# 6.3100 April 15, 2026 Lecture:
# LQR and Robustness to Modeling Errors

## Last Lecture

- Phase margin as a measure of robustness to relative modeling errors.
- Robustness guarantees can be assured by analyzing frequency responses ("zero exclusion principle").
- Expressing robustness in terms of closed loop frequency responses.

## Outline for Today

- Guaranteed robustness of LQR to relative modeling errors.
- Other robustness guarantees addressed by LQR.
- Introduction to state estimation.

## Recall: Robustness to Relative Modeling Error

Recall that we are working with a feedback control design setting described by the block diagram below:

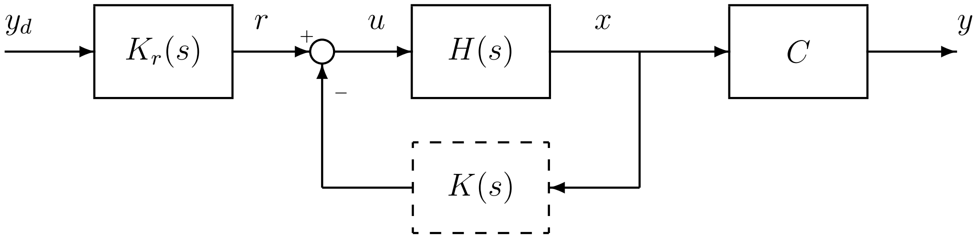

where $y(t)$ is the controlled output, $y_d(t)$ is the desired value of the controlled output, and $u(t)$ is the command.

**In the output feedback case:** $x(t) = y(t)$ is the same as the controlled output, so $C = 1$, the block with transfer function $H(s)$ describes how the controlled output $y(t)$ depends on command $u(t)$, and $K_r(s)$ is the same as $K(s)$, 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):** $x(t)$ is the $n$-dimensional state of a the state space model

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

hence $H(s) = (sI - A)^{-1}B$ (an $n$-by-1 matrix), $K(s) = K$ is a constant 1-by-$n$ matrix (the row matrix of state feedback gains), and $K_r$ is a constant scalar calculated, as a rule, to make the closed loop steady state response to $y_d(t) \equiv 1$ to be 1.

To address imperfection our modeling efforts, we assume that the frequency response $H(j\omega)$ is known, at best, with some *relative error* $\delta$, as in

$$H(j\omega) = (1 + \delta)H^0(j\omega), \qquad \delta \in \Delta,$$

where $H^0(s)$ is a fixed "nominal" transfer function, and $\delta$ is a complex number which is allowed to depend on $\omega$, but ranges over a subset set $\Delta = \{\delta\}$ of complex numbers (typically, $\Delta$ is connected and contains zero). This "relative modeling error" scenario is described by the block diagram below:

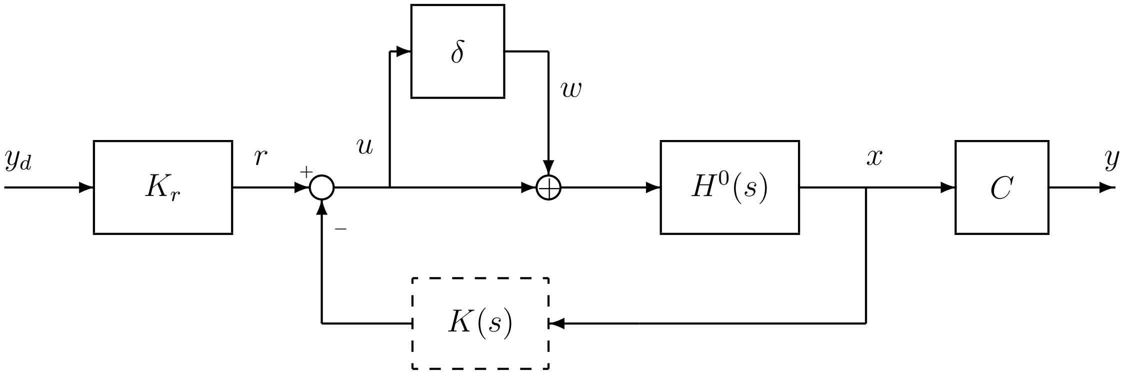

Under mild assumptions, stability of the closed loop follows from stability of the nominal closed loop and the "zero exclusion inequality"

$$1 + (1 + \delta)K(j\omega)H^0(j\omega) \neq 0 \qquad \text{for all} \quad \omega \in (-\infty, \infty), \quad \delta \in \Delta.$$

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

*Phase margin* is an example of such zero exclusion inequality. Recall that, 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 $\phi \in [-\phi_0, \phi_0]$, which is equivalent to the zero exclusion principle 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 $\Delta = \Delta_{gm[r_0]} = \left[\frac{1}{r_0} - 1, r_0 - 1\right]$.

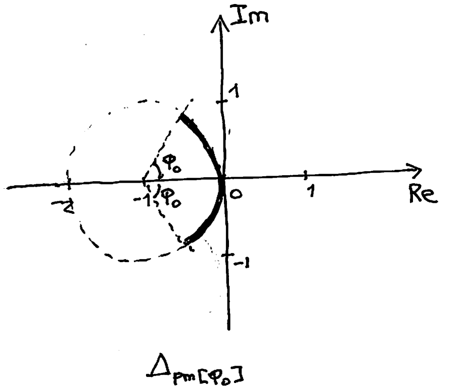

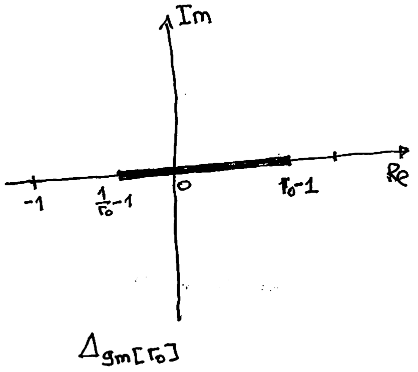

## Guaranteed Robustness of LQR Controllers

It turns out that, as long as the weight matrices $Q$ and $R$ have all positive eigenvalues (i.e., no ridiculous "negative" weights), and $N = 0$ (i.e., there are no cross-terms between $x$ and $u$ in the cost), the LQR controller come with a significant guaranteed degree of robustness to relative modeling error:

If $N = 0$ and $Q$ and $R$ have all positive eigenvalues then the LQR-optimal feedback controller $K$ is robust to the relative error described by $\Delta = \Delta_{lqr} = \{\delta : \ \mathrm{Re}(\delta) \geq -0.5\}$.

In particular, this level of robustness guarantees better than 60° ($\phi_0 = \pi/3$) of phase margin and better than $r_0 = 2$ of gain margin. Moreover, since $\Delta_{lqr}$ indicates that $1 + \mu K H^0(j\omega) \neq 0$ for all complex $\mu$ with $\mathrm{Re}(\mu) \geq 0.5$, for every $h > 0.5$ the scaled LQR controller $K_h = hK$ will have phase margin larger than $\arccos(1/2h)$ (approaches 90° quickly as $h$ increases) and gain margin larger than $2h$.

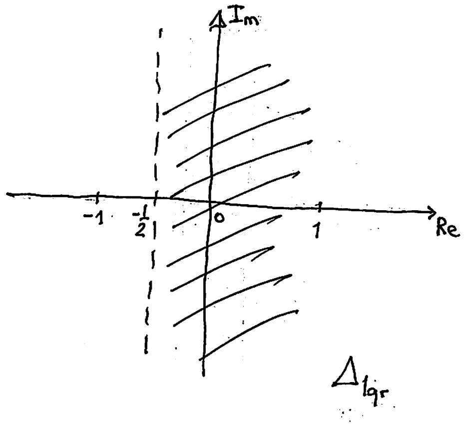

If robustness to relative modeling error was the only feedback design objective, sufficiently scaled up LQR controller would provide very good results. Of course, there are usually multiple conflicting design objectives and constraints. For example, maximal value of control in response to a unit step in desired output would limit the scaling that can be applied to an LQR controller to improve robustness to relative modeling error. To this end, it would help to design the LQR controller which exerts minimal possible control effort to stabilize, which can be achieved by making LQR weight $Q$ negligible as compared to $R$. (Actually, for the maglev model, it is completely acceptable to set $Q = 0$ and $R = 1$, to get the ultimate "expensive control" feedback gains vector $K$.) For example, in the maglev setup with $\lambda_E = -120$, $\gamma = 2000$, and $\gamma_{da/dy} = 900$, the "expensive control" controller scaled by a factor of 3 yields slightly better than 80° phase margin, as compared to less than 30° for the PD controller with $K_p = 3$ and $K_d = 0.1$, and to about 70° for the pole placement full state feedback, all with comparable maximal control value of the unit step response. See the plots below (LQR is left, PD is center, pole placement on the right):

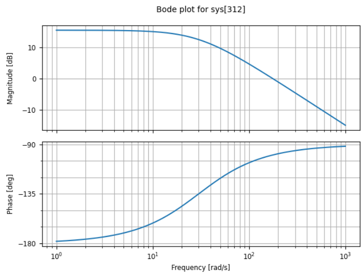

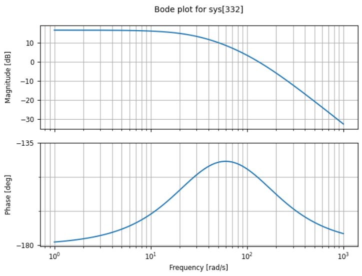

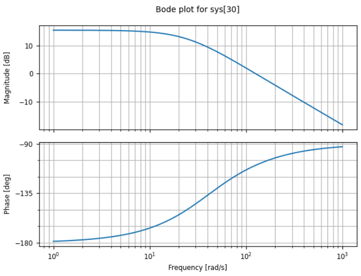

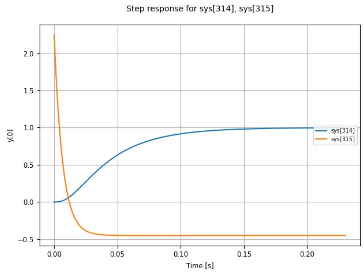

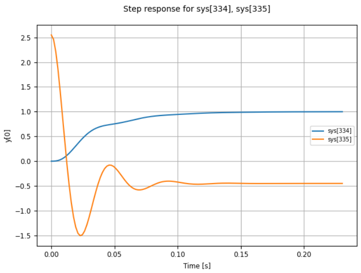

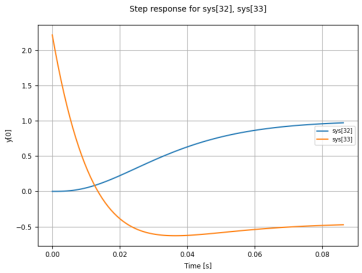

## LQR and Robustness to Multiple Modeling Errors

Consider, for simplicity, the case when the relative modeling error $\delta$ is only limited by its absolute value, i.e., when

$$\Delta = \Delta_{d[r]} = \{\delta \in \mathbb{C} : \ |\delta| \leq r\}$$

for some $r \in (0,1)$. Re-writing the zero exclusion inequality for the relative modeling error

$$1 + (1+\delta)K(j\omega)H^0(j\omega) \neq 0 \quad \text{for } \delta \in \Delta_{d[r]}$$

in terms of the closed loop transfer function from $w$ to $u$:

$$G(j\omega) = -\frac{K(j\omega)H^0(j\omega)}{1 + K(j\omega)H^0(j\omega)} \neq 1/\delta \quad \text{for } \delta \in \Delta_{abs[r]},$$

yields the simple comdition $|G(j\omega)| < 1/r$, which suggests that increasing the LQR penalty on the control signal $u$ will result in improved robustness to relative error limited only by its absolute value.

Modeling errors do not have to limited to scalar multipliers of $H(s)$, as is the case of the relative modeling error discussed so far. For example, relative error in determining the $\gamma_{da/dy} = 900$ coefficient is shown on the block diagram of a part of the maglev model shown below:

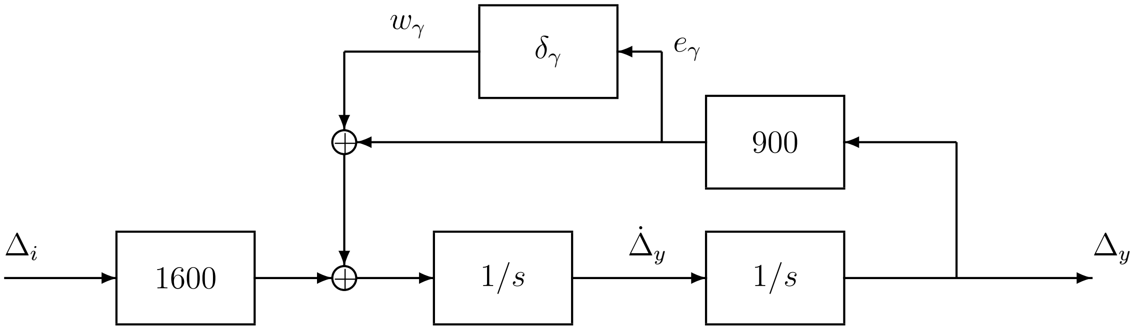

In order to address robustness to absolute-value-bounded small perturbations here, one can penalize the closed loop transfer function from $w_\gamma$ to $e_\gamma$ in the closed loop system (with $\delta_\gamma$ removed).

Ultimately, one may need to assure robustness to multiple simultaneous modeling errors, as, for example, will be the case of the maglev system when the relative modeling error in determining the command-to-current transfer function $\frac{-\lambda_E \gamma_{di/dc}}{s-\lambda_E}$ is to be addressed simultaneously with the relative error in modeling the $\gamma_{da/dy}$ coefficient. In this case the LQR objective should penalize transfer functions from *all* uncertainty outputs $w, w_\gamma$ to *all* uncertainty inputs $u, e_\gamma$, i.e., absolute values of closed loop frequency responses $G_{w\to u}$, $G_{w\to e_\gamma}$, $G_{w_\gamma\to u}$, and $G_{w_\gamma\to e_\gamma}$ are to be pushed down together.

## What if Not All States Are Available for Measurement?

The formulation of LQR optimization assumes that all components of the model's state are available for measurement For example, the full state feedback block diagram for the maglev system shown below:

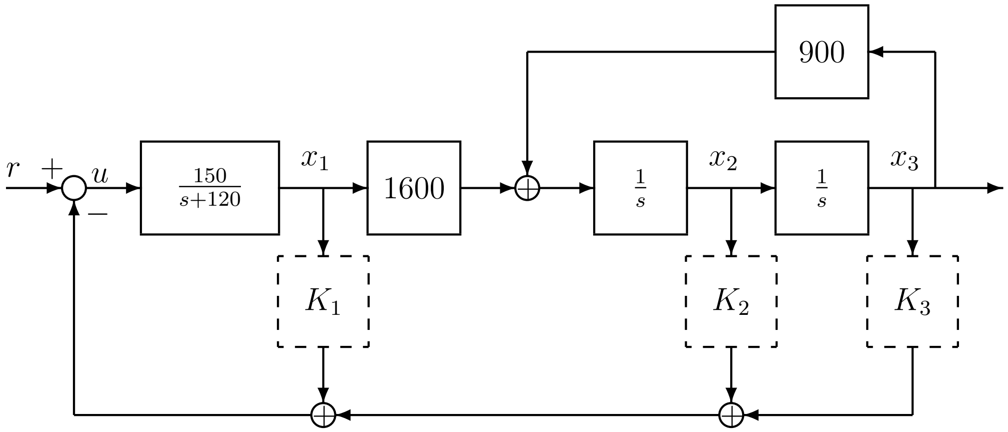

assumes that states $x_1, x_2, x_3$ are available for control action computation. In practice, only a few state components are measured: two ($x_1$ and $x_3$) in our maglev experiment, only one (angular position) in our propeller arm control setup. Worse, all measurements involve errors, typically referred to as "noises", which means that trying to recover non-measured states from the measured ones by (numerical) differentiation will incur significant amplification of noise.

It turns out that there is a systematic way of recovering unmeasured states from the measured ones which does not lead to any unnecessary noise amplification. This is done by feeding the control command and all measurements as inputs into a specially designed state space model, the states of which can then serve as *estimates* of the true states of the original system model. Moreover, this "estimator" system can be optimized via the very same LQR algorithm we have discussed, just applied to a slightly weird state space model (actually, to optimmize the estimator, LQR is applied to the model that is "dual" to the original one). The estimates usually torn out to be so good, that, for state feedback, it is usually better to use estimates of the measured states, instead of the actual measurements! Our objective for the next few lectures is to study the format of state estimation, the type of modeling that preceeds its optimization, the use of LQR in optimizing the coefficients of the estimator, as well as limitations of the whole paradigm.

## Example: Recovering the Middle Signal in a Chain of Integrators

Consider the (rather common) scenario when both the input $r$ and the output $p_2$ of a chain of two pure integrators are measured (with noises $v_1$ and $v_2$, respectively), while the output $p_1$ of the first integrator is not, as described by the block diagram

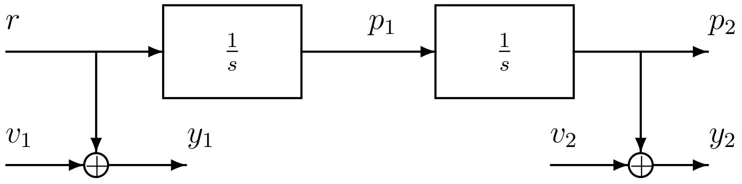

or, equivalently, equations

$$\begin{aligned}
\dot{p}_1(t) &= r(t), \\
\dot{p}_2(t) &= p_1(t), \\
y_1(t) &= r(t) + v_1(t), \\
y_2(t) &= p_2(t) + v_2(t).
\end{aligned}$$

For example, this is the case for our maglev system on the earlier block diagram, with $r = 1600x_1 + 900x_3$, $p_2 = x_3$, $p_1 = x_2$. Since both $x_1$ and $x_3$ are measured (with some noise), it is little stretch to state that $r$ and $p_2$ are measured. How to estimate $p_1(t)$, as accurately as possible, from signals $y_1(t)$ and $y_2(t)$?

Let us denote the desired estimate of $p_1(t)$ by $\hat{p}_1(t)$. We are trying to design an estimator as a linear dynamical system with constant coefficients, two inputs ($y_1$ and $y_2$), and single output $\hat{p}_1$. We will require the transfer function $H_r(s)$ from $r$ to the resulting estimation error $e = \hat{p}_1 - p_1$ to be zero, and, for simplicity, will quantify quality of estimation by the sum of maximal absolute values of the frequency responses $H_k(j\omega)$ from each of the noises $v_k$ to the estimation error $e$. (The more appropriate measure of estimation quality would be the sum of the *integrals* of squares of absolute values of the frequency responses.)

Two "direct" approaches to creating $\hat{p}_1(t)$ fail quite miserably:

**Using $\hat{p}_1(t) = \dot{y}_2(t)$.** Since $p_1$ is the derivative of $p_2$, and $y_2$ is the measured value of $p_2$, it is natural to try the derivative of $y_2$ in the role of $\hat{p}_1$. While the resulting transfer functions from $r$ and $v_1$ to $e$ are zero, the transfer function $H_2(s)$ from $v_2$ to $e$ is $s$, hence the maximal value of $|H_2(j\omega)| = |\omega|$ over all real $\omega$ is infinity.

**Using $\hat{p}_1(t)$ as an integral of $y_1(t)$** Since $r$ is the derivative of $p_1$, and $y_1$ is the measured value of $r$, it is natural to try the integral of $y_1$ in the role of $\hat{p}_1$. While the resulting transfer functions from $r$ and $v_2$ to $e$ are zero, the transfer function $H_1(s)$ from $v_1$ to $e$ is $1/s$, hence the maximal value of $|H_1(j\omega)| = 1/|\omega|$ over all real $\omega$ is infinity.

It turns out there are much better ways to estimate $p_1$, which will be used in the next lecture to introduce the general framework of linear state estimation.