then there's also a slow feedback loop trying to keep the whole thing oscillating in the center. Everybody follow that? So there's the primary feedback loop, which is very analogous to that guy. But then there's a slow feedback loop saying wait a minute, you're spending too much time to the left. I'm going to move it a little bit in the wrong way so that you migrate toward the right. But the result of that is very stable. How stable?

LORENZO: We'll find out.

DENNIS FREEMAN: I didn't put a battery on the top of mine.

[LAUGHTER]

I could have! Well, I could've put it there. So now we put the--

[LAUGHTER]

AUDIENCE: I think someone should sell this.

DENNIS FREEMAN: A volunteer!

[LAUGHTER]

Everybody in favor? It's overwhelming. OK. So that's all for today. The idea is that we can use feedback to do a whole bunch of things. And maybe even more tomorrow. Oh well now that's just a joke. See you tomorrow at the exam. No recitation. We'll stick around for two or three minutes. You can look at these more if you want to.