AI Enabled Control Engineering

Professor Xun Huang

Aeronautics and Astronautics
College of Engineering
Peking University

huangxun@pku.edu.cn
https://xunger99.github.io/xunger/

Recap

- Closed-loop transfer function.
- Characteristic equation.
- Poles and zeros.
- Stability and BIBO stability.
- Routh criterion and root locus.

Lecture 4

- Frequency response.
- Bode plot and Nyquist plot.
- Gain margin and phase margin.
- State-space modelling perspective.
- Controllability and observability.
- State feedback perspective.
- Limitations of purely classical methods.
- A first introduction to reinforcement learning for control.

Why go beyond poles on the complex plane?

So far, we have mainly used

- characteristic equations,
- pole locations,
- Routh criterion,
- root locus.

These are powerful, but they are not the whole story.

We also want to know:

- how a system responds to sinusoidal inputs,
- how robust the closed loop is,
- how to describe multi-state systems directly,
- how to design control when a precise model is hard to obtain.

Frequency response concept

Frequency response studies how a system reacts to sinusoidal signals of different frequencies.

If the input is

$$u(t) = A \sin(\omega t),$$

then the output has

- a different amplitude,
- a different phase,
- but the same frequency in an LTI system.

So frequency response tells us which frequencies are amplified or attenuated.

Bode Plots

These are two plots; the magnitude of $GH(i\omega)$ and the phase of $GH(i\omega)$. The frequency ω is the major variable used to assess performance of a system with a Bode plot. It is normally logarithmic that the x-axis is presented in a logarithmic form. The y-axis is normally presented in a decibel form.
Thus

$$db = 20\log_{10}(GH(i\omega)).$$

Thus for example if $|GH(i\omega)| = 10 \Rightarrow 20\log(10) = 20db$. The plot of $\log(\omega)$ against db is called the Bode Magnitude Plot. The plot of phase angle against $\log(\omega)$ is called the Bode Phase plot.

Bode Plots

In the following figure the function is:

$$G(s) = s.$$

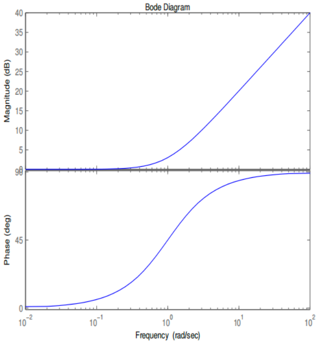

Figure: The Bode plot of a simple zero

Bode Plots

In the following figure the function is:

$$G(s) = \frac{1}{s}.$$

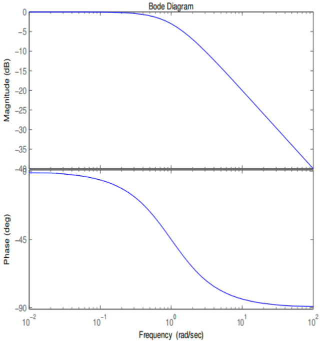

Figure: The Bode plot of a simple pole.

Bode Plots

In the following figure the function is:

$$P(i\omega) = \frac{100\left(1 + i\frac{\omega}{10}\right)}{1 + i\omega}.$$

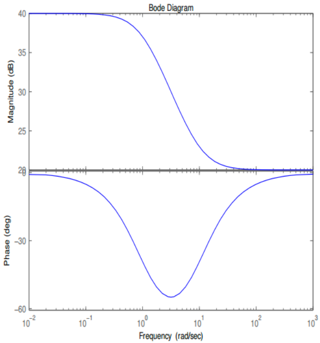

Figure: The Bode plot of example in total.

Bode Plots

- The Bode plots of multiple poles, zeroes, and time delay?
- The Bode plots tell us information of stability.
- Bandwidth of a system is the range of frequencies of the input over which the system will respond satisfactorily.

Bode view of the cart-pole

For the linearized pole-angle channel

$$G_\theta(s) = \frac{\Theta(s)}{U(s)} = -\frac{1.4634}{s^2 - 15.7756},$$

we study its frequency response by substituting

$$s = i\omega.$$

Then

$$G_\theta(i\omega) = -\frac{1.4634}{-\omega^2 - 15.7756}.$$

This tells us how the pole-angle channel reacts to sinusoidal forcing at different frequencies.

Frequency Domain Specifications

Gain margin is a measure of the relative stability of the magnitude of the reciprocal of the open loop transfer function ω_π at which the phase angle is -180°. The gain margin is therefore

$$K = \frac{1}{|G(i\omega_\pi)H(i\omega_\pi)|},$$

and

$$\angle G(i\omega_\pi)H(i\omega_\pi) = -180^\circ.$$

This is also called the phase crossover frequency ω_π.

Frequency Domain Specifications

Phase margin is a measure of the relative stability of a control system, defined as the phase angle, ϕ_1 of the **open Loop** transfer function at a **Unity Gain**.

$$\phi_{PM} = \left(180^\circ + \angle G(i\omega_1)H(i\omega_1)\right),$$

where

$$|G(i\omega_1)H(i\omega_1)| = 1.$$

Here ω_1 is the cross over frequency.

Frequency Domain Specifications

Example

Determine the gain margin for a system where

$$GH(i\omega) = \frac{1}{(s+1)^3},$$

$$GH(i\omega) = \frac{1}{(i\omega+1)^3} = \frac{1}{(\omega^2+1)^{\frac{3}{2}}}\angle(-3\arctan(\omega))$$

Thus,

$$-3\arctan(\omega) = -\pi \Rightarrow \tan\left(\frac{\pi}{3}\right) = \omega = 1.732.$$

Therefore the gain margin is:

$$\text{Gain Margin} = \frac{1}{GH(i\omega)} = 8.$$

Frequency Domain Specifications

The phase margin is found as follows:

$$|GH(i\omega)| = \frac{1}{(\omega^2+1)^{\frac{3}{2}}} = 1 \Rightarrow \omega = \omega_1 = 0.$$

$$\therefore \ \phi_{PM} = 180^\circ + 3\arctan(0) = 180^\circ \equiv \pi.$$

Nyquist plot

Nyquist plot is another frequency-domain stability tool.
It is obtained by plotting

$$G(i\omega)H(i\omega)$$

in the complex plane as ω varies.

It is useful for studying:

- gain margin,
- phase margin,
- closed-loop stability from open-loop frequency response.

Compared with Bode plot:

- Bode separates magnitude and phase,
- Nyquist keeps them together in the complex plane.

Classical frequency tools: what they do well

Classical frequency tools are excellent for

- SISO transfer functions,
- loop shaping,
- robustness intuition,
- stability margins.

For many practical engineering systems, Bode and Nyquist remain indispensable.

However, they are less natural when

- the system has many internal states,
- the system is MIMO,
- we want direct state feedback design.

Why state-space modelling?

The cart-pole is naturally described by the state vector

$$x = \begin{bmatrix} p & \dot{p} & \theta & \dot{\theta} \end{bmatrix}^\top .$$

This immediately suggests a state-space viewpoint:

$$\dot{x} = Ax + Bu, \qquad y = Cx + Du.$$

State-space modelling is useful because it

- keeps track of internal states directly,
- naturally handles multi-state systems,
- supports state feedback and observer design.

State-space modelling perspective

A continuous-time linear dynamical system has the form

$$\dot{x}(t) = Ax(t) + Bu(t), \qquad y(t) = Cx(t) + Du(t),$$

where

- $x(t) \in \mathbb{R}^n$ is the state,
- $u(t) \in \mathbb{R}^m$ is the input,
- $y(t) \in \mathbb{R}^p$ is the output.

This is also called an m-input, n-state, p-output LDS.

SISO vs MIMO

In transfer-function language, we often think in terms of SISO.

But in state-space language, we naturally distinguish:

- SISO: single-input, single-output,
- MIMO: multiple-input, multiple-output.

For the cart-pole:

- one can view it as SISO if only one output is emphasized,
- but physically it already has multiple meaningful outputs such as p and θ.

So state-space language is often closer to the real system structure.

Cart-pole state-space model

For the linearized cart-pole near the upright equilibrium,

$$\dot{x} = Ax + Bu, \qquad x = \begin{bmatrix} p \\ \dot{p} \\ \theta \\ \dot{\theta} \end{bmatrix}.$$

With the parameters used in Gymnasium CartPole,

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & -0.7171 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 15.7756 & 0 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ 0.9756 \\ 0 \\ -1.4634 \end{bmatrix}.$$

This form is the starting point for controllability, observability, and state feedback.

Controllability

A system is controllable if we can drive the state to desired locations by suitable input. For

$$\dot{x} = Ax + Bu,$$

the controllability matrix is

$$\mathcal{C} = \begin{bmatrix} B & AB & A^2B & \cdots & A^{n-1}B \end{bmatrix}.$$

The system is controllable if

$$\mathrm{rank}(\mathcal{C}) = n.$$

Intuition:

- can the actuator influence all important state directions?

Observability

A system is observable if we can infer the state from measured outputs. For

$$\dot{x} = Ax + Bu, \qquad y = Cx + Du,$$

the observability matrix is

$$\mathcal{O} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}.$$

The system is observable if

$$\mathrm{rank}(\mathcal{O}) = n.$$

Intuition:

- do the sensors reveal all important state directions?

Why go beyond purely classical methods?

Classical control is powerful and elegant.

- It gives stability guarantees.
- It has clear mathematical interpretation.
- It works very well near known operating points.
- It is often efficient, safe, and easy to deploy.

However, classical control usually relies on

- a reasonably accurate model,
- a fixed control structure,
- carefully selected parameters.

When these assumptions become weak, we may need more flexible methods.

Classical control vs AI-based control

A rough comparison is:

- **Classical control:**
  - starts from a model,
  - emphasizes stability and analysis,
  - designs a controller explicitly.
- **AI-based control:**
  - starts from data or interaction,
  - emphasizes performance improvement,
  - learns a control policy from experience.

In practice, these two viewpoints are often complementary rather than contradictory.

A first reinforcement learning viewpoint

In reinforcement learning, we describe the control problem using

- state s,
- action a,
- reward r,
- next state s'.

For cart-pole:

- state: position, velocity, angle, angular velocity,
- action: push left or push right,
- reward: keep the pole upright as long as possible.

The controller is no longer written directly as a formula such as $u = -Kx$. Instead, it is learned as a policy from interaction.

What comes next?

So far, we have seen three viewpoints for the same control problem:

- transfer-function viewpoint,
- state-space viewpoint,
- reinforcement-learning viewpoint.

For the inverted pendulum, these viewpoints answer different questions:

- Is the system stable?
- Can the state be controlled and observed?
- Can a good policy be learned from data?

In the next lecture, we will move from this introduction to an actual RL algorithm on CartPole.

Summary

- Frequency response, Bode plot, and Nyquist plot.
- Gain margin and phase margin.
- State-space modelling perspective.
- Controllability and observability.
- Why state feedback is natural for multi-state systems.
- Limitations of purely classical methods.
- A first bridge from classical control to AI-based control.