Lecture 11

LQG Control

In this lecture, we study the properties of the linear-quadratic Gaussian (LQG) controller and give several design examples.

11.1 The LQG Controller

In the previous two lectures we have derived the optimal state feedback and the optimal observer for a linear system driven by white noise and with a quadratic cost function. Combining the two solutions, we obtain an optimal output feedback controller—the LQG controller (see Figure 11.1). The result can be summarized in the following theorem:

Theorem 11.1—LQG Control
Given a stabilizable and detectable linear plant

$$\begin{aligned}\dot{x}(t) &= Ax(t) + Bu(t) + w_1(k),\\ y(t) &= Cx(t) + w_2(t),\end{aligned}$$

where w_1 and w_2 are white noise processes with the intensities $\Phi_w = \begin{pmatrix} R_1 & R_{12} \\ R_{12}^T & R_2 \end{pmatrix} > 0$, consider controllers of the form

$$\begin{aligned}\dot{\hat{x}}(t) &= A\hat{x}(t) + Bu(t) + K\big(y(t) - C\hat{x}(t)\big),\\ u(t) &= -L\hat{x}(t).\end{aligned}$$

The stationary variance

$$\mathrm{E}\left(x^T Q_1 x + 2x^T Q_{12} u + u^T Q_2 u\right)$$

is minimized for the $S = S^T > 0$ and $P = P^T > 0$ that solve

$$\begin{aligned} L &= Q_2^{-1}(SB + Q_{12})^T \qquad K = (PC^T + R_{12})R_2^{-1}\\ 0 &= Q_1 + A^T S + SA - (SB + Q_{12})Q_2^{-1}(SB + Q_{12})^T\\ 0 &= R_1 + AP + PA^T - (PC^T + R_{12})R_2^{-1}(PC^T + R_{12})^T \end{aligned}$$

The minimal variance is given by $\mathrm{tr}(SR_1) + \mathrm{tr}[PL^T(B^T SB + Q_2)L]$ □

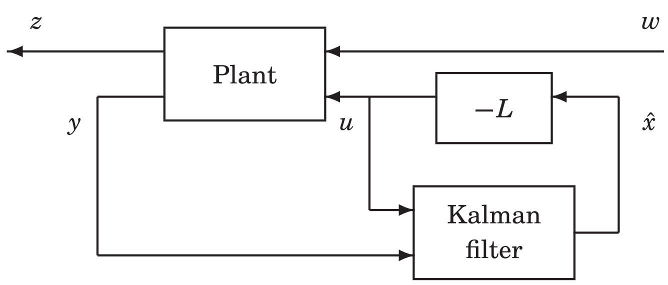

Figure 11.1 An optimal output feedback controller—an LQG controller—consists of an optimal state feedback $-L$ from an optimal state estimate $\hat{x}$ (given by a Kalman filter).

Compared to the stochastic interpretation of the linear-quadratic controller (see Lecture 9), we note that the minimum variance is increased by $\operatorname{tr}[PL^T(B^TSB+Q_2)L]$. This term represents the cost of having noisy measurements and an observer instead of having direct access to the plant state.

The fact that the optimal state feedback and the optimal observer are independent of each other is known as the *separation principle*. For the LQG controller this is true since the optimal state feedback gain L is independent of the state uncertainty, and that the optimal Kalman filter gain K is independent of the control objective. The controller transfer function (from $-y$ to u) of the LQG controller is given by

$$C_{\text{LQG}}(s) = L(sI - A + BL + KC)^{-1}K$$

The controller will have the same number of states as the plant model.

EXAMPLE 11.1—LQG CONTROL OF AN INTEGRATOR
Consider the problem of minimizing $\mathrm{E}(Q_1x^2 + Q_2u^2)$ for

$$\begin{cases} \dot{x}(t) = u(t) + w_1(t) \\ y(t) = x(t) + w_2(t) \end{cases} \qquad \Phi_w = \begin{pmatrix} R_1 & 0 \\ 0 & R_2 \end{pmatrix}$$

The observer-based controller

$$\begin{cases} \dfrac{d}{dt}\hat{x}(t) = A\hat{x}(t) + Bu(t) + K[y(t) - C\hat{x}(t)] \\ u(t) = -L\hat{x}(t) \end{cases}$$

is optimal with K and L computed as follows (see the previous two lectures):

$$\begin{aligned} 0 &= Q_1 - S^2/Q_2 \quad \Rightarrow \quad S = \sqrt{Q_1Q_2} \quad \Rightarrow \quad L = S/Q_2 = \sqrt{Q_1/Q_2} \\ 0 &= R_1 - P^2/R_2 \quad \Rightarrow \quad P = \sqrt{R_1R_2} \quad \Rightarrow \quad K = P/R_2 = \sqrt{R_1/R_2} \end{aligned}$$

The optimal controller in this example is a proportional controller with a low-pass filter. Note that the state feedback gain L only depends on A, B, Q_1 and Q_2, while the Kalman gain K only depends on A, C, R_1 and R_2. □

11.2 Tuning the LQG controller

The LQG controller is a function of the given problem parameters. It can be tuned by adjusting the cost matrix $Q = \begin{pmatrix} Q_1 & Q_{12} \\ Q_{12}^T & Q_2 \end{pmatrix}$ and the noise matrix $R = \begin{pmatrix} R_1 & R_{12} \\ R_{12}^T & R_2 \end{pmatrix}$.

Tuning of Q

Only in rare instances does a quadratic cost function follow directly from the design specification. In most cases, the cost function must be iteratively tuned by the designer to achieve the desired closed-loop behavior. One possible starting point is to only penalize the outputs $y = Cx$ and the inputs u. We can then put $Q_1 = C^TC$, $Q_2 = \rho I$, and $Q_{12} = 0$. For cases where the state variables have a physical interpretation, another option is let the diagonal elements be equal to the inverse value of the square of the allowed deviations:

$$Q_1 = \begin{pmatrix} \dfrac{1}{(x_1^{\max})^2} & \cdots & 0 \\ \vdots & \ddots & \\ 0 & & \dfrac{1}{(x_n^{\max})^2} \end{pmatrix}, \quad Q_2 = \begin{pmatrix} \dfrac{1}{(u_1^{\max})^2} & \cdots & 0 \\ \vdots & \ddots & \\ 0 & & \dfrac{1}{(u_m^{\max})^2} \end{pmatrix}, \quad Q_{12} = 0$$

Besides the above, some general guidelines can be given:

- To achieve higher bandwidth (more aggressive control), decrease Q_2 or increase Q_1.
- To increase the damping of a state x_j, add penalty on $\dot{x}_j^2$.
- (Advanced:) To make a state x_j behave more like $\dot{x}_j = -\alpha x_j$, add penalty on $(\dot{x}_j + \alpha x_j)^2$.

Note that last two options typically introduce cross-terms in the cost function.

EXAMPLE 11.2—LQ CONTROL OF THE FLEXIBLE SERVO
Consider linear-quadratic control of the flexible servo used in Lab 1 (see the lab manual for parameter values):

$$m_1 \frac{d^2 y_1}{dt^2} = -d_1 \frac{dy_1}{dt} - k(y_1 - y_2) + u(t)$$
$$m_2 \frac{d^2 y_2}{dt^2} = -d_2 \frac{dy_2}{dt} + k(y_1 - y_2)$$

Introducing the state vector $x = \begin{pmatrix} y_1 & \dot{y}_1 & y_2 & \dot{y}_2 \end{pmatrix}^T$, the plant can be written in state-space form as

$$\dot{x} = Ax + Bu$$
$$y_1 = C_1 x$$
$$y_2 = C_2 x$$

The impulse response of the open-loop system is shown in Figure 11.2(a). We would like to control the plant to make the output y_2 settle to zero within 2 seconds and without too much oscillations.

In the first design iteration, we penalize the output y_2 and u equally, setting $Q_1 = C_2^T C_2$ and $Q_2 = 1$. The resulting closed-loop response is shown in Figure 11.2(b). The system is a bit too fast—especially the control signal is out of bounds (± 10 for the real process).

We can slow the system down by increasing the penalty on u: In the second iteration we keep $Q_1 = C_2^T C_2$ and set $Q_2 = 100$. The new response is shown in Figure 11.3(a). The speed is now appropriate, but the oscillations need to be better damped.

In the third iteration, we add a penalty on $\dot{y}_2 = C_2 \dot{x} = C_2(Ax + Bu)$ to improve the damping. In this case $C_2 B = 0$, so the new weighting matrices are chosen as $Q_1 = C_2^T C_2 + 0.1(C_2 A)^T (C_2 A)$, $Q_2 = 100$. The closed-loop response is seen in Figure 11.3(b). The damping is now better, but since the relative size of Q_2 is now smaller, the controller is again a bit too aggressive. In a final iteration (not shown), the relative penalty on y_2, $\dot{y}_2$ and u can be fine-tuned. □

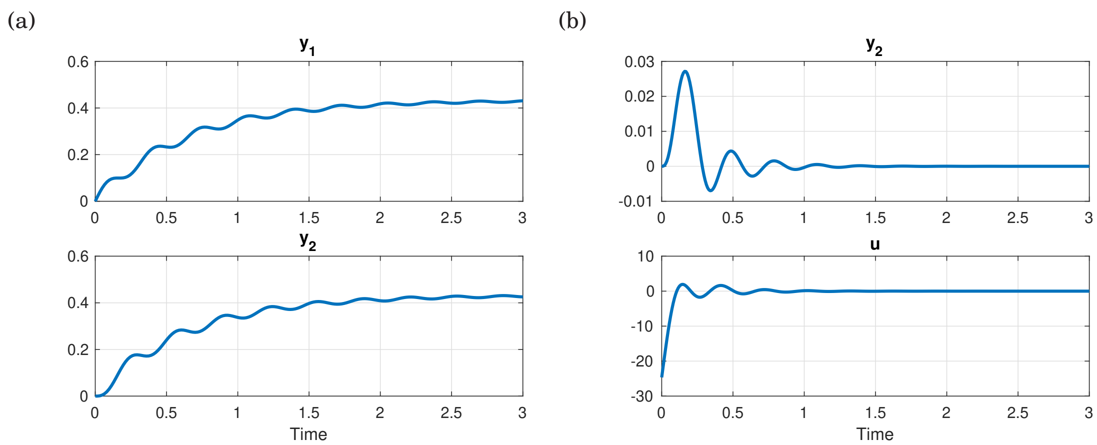

Figure 11.2 Flexible servo example: (a) Open-loop impulse response, (b) first design iteration.

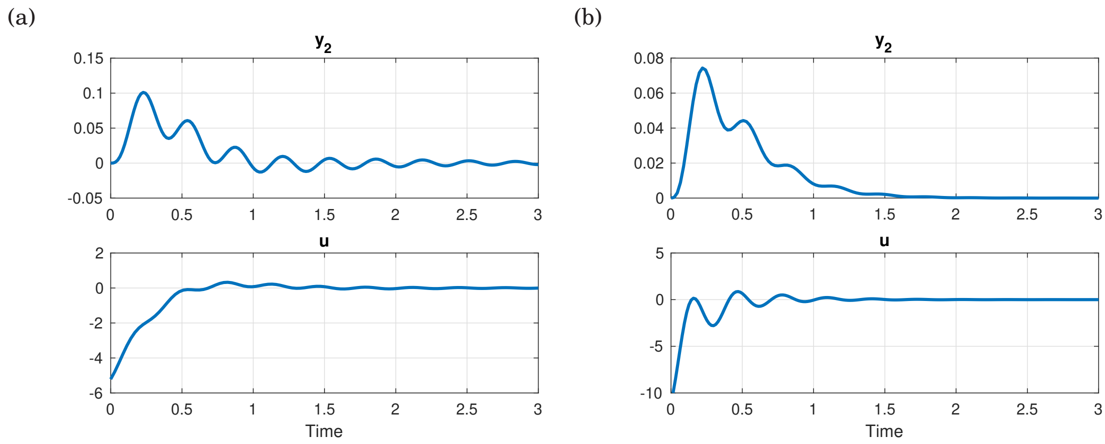

Figure 11.3 Flexible servo example: (a) Second design iteration, (b) third design iteration.

Tuning of R

When designing the Kalman filter, the properties of the real noise are seldom known. Noise can also used to model process uncertainty and nonlinearities, for which there is often not a good model available. As a starting point for tuning the filter, one can put $R_1 = BB^T$ (input noise) or $R_1 = I$ (noise on all states), $R_{12} = 0$ and $R_2 = \rho I$. If the controller is too sensitive to measurement noise, one should increase R_2 or, equivalently, decrease R_1. Making R_2 too large can however impair the robustness of the system, if the Kalman filter becomes too slow compared to the state feedback. The problem of robustness is discussed in the next section.

11.3 Robustness of LQG

Recall from Lecture 9 that LQ state feedback control has remarkable robustness properties. Under certain assumptions, the maximum sensitivity is guaranteed not to be larger than 1. Unfortunately, this property no longer holds when the LQ controller is combined with a Kalman filter into an LQG controller. This was famously shown in a 1978 paper by John Doyle with the title "Guaranteed Margins for LQG Regulators." The abstract plainly stated, "There are none." Even for plants without fundamental limitations, it is possible to construct examples where the stability margins become arbitrarily small.

EXAMPLE 11.3—[DOYLE & STEIN, 1979]
Consider the minimum-phase SISO plant

$$\dot{x} = \begin{pmatrix} -4 & -3 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix} u + \begin{pmatrix} 61 \\ -35 \end{pmatrix} w_1$$
$$y = \begin{pmatrix} 1 & 2 \end{pmatrix} x + w_2$$

where the white noise process $w = \begin{pmatrix} w_1 & w_2 \end{pmatrix}^T$ has unit intensity, $R = I$. An LQG controller is designed to minimize

$$J = \mathrm{E}\left(80\, x^T \begin{pmatrix} 1 & \sqrt{35} \\ \sqrt{35} & 35 \end{pmatrix} x + u^2\right)$$

The resulting closed-loop system is stable, with control poles in $-7 \pm 2i$ and observer poles in $-7.02 \pm 1.95i$. Plotting the Nyquist curve of the loop gain (see Figure 11.4(a)) reveals very poor stability margins: $M_s = 4.8$ and $\varphi_m = 14.8°$. In this example, the Kalman filter interferes with the state feedback to create an unrobust system. □

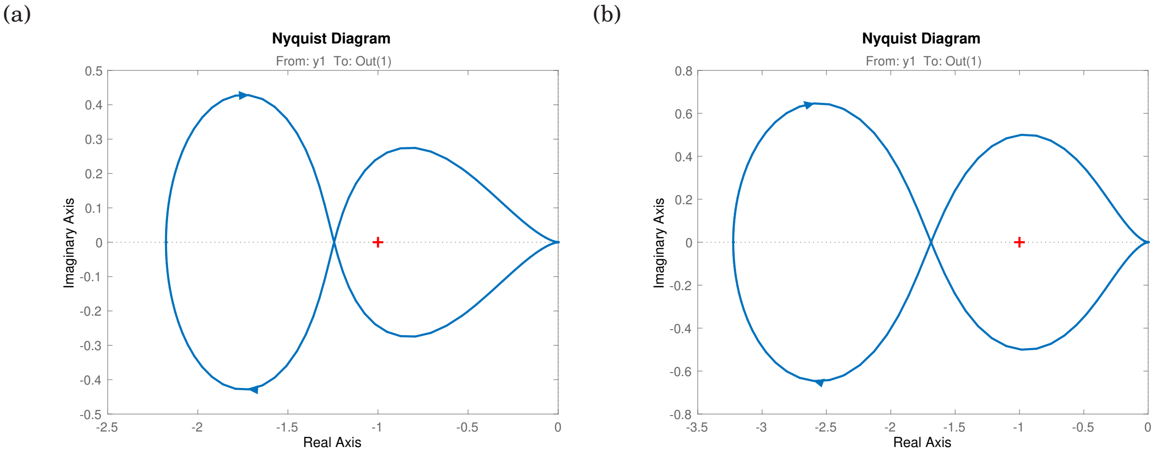

Figure 11.4 Doyle–Stein example: (a) Original example (b) Example with loop transfer recovery.

Loop transfer recovery

Much research has been spent on trying to improve the robustness of LQG. A design method known as *loop transfer recovery* (LTR) suggests that the robustness of an LQG controller can be improved by modifying the design weights Q or R. In *output LTR*, an extra penalty term proportional to C^TC is added to Q_1. In *input LTR*, an extra noise term proportional to BB^T is added to R_1. Both of these methods tries to make the loop transfer function more similar to the state feedback (LQ) case. The price is that the resulting controller will have higher gain, which gives more amplification of measurement noise and larger control signals.

Example 11.4—Doyle & Stein with LTR

Again consider the example from [Doyle & Stein, 1979]. To improve the robustness, we apply output LTR and modify the Q_1 matrix according to

$$Q_1^{new} = Q_1 + 400C^TC.$$

The resulting loop gain in shown in Figure 11.4(b). The Nyquist curve now avoids the critical point, giving better robustness ($M_s = 2.0$, $\varphi_m = 30°$). □

11.4 Integral action, reference values

In the previous lecture we introduced noise shaping as a way to introduce integral action in the Kalman filter (and hence also in the observer-based controller). Another option is to introduce integral action in the state feedback part of the controller. The idea is to extend the process model by adding explicit integrators as

$$\dot{x}_i = r - y$$

where r is the vector of reference signals and y is the vector of measured outputs. (If an output is not measurable, an estimate from the Kalman filter can be used instead.) This provides a nice way to introduce tracking of reference signals in the control system. With one integrator per output, the system can track reference signals without stationary error also in the presence of constant load disturbances.

Given a plant model $(A, B, C, 0)$, the methodology above gives an extended plant model

$$\begin{pmatrix} \dot{x} \\ \dot{x}_i \end{pmatrix} = \begin{pmatrix} A & 0 \\ -C & 0 \end{pmatrix} \begin{pmatrix} x \\ x_i \end{pmatrix} + \begin{pmatrix} B \\ 0 \end{pmatrix} u + \begin{pmatrix} 0 \\ I \end{pmatrix} r + \begin{pmatrix} I \\ 0 \end{pmatrix} w_1.$$

The Q_1 matrix is analogously extended as

$$Q_{1e} = \begin{pmatrix} Q_1 & 0 \\ 0 & Q_i \end{pmatrix},$$

where Q_i is the penalty on the integral states. We can then design an extended state feedback law using LQ theory:

$$u = -\begin{pmatrix} L & L_i \end{pmatrix} \begin{pmatrix} x \\ x_i \end{pmatrix}.$$

The control can also be extended with direct feedforward from r to further shape the command signal response.

Feedforward from the reference signal can also be used to introduce setpoint tracking without integral action. Introducing the control law

$$u = -Lx + L_r r,$$

the closed-loop transfer function is given by

$$G_{yr}(s) = C(sI - A + BL)^{-1} B L_r$$

Assuming a square plant (equal number of inputs and outputs), we can achieve a static closed-loop gain of I by selecting the feedforward gain as

$$L_r = [C(BL - A)^{-1} B]^{-1}$$

A reference filter to further shape $G_{yr}(s)$ can be added if needed.

Example 11.5—LQG control of DC-servo

As a final example, we develop a complete LQG controller for a simple DC servo model, including integral action and reference tracking. We will return to this example in future lectures. Assuming the basic control loop in Figure 11.5, the process is given by

$$P(s) = \frac{20}{s(s+1)},$$

and the white noise processes v_1 and v_2 are independent with intensities $R_1 = R_2 = 1$. The goal is to find the LQG controller that minimizes

$$J = \mathrm{E}\left(z^2 + u^2\right)$$

A state-space model of the plant is given by

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \overbrace{\begin{pmatrix} 0 & 0 \\ 1 & -1 \end{pmatrix}}^{A} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \overbrace{\begin{pmatrix} 20 \\ 0 \end{pmatrix}}^{B} u + \overbrace{\begin{pmatrix} 20 \\ 0 \end{pmatrix}}^{G} v_1$$

$$y = \underbrace{\begin{pmatrix} 0 & 1 \end{pmatrix}}_{C} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + v_2 \qquad z = x_2$$

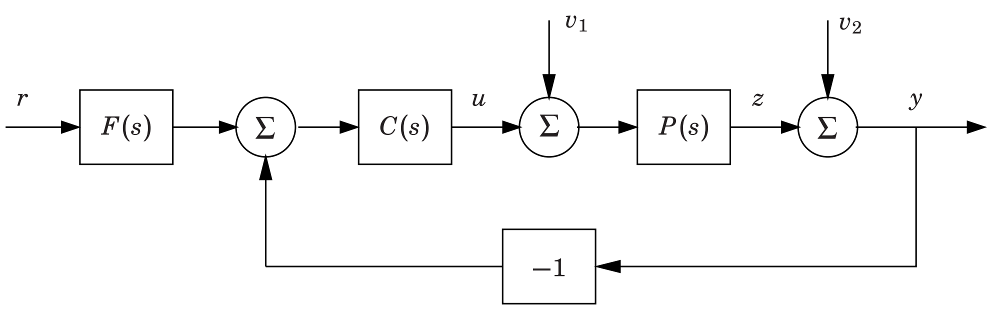

Figure 11.5 Basic control loop for LQG control of a DC servo.

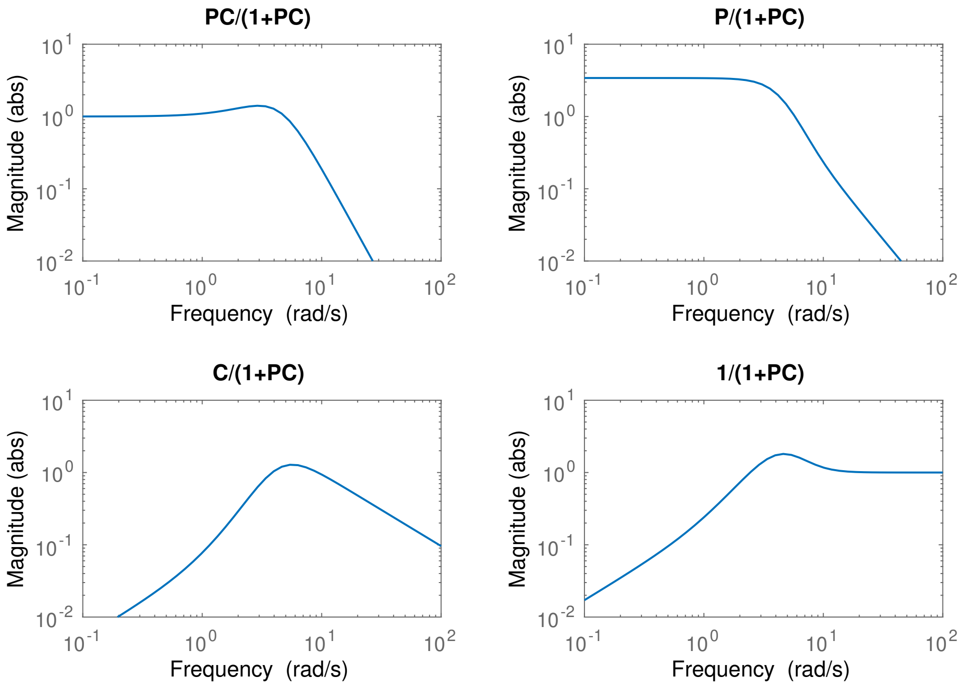

Figure 11.6 "Gang of four" for the LQG-controlled DC servo.

Solving the algebraic Riccati equations for the LQ state feedback and the Kalman filter in Matlab gives the optimal controller

$$\begin{aligned} \dot{\hat{x}} &= (A - BL - KC)\hat{x} + Ky \\ u &= -L\hat{x} \end{aligned}$$

with

$$L = \begin{pmatrix} 0.2702 & 0.7298 \end{pmatrix}, \qquad K = \begin{pmatrix} 20.0000 \\ 5.4031 \end{pmatrix}.$$

The result of any optimization-based controller design must always be carefully examined to see that it is reasonable. The "Gang of four" Bode magnitude diagrams for the LQG-controlled system are shown in Figure 11.6. It is seen that the bandwidth of the closed-loop system is close to 10 rad/s, and that the maximum sensitivities look reasonable. However, the closed-loop transfer function from plant disturbance to output, $P/(1+PC)$, reveals poor low-frequency disturbance rejection properties. To remedy this problem, we add integral action to the controller. Adding an explicit integrator, $\dot{x}_i = r - y$, we get the extended process model

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_i \end{pmatrix} = \overbrace{\begin{pmatrix} 0 & 0 & 0 \\ 1 & -1 & 0 \\ 0 & -1 & 0 \end{pmatrix}}^{A_e} \begin{pmatrix} x_1 \\ x_2 \\ x_i \end{pmatrix} + \overbrace{\begin{pmatrix} 20 \\ 0 \\ 0 \end{pmatrix}}^{B_e} u + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} r + \begin{pmatrix} 20 \\ 0 \\ 0 \end{pmatrix} v_1$$

Minimization of the extended cost function $\mathrm{E}\left(x_2^2 + Q_i x_i^2 + u^2\right)$ with $Q_i = 0.01$ then gives the optimal state feedback

$$u = -L_e \begin{pmatrix} \hat{x} & x_i \end{pmatrix}$$

where

$$L_e = \begin{pmatrix} 0.2751 & 0.7569 & -0.1 \end{pmatrix}$$

In the output feedback controller, we can use the same Kalman filter as before, since x_i is known by the controller. Plots of the new "Gang of four" Bode diagrams are shown in

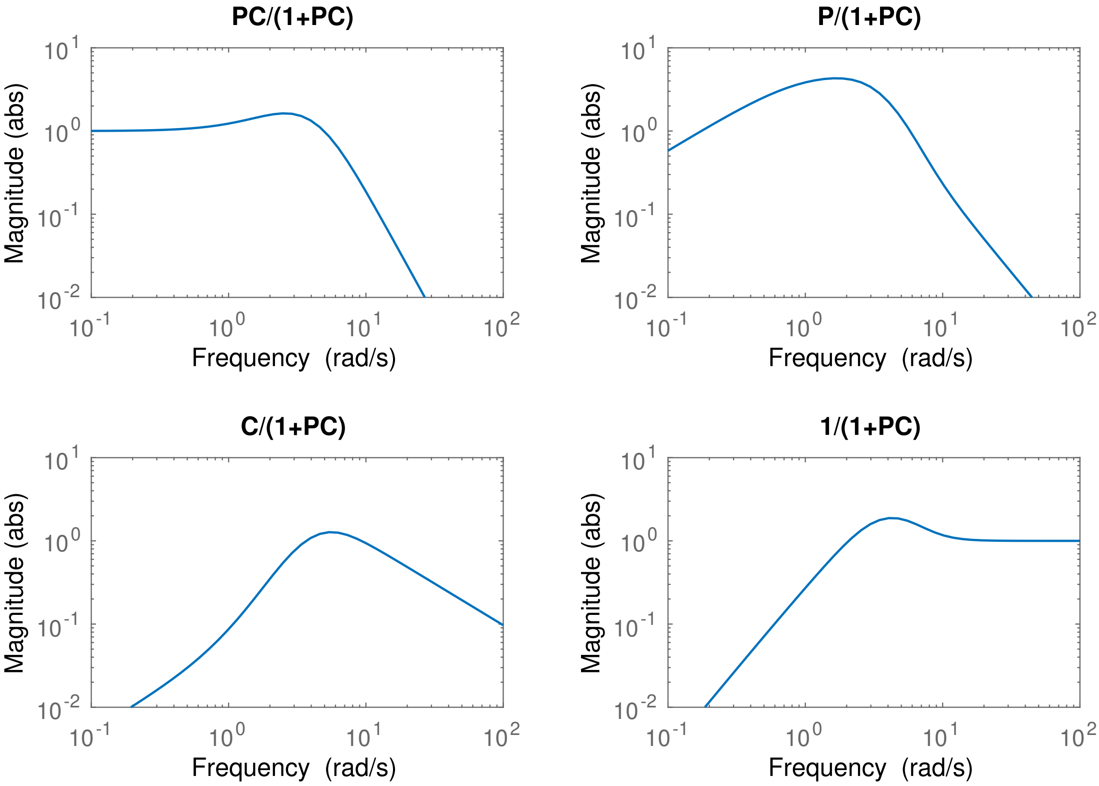

Figure 11.7 "Gang of four" for the LQG controller with integral action

Figure 11.7. It is seen that $P/(1+PC)$ now tends to zero for low frequencies, which shows that the controller has indeed integral action. The speed of integration, which influences both the disturbance rejection and the reference tracking speed, can be adjusted by tuning of the integral state penalty Q_i. □

11.5 Summary of LQG

The LQG methodology provides a systematic way to design stabilizing controllers for any linear plant model, including MIMO systems. The theory is well developed, and there exist analytical solutions as well as efficient software to produce the controllers. The observer structure ties nicely to reality if the states of the model have a correspondence in the physical system. Extensions with noise shaping, integral action and reference tracking are straightforward.

There are some potential disadvantages: First, a detailed state-space model of the plant is required. This can be very time-consuming to obtain. Second, the controller has the same order as the plant, which can lead to very high-order controllers if the model has many states. Third, it can be difficult to tune the cost and noise design weights. Often, several iterations are needed, together with insights of the physical process and its limitations. As usual, any fundamental limitations due to right half-plane zeros, poles, and time delays must be respected. Finally, the LQG controller has no guaranteed robustness, so the results must always be carefully checked.

The robustness problems with LQG led to the development of robust control theory in the 1980s. In the branch known as H_∞ optimal control, the objective is to minimize the maximum closed-loop gain from disturbances w to performance outputs z:

$$\text{Minimize} \quad \sup_\omega \|G_{zw}(i\omega)\|$$

This formulation ties closer to the shaping of the "Gang of four" closed-loop transfer functions and gives much better control of the stability margins M_s and M_t. Similarly to LQG, the solution can be obtained by solving a couple of algebraic Riccati equations.

In general, controller optimization can use a mix of objectives for the closed-loop system, including H_2 (LQG) and H_∞ criteria, as well as various constraints. This more general setting will be the topic of the final section of the course.