


Robustness analysis

Christophe Salzmann & Colin Jones

Automatic Control Lab



Understand Nyquist criterion, robustness and margins

- Gain margin (critical frequency)
- Phase margin (crossover frequency)
- Delay margin

Build intuition for robustness via experimentation

- Critical gain
- Model mismatch

TP Control Systems

Hello and welcome to your hands-on session focused on robustness analysis. Our objective today is to understand the Nyquist criterion, and through this, to get a feeling for what robustness means. Of course, we also want to understand the margin, which is the measure of what it means to be robust or non-robust. There are three margins we're going to look at today: the gain margin, the phase margin, and the delay margin. We then want to build an intuition of what robustness means by interacting with the experimental platform. Here, we're going to look at a number of different points, specifically at critical gain and at the impact of model mismatch on the controller performance.



Notes

Summary



0m 05s

Simplified Nyquist criterion → Stable ?



Simplified Nyquist criterion

If open-loop system is **stable** (poles of $K(z)H(z)$ inside the unit circle) & -1 point lies to the left of the Nyquist curve, then the closed-loop system is stable

Example

Electrical drive in position mode with $K(z) = 1$
(proportional controller)

$$H(z) = \frac{0.0024z + 0.0023}{z^2 - 1.8521z + 0.8521}$$

TP Control Systems


The main tool that we're going to use today is the simplified Nyquist criterion. This tells us that if the open-loop system is stable, named at the poles of K times H or open-loop gain are all inside the unit circle, and the Nyquist point, the -1 point lies completely to the left of the Nyquist curve, then we know that our closed-loop system is going to be stable. Over here, we can see that this point here is our -1 Nyquist point. If this is our Nyquist plot that we have over here, we can see from our example that in fact, the system is stable, in open-loop stable. We can tell immediately that this system is going to be closed-loop stable once we actually close the loop.

Notes

Summary



0m 40s



Uncertainties

Due to model mismatch such as gain or delay mismatch, noise, etc.

Robustness

How far the Nyquist curve is from the -1 pt
→ The farther the better

Characterization

Margins measure the distance from the -1 pt

TP Control Systems


What we want to focus on today is the question of robustness. We want to be robust to a number of different uncertainties that can come and impact our system. These uncertainties can come from things like model mismatch. Maybe we've done a very poor job of modelling the dynamics of our system. Perhaps we just didn't have enough time or enough budget to go and build a full high-fidelity model of our system, which is normally the case. Or perhaps that system is changing with time. Maybe we're building thousands of these different systems and everyone is going to be slightly different. Or maybe we have a delay mismatch. Maybe we're controlling something over the Internet and as a result, the delay varies randomly with time. There's a huge range of different sources for uncertainty. But what we want to know is whether our controller is still going to work, even though the real system is quite a bit different from the system that, in fact, we designed our controller for. The way that we're going to measure robustness is to measure how far the Nyquist curve is, so over here, our Nyquist curve, how far this is from the Nyquist point, from the -1 point.

Notes

Summary



1m 24s



Uncertainties

Due to model mismatch such as gain or delay mismatch, noise, etc.

Robustness

How far the Nyquist curve is from the -1 pt
→ The farther the better

Characterization

Margins measure the distance from the -1 pt

TP Control Systems

The basic idea is that our Nyquist plot, this line here, this is the plot of the model of our system. Of course, this is an approximation for what's going to happen in the real world. But the basic idea is if we have a model like this, perhaps the real system is going to lie somewhere not too far away. If the original model of our system is a long way away from the -1 point, then if our model is representative of our real system, we can hope that the real system is also not going to cross -1 and therefore also be stable. Now, the way that we characterise the level of robustness that our system has is, of course, through margins. The margins are effectively just different ways of measuring how far our model Nyquist curve is from the -1 point.


Notes

Summary



2m 30s

Phase and Gain margins (Nyquist)



Gain margin

Gain increase before instability, expressed in dB
as $-20\log_{10}(a)$

Phase margin ϕ

Phase difference before instability

TP Control Systems

We're going to look at three different margins today. The first two are the most common, the gain margin and the phase margin. The gain margin is a measure of how wrong we can be in terms of the gain of our system before we become unstable.


Notes

Summary



3m 20s

Phase and Gain margins (Nyquist)



Gain margin

Gain increase before instability, expressed in dB
as $-20\log_{10}(a)$

Phase margin ϕ

Phase difference before instability

TP Control Systems

You can see that if this is our Nyquist curve here, this black line, you can see that if we doubled this particular, the gain of our system, it would grow, it would look exactly the same, it would just get bigger. There's a magic number, namely $1/a$. If this value here is a , then the magic number is $1/a$, which will cause our gain or our Nyquist plot to go exactly through the -1 point, and therefore for our system to become unstable. The gain margin is simply a measure of $1/a$, which we write in terms of decibels as $-20\log_{10}(a)$.


Notes

Summary



3m 42s

Phase and Gain margins (Nyquist)



Gain margin

Gain increase before instability, expressed in dB
as $-20\log_{10}(a)$

Phase margin ϕ

Phase difference before instability

TP Control Systems

The second margin we're going to look at is the phase margin. Well, the first one, the gain margin was a measure of how wrong we can be in gain before we become unstable. The phase margin is simply a measure of how wrong we can be in phase before we've become unstable. To measure the gain margin, all we need to do is to look at the unit circle. You can see the unit circle and the intersection with our curve right here. We can see that we can be wrong. This whole curve can rotate by up to 5 degrees before this point ends up going directly through the -1 point. Exactly this measure here, ϕ , is the measure of our gain margin or how wrong we can be in phase at the -1 point before we become unstable.


Notes

Summary



4m 19s

Gain margin & critical frequency



Critical frequency

Frequency ω at which the **phase** of $K()H()$ becomes less than -180°

Gain margin

Difference between 0 dB and the magnitude in dB at the critical frequency

TP Control Systems

How do we read these margins off of a Bode plot? We've seen it on the Nyquist plot. Now we want to look at it on the Bode plot. The first one, of course, is the gain margin. The gain margin is a measure of how wrong we can be in gain before our system becomes unstable. Unstable means that we are going through the -1 point. -1 has a phase of -180 degrees over here, and it has a gain of one, which, of course, is 0 decibels up here. The first thing we do is we start from the phase of -180 degrees and we determine our critical frequency, which is simply where our Bode plot, our phase plot in this case, goes through this -180 degree point. Now, for this particular figure, just focus on the black lines. We can ignore the green and the red. Once we have our critical frequency, we simply go up and we can determine the magnitude of our Bode plot at that particular frequency. If that magnitude happened to be going through a gain of 0 decibels, we would exactly be going through Nyquist point, through the -1 point. This measure right here is our gain margin. The convenient thing about reading it off of a Bode plot as opposed to a Nyquist plot, is of course, it's already in decibels.


Notes

Summary



5m 07s

Phase margin & crossover frequency



Crossover frequency

Frequency ω at which the **magnitude** of $K()H()$ is 0 dB

Phase margin

Difference between the phase at the crossover frequency and -180°

TP Control Systems

This number exactly is our gain margin. The phase margin is equally easy to read off of a Bode plot. In this case, we want to know if we're at magnitude 1, how wrong can we be in phase before we end up going through the -1 point? The first thing we do is we look at the 0 dB point or the gain of 1, and we find where our plot goes through the 0 decibels point, and this is our crossover frequency. We then follow this down to the phase plot, and we calculate the phase at this particular point. We can see if that phase was -180 , then this would exactly be going through the -1 point. Our phase margin is this measure. It's the distance from the phase of our system to the -180 point.


Notes

Summary



6m 34s

Delay margin



Delay margin

- Conversion of phase margin from angle to time
- $$\text{Delay margin} = \frac{\text{Phase margin [radians]}}{\text{Crossover frequency}}$$
- Represents the largest pure delay that can be neglected before the system becomes unstable

TP Control Systems

The last margin that we want to look at today is the delay margin. This is effectively another way of interpreting the phase margin. Effectively, we're just taking the phase margin and we're converting it from radians to time. We define the delay margin as the phase margin measured in radians divided by the crossover frequency, which is measured in radians per second. The delay margin is measured in seconds. Effectively, this is just representing the largest pure delay that we can add to our system before it becomes unstable. A delay is just simply going to rotate our system up like this.

Notes

Summary



7m 32s

Experiment 1: System model in position


Using the remote experimentation tool

- Apply a step within the linear range
(in position and open-loop mode)

- Save the measurements

Using the temporal fit tool

- Load the measurements and fit the model
- Save the **model** using the **save** button



Note: very similar to what has been done in M4

TP Control Systems

We're now ready to start with the experiments today. The first thing we want to do is to build a system model in the position mode. Go ahead and connect the remote experimentation tool, apply a step that is within your linear range. Do this in position mode and in open-loop mode. Save your measurements and open the temporal fit tool, load your measurements and then fit a model from the step response. Save this model so that you have the two values, tau and k that you need in order to represent your model. Now, you should be able to do this very quickly because, of course, we've done something very similar to this before in the fourth module.

Notes

Summary



8m 11s

PID LOOP SHAPING tool

Choose the plant model

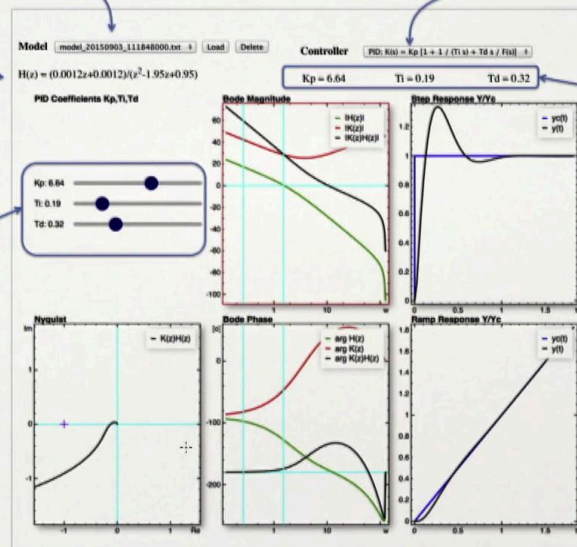
Selected model

Choose the controller to be used (P, PI, PD, PID)

Change the controller parameters

Selected controller parameters

Note that the model parameters you identified previously are digitized with $h = 0.04s$



Tool: Play interactively with the controller parameters and display the resulting information/plots

TP Control Systems

For the next part of this experiment, we're going to use the PID loop shaping tool. This tool looks like this. The first thing you want to do is to load the plant model over here. You'll notice that the model shows up over here in discrete time. The tool will also automatically take the model that you fit in continuous time and sample it at 40 milliseconds in order to generate a discrete time model. You can choose the control structure that you want to see up here, whether it's P, PI, PD, or PID. Then the key piece that you play with is over here. We can change the three controller parameters: proportional gain, integral gain, and derivative gain. Now, what you're seeing and what will change interactively as you change these values is the Bode plot magnitude and phase over here, a step response, and a ramp response over here. The primary piece that we're focused on today is the Nyquist plot over here.

Notes

Summary



8m 50s

Experiment 2 – Find critical gain

Using the model (in position) previously identified

- Compute P-controller to make the system unstable

Using the PID tool

- Find experimentally the K_p that will make the system unstable

Connect to remote experiment

- Validate the gain that destabilizes the system

You may find slight variation due to saturation

- Observe that U and Y are 180° phase shifted

Once observed, set the gain to a "normal" value (i.e. <5)

TP Control Systems

In the second experiment, what we want to do is to find the critical gain. Use the model that you derived previously and calculate a P controller that will cause this to just become unstable. Then open the PID tool and do this again, but now do it experimentally by adjusting the proportional gain until the system becomes unstable. Go ahead and connect to the remote experiment, and your goal is to validate that the gain that you measured before up here, which, of course, both of the gains that you saw before should be the same, and validate that this gain does in fact destabilise the system. Now, as a result of saturation, you might find that there is, in fact, a slight variation, a slight difference between what you have up here and what you see in the experiment. But that's simply because we're not modelling the non-linear effect of saturation in this portion up here. Now, the key thing you want to observe is that your input and your output are in fact phase shifted by 180 degrees. Of course, this is because you should be going through -1 point, which has a phase of -180. Now, once you have in fact observed this, please do set the gain back to a normal value. Don't leave the system running and cycling through an unstable mode. This is really not very good for the system itself.

Notes

Summary



9m 45s


Experiment 3 – System mismatch (delay)

Using the model (in position) previously identified and the PID loop shaping tool

- Find the K_p to have a phase margin of $\sim 60^\circ$
- Compute the delay margin

Connect to remote experiment
(closed loop & position mode)

- Set K_p to the value found above
- Apply a step input and observe the response
- Add a perturbation (delay in the command > 0)

Increase the delay 1 by 1 Ysc 10 ☒ P clr buff 
Perturbation, delay 0

- Change the set point value and observe the response
 - Is the minimal destabilizing gain equal to the delay margin?

TP Control Systems

In the third experiment, what we want to get a feel for is what happens when our model and our system are different. Specifically, we want to look at the delay margin. Take the model that you identified previously and put it into the PID loop shaping tool. Find the proportional gain so that you have a phase margin of 60 degrees, roughly. Then go ahead and compute the delay margin. Now, connect to the remote experiment in closed loop and in position mode. Set the proportional gain that you found previously where we have a phase margin of 60 degrees, and apply a step input and see what happens. What we want to do is now add a perturbation in terms of a delay. You can increase the delay slowly and see what the response is going to be.

Notes

Summary



11m 08s

Experiment 3 – System mismatch, cont'

Using the model (in position) previously identified and the PID loop shaping tool:

- Use the K_p for a phase margin of $\sim 60^\circ$
- Enable perturbation (delay in the command)
- Increase the delay until $K()H()$ becomes unstable
- Observe its effect in the Nyquist plot

TP Control Systems

Your goal is to increase the delay until your system becomes unstable to observe what impact this has on the Nyquist plot. Are we observing that the system is becoming unstable in the Nyquist plot as we're also seeing it on the real system? Finally, is the delay margin that you measured previously or that you calculated previously equal to the delay that it takes in order to destabilise your system?

Notes

Summary



12m 01s

Conclusions

What you have learned today

- Robustness margins can be interpreted and determined with the Bode and Nyquist plots of $K()H()$
- Margins tells you how close you are to instability
- System mismatch such as delay can destabilize the system

TP Control Systems

Today, we focused on robustness margins. We saw that we can readily interpret Bode plots and Nyquist plots in order to read these margins off of these types of plots. The point of the margin is it tells us how close our system is to becoming unstable. A large robustness margin means that we can be extremely incorrect in terms of our model and our system will still be stable. A small margin means that if our system is slightly different from the model that we've taken, in fact, it will become unstable very, very easily. We then saw, of course, that a system mismatch, such as things like delay, noise, can certainly destabilise our system, and this is why robustness margins are important.

Notes

Summary



12m 29s