6.3100 April 13, 2026 Lecture: Robustness to Modeling Errors

April 14, 2026

Last Lecture

Tools for thinking in both time and frequency domain terms.

- Transfer function from a state space model:

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

(input $w(t)$, state $x(t)$, output $y(t)$) has transfer function

$$H(s) = D + C(sI - A)^{-1}B.$$

On particular, the degree of the denominator of $H(s)$ equals the dimension of $x(t)$.

- Parceval identity: assuming all eigenvalues of $A - BK$ have negative real part, for the state space model

$$\dot{x}(t) = Ax(t) + Bu(t) + B_1 w(t), \quad u(t) = -Kx(t), \quad z(t) = Cx(t) + Du(t)$$

(control input $u(t)$, noise/disturbance input $w(t)$, state $x(t)$, cost output $z(t)$) let $z_{ZIR}(t)$ be the output with zero noise $w(t) \equiv 0$ and initial condition $x(0) = B_1$. Also, let $H_{w \to z}(s)$ be the transfer function from w to z. Then

$$\int_0^\infty |z_{ZIR}(t)|^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |H_{w \to z}|^2 d\omega.$$

Outline for Today

- Phase margin as a measure of robustness to modeling errors.
- Zero exclusion principle.
- Expressing phase margin and other robustness measures in terms of closed loop frequency response.

Re-Visiting Phase Margin

So far, all continuous time feedback control scenarios considered in 6.3100 are covered by the block diagram below:

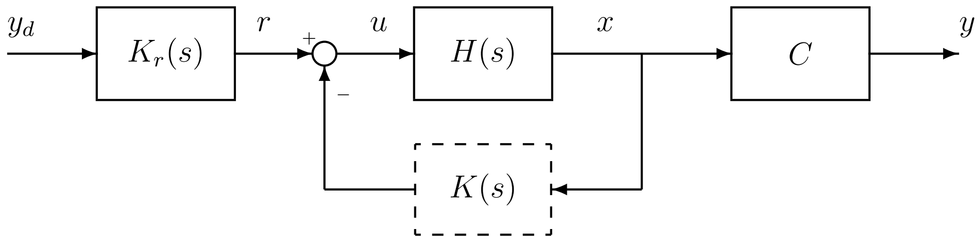

where $y(t)$ is the controlled output, $y_d(t)$ is the desired value of the controlled output, and $u(t)$ is the command:

In the output feedback case: $C = 1$ (i.e., $x = y$), the block with transfer function $H(s)$ describes how the controlled output $y(t)$ depends on command $u(t)$, and we use $K_r(s) \equiv K(s)$, where the block with transfer function $K(s)$ describes the controller (e.g., PID or lead), practically applied to the difference $e(t) = y_d(t) - y(t)$.

In the full state feedback case (as in pole placement or LQR): we use $H(s) = (sI - A)^{-1}B$ (an n-by-1 matrix, where n is the dimension of $x(t)$), so that the block with transfer function $H(s)$ describes how the state $x(t)$ of the controlled plant depends on command $u(t)$, $K(s) = K$ is a constant 1-by-n matrix (the row matrix of state feedback gains), and K_r is a constant scalar calculated to make the closed loop steady state response to $y_d(t) \equiv 1$ to be 1.

For both scenarios, assuming that the feedback system on the block diagram is stable, the *phase margin* is defined (in radians) as the largest ϕ_0 such that $1 + e^{j\phi}K(j\omega)H(j\omega) \neq 0$ for all $\phi \in (-\phi_0, \phi_0)$.

Recall that our original motivation for considering phase margin was "to keep $K(s)H(s)$ as far from -1 as possible". Let us re-visit both our motivation and the means to address it. In particular, the following questions surface naturally here:

(1) Why is it sufficient to look at the imaginary axis only? After all, for stability, we need $K(s)H(s) \neq -1$ for all s with $\mathrm{Re}(s) \geq 0$, not just for s on the imaginary axis, as zeros of $1 + K(s)H(s)$ are the natural frequencies of the closed loop system.

(2) If the goal is to keep $K(j\omega)H(j\omega)$ away from -1, a very straightforward mathematical measure of this is to keep $\left|\frac{1}{1+K(j\omega)H(j\omega)}\right|$ small (this would also fit nicely with LQR's ability to "push down" absolute values of selected closed loop transfer functions, as $G(s) = \frac{1}{1+K(s)H(s)}$ is the closed loop transfer function from r to u. Why not using this alternative?

(3) Which specific "system level" property is addressed by the phase margin?

Robustness to Relative Errors in the Plant Transfer Function

Stability of the closed loop system shown on the previous block diagram is equivalent to two conditions:

(a) the unstable poles of $K(s)H(s)$ are the same (counting multiplicity) as the unstable poles of $H(s)$ and $K(s)$, combined;

(b) $1 + K(s)H(s) \neq 0$ for all complex scalars s with $\mathrm{Re}(s) \geq 0$.

Here condition (a) helps to reject the "ridiculous" scenarios, such as trying to use a *zero* controller $K(s) \equiv 0$ to stabilize an *unstable* $H(s)$. Practically, condition (b) is the important one.

But what if we are *not sure* what $H(s)$ is? For example, when modeling the maglev system, even if we accept that the command-to-position transfer function equals

$$H(s) = \frac{-\lambda_E \gamma}{(s - \lambda_E)(s^2 - \gamma_{da/dy})},$$

there is little confidence that we know the values of λ_E, γ, and $\gamma_{da/dy}$ accurately. As a result, one can argue that $H(s)$ is known, at best, with some *relative error* δ, as in

$$H(s) = (1 + \delta)H^0(s),$$

where $H^0(s)$ is a fixed "nominal" transfer function, and δ is a complex number which is allowed to depend on s, as shown on the block diagram below.

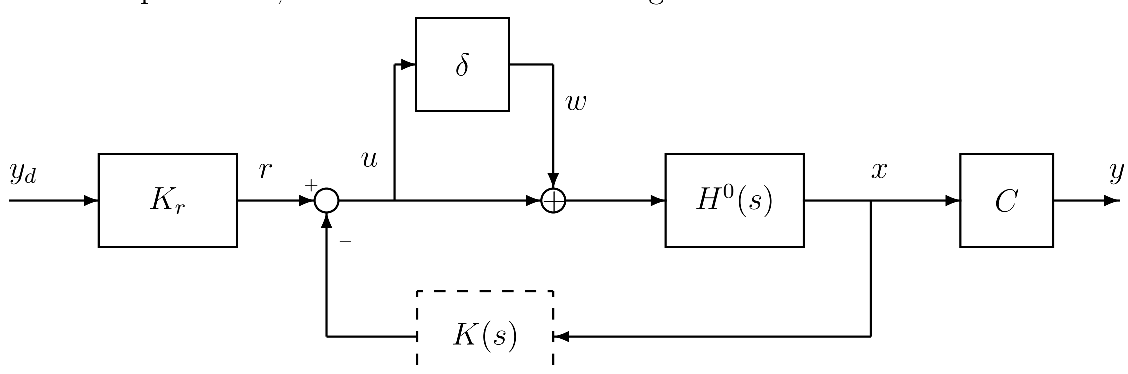

Since tracking the exact range of the relative error δ as a function of s is a difficult task (and using such super accurate knowledge of uncertainty is even more involved), one usually settles for a rough assumption that δ ranges over a nicely defined set Δ of complex numbers, independently of s, at least for all s on the imaginary axis $s = j\omega$. The following mathematical result, sometimes referred to as the "zero exclusion principle", claims that, under minor assumptions, to verify that $1 + K(s)H(s) \neq 0$ for $\mathrm{Re}(s) \geq 0$, it is sufficient to check that $1 + (1 + \delta)K(s)H^0(s) \neq 0$ on the imaginary axis for all $\delta \in \Delta$:

Zero Exclusion Principle: If $L(s)$ and $L^0(s)$ are two strictly proper transfer functions with the same number of unstable poles, and Δ is a closed connected subset in the complex plane, such that

(a) $1 + L^0(s) \neq 0$ for $\mathrm{Re}(s) \geq 0$;

(b) for every $\omega \in (-\infty, \infty)$, $L(j\omega) = (1+\delta)L^0(j\omega)$ for some $\delta \in \Delta$;

(c) $1 + (1+\delta)L^0(j\omega) \neq 0$ for all $\omega \in (-\infty, \infty)$ and $\delta \in \Delta$;

then $1 + L(s) \neq 0$ for $\mathrm{Re}(s) \geq 0$.

In the zero exclusion principle, $L^0(s)$ stands for $K(s)H^0(s)$, the "nominal" open loop transfer function, $L(s) = K(s)H(s)$ is the "true" transfer function. Condition (a) means that feedback $K(s)$ stabilizes the nominal model, condition (b) means that Δ constraints the relative error of approximating $H(s)$ by $H^0(s)$ *on the imaginary axis*, and condition (c) means that, within the relative error bounds, no zeros of $L(j\omega)$ are possible on the imaginary axis. The "zero exclusion" refers to the imaginary axis acting as a barrier for the zeros of $1 + L(s)$: all zeros of $1 + L^0(s)$ have negative real part, and as $L^0(s)$ is "deformed" into $L(s)$, no zeros may slip into the right half plane $\mathrm{Re}(s) > 0$ without first crossing the imaginary axis, which is prohibited by (c).

How Phase/Gain Margins Relate to Relative Errors

By definition, phase margin of controller $K(s)$ which stabilizes nominal system $H(s) = H^0(s)$ is larger than $\phi_0 > 0$ (in radians per second) if and only if $1 + e^{j\phi}K(j\omega)H^0(j\omega) \neq 0$ for all $\omega \in (-\infty, \infty)$ and all $\phi \in [-\phi_0, \phi_0]$. Equivalently, this means that the conditions of zero exclusion principle are satisfied for the relative uncertainty set $\Delta = \Delta_{pm[\phi_0]} = \{e^{j\phi} - 1 : \ \phi \in [-\phi_0, \phi_0]\}$.

Similarly, if we accept the (slightly non-conventional) definition that the gain margin (of the same controller $K(s)$ stabilizing the same nominal system $H(s) = H^0(s)$) is larger than $r_0 > 1$ if and only if $1 + rK(j\omega)H^0(j\omega) \neq 0$ for all $\omega \in (-\infty, \infty)$ and all $r \in [1/r_0, r_0]$, then equivalently this means that the conditions of zero exclusion principle are satisfied for the relative uncertainty set $\Delta = \Delta_{gm[r_0]} = \left[\frac{1}{r_0} - 1, r_0 - 1\right]$.

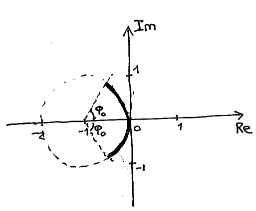

$\Delta_{pm[\varphi_0]}$

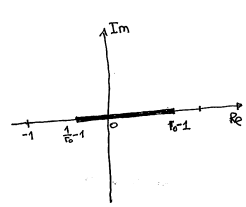

$\Delta_{gm[r_0]}$

Robustnes to Modeling Error Expressed in Closed Loop Terms

So far, we have formulated robustness to relative modeling error in terms of the nominal *open loop* frequency response $K(j\omega)H^0(j\omega)$. This makes it inconvenient to address robustness via LQR optimization, formulated in terms of the *closed loop* frequency responses.

It turns out that it is easy to re-formulate robustness to relative modeling uncertainty in terms of the closed loop transfer functions, which in turn provides a hint for how to set up LQR penalties. Indeed the condition

$$1 + (1+\delta)K(j\omega)H^0(j\omega) \neq 0 \quad \text{for } \delta \in \Delta$$

can be re-written as

$$G(j\omega) = -\frac{K(j\omega)H^0(j\omega)}{1 + K(j\omega)H^0(j\omega)} \neq 1/\delta \quad \text{for } \delta \in \Delta.$$

Now the transfer function $G = -\frac{KH^0}{1+KH^0}$ is the closed loop transfer function from w to u (with the δ block removed). When 0 is in the interior of the uncertainty set Δ (which is always reasonable to assume), the set $\Delta^{-1} = \{1/\delta : \ \delta \in \Delta\}$ has zero in the interior of its *complement* $\mathbb{C}\backslash\Delta^{-1}$. Hence improving robustness to relative error in modeling $H(s)$ means that $G(j\omega)$ should be pushed to be as small as possible, which suggests that u should be penalized proportionally to the amount of the uncertainty in modeling H.

This observation can be extended further. The description of the modeling error in $H(s) = (1+\delta)H^0(s)$ can be given in the form of the block diagram

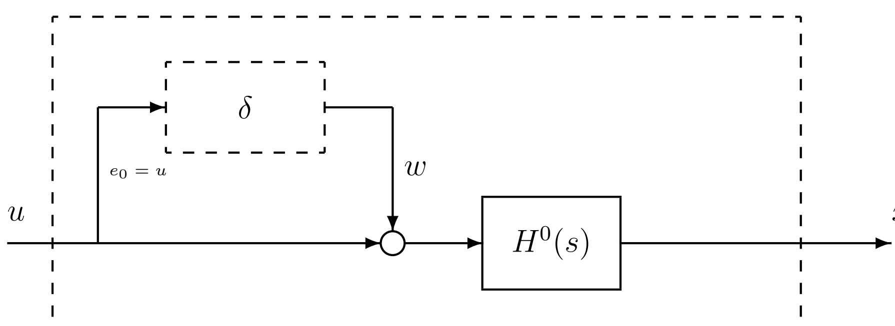

What we just discovered is that, in order to improve robustness of the feedback system with respect to the modeling error represented by δ, one has to push down the absolute value of the frequency response $G(j\omega)$ from what-would-be the output w of the δ block to what-would-be the input e_0 of the δ block, with the δ block removed, as shown on the block diagram below:

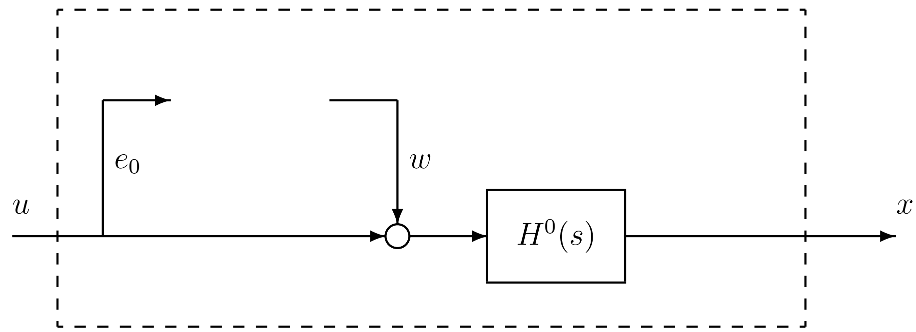

LQR optimization is well suited to push down absolute values of closed loop frequency responses, as was discussed in the previous lecture.

The observation above can be made more universal: wherever one can represent some modeling error by an uncertain gain δ (centered around its nominal value $\delta = 0$) connecting some output e_k to some external input w_q, robustness to this error is improved by trying to push down the absolute value of the closed loop frequency response from w_q to e_k. We will see that this task is well suited to be addressed by LQR and its generalizations.