State observers.
Dynamic output Stabilisation

Elena VANNEAUX

elena.vanneaux@ensta-paris.fr

Course grade breakdowns

Labs - 50%

Final project - 50 %

State feedback design

Linear state space control theory involves modifying the behaviour of an m-input, p-output, n-state system

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \\ \mathbf{y}(t) &= \mathbf{C}\mathbf{x}(t), \end{aligned} \tag{OL}$$

which we call **the plant**, or **open loop state equation**, by application of a control law of the form

$$\mathbf{u}(t) = \mathbf{N}\mathbf{r}(t) - \mathbf{K}\mathbf{x}(t), \tag{U}$$

in which $\mathbf{r}(t)$ is the new (reference) input signal. The matrix $\mathbf{K}$ is the state feedback gain and $\mathbf{N}$ the feedforward gain.

Is Full State Feedback Always Available?

In a typical system, measurements are provided by sensors:

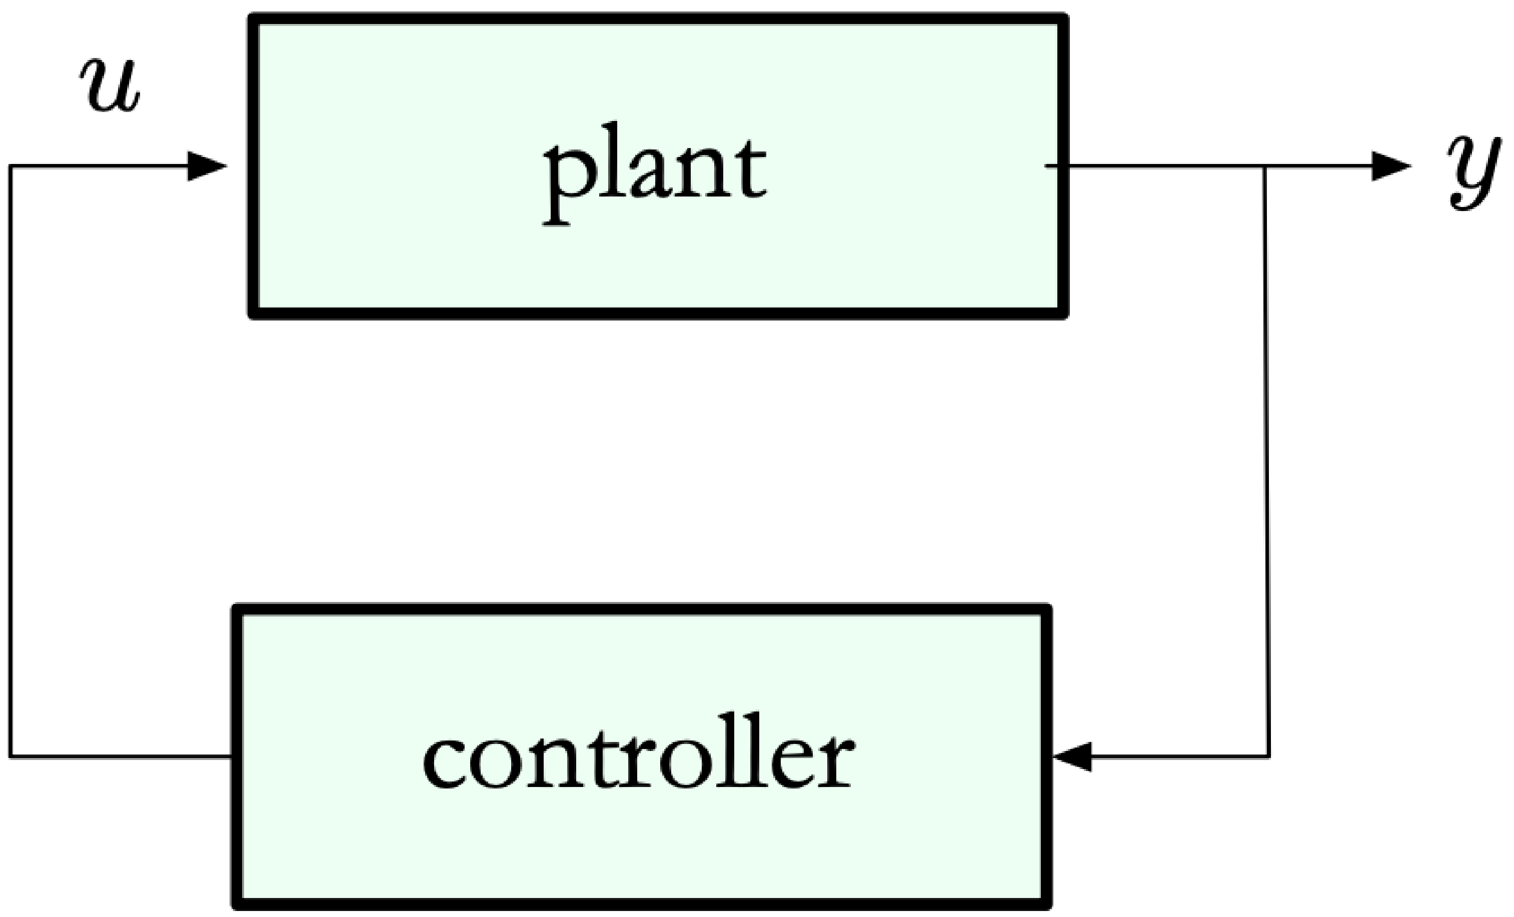

Full state feedback $u = -Kx$ is *not implementable*!!

Is Full State Feedback Always Available?

In a typical system, measurements are provided by sensors:

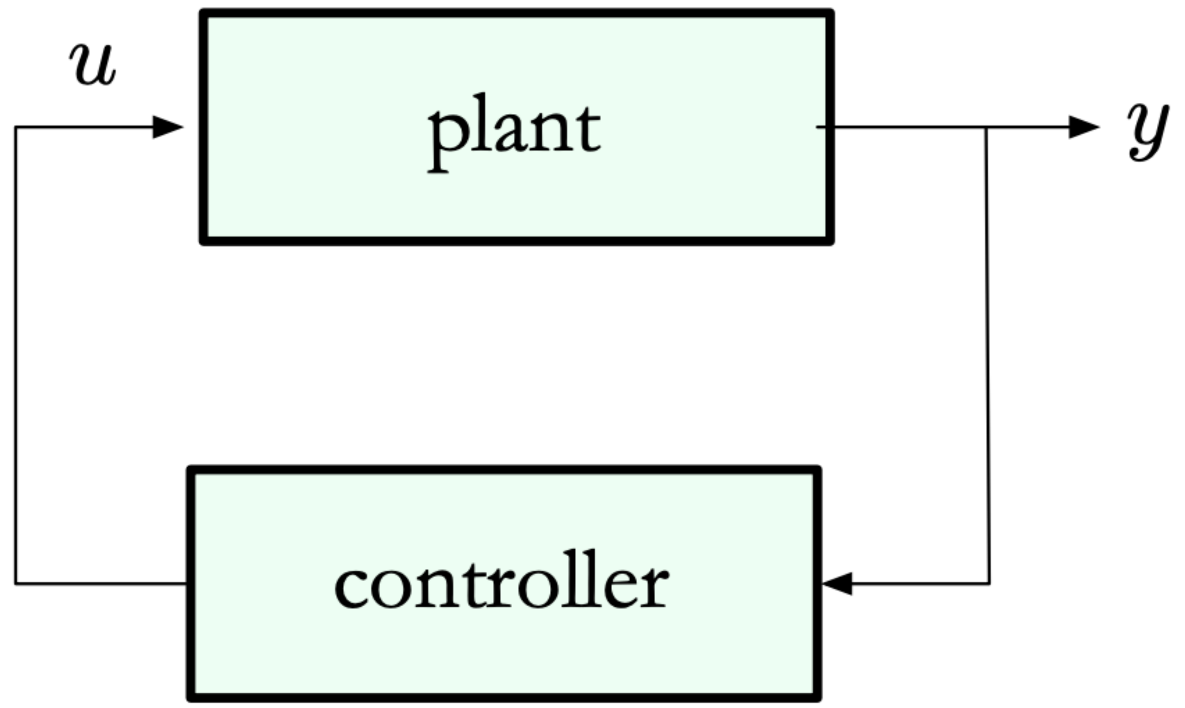

Full state feedback $u = -Kx$ is *not implementable*!!

When Full State Feedback Is Unavailable ...

... we need an Observer!!

Is Full State Feedback Always Available?

In a typical system, measurements are provided by sensors:

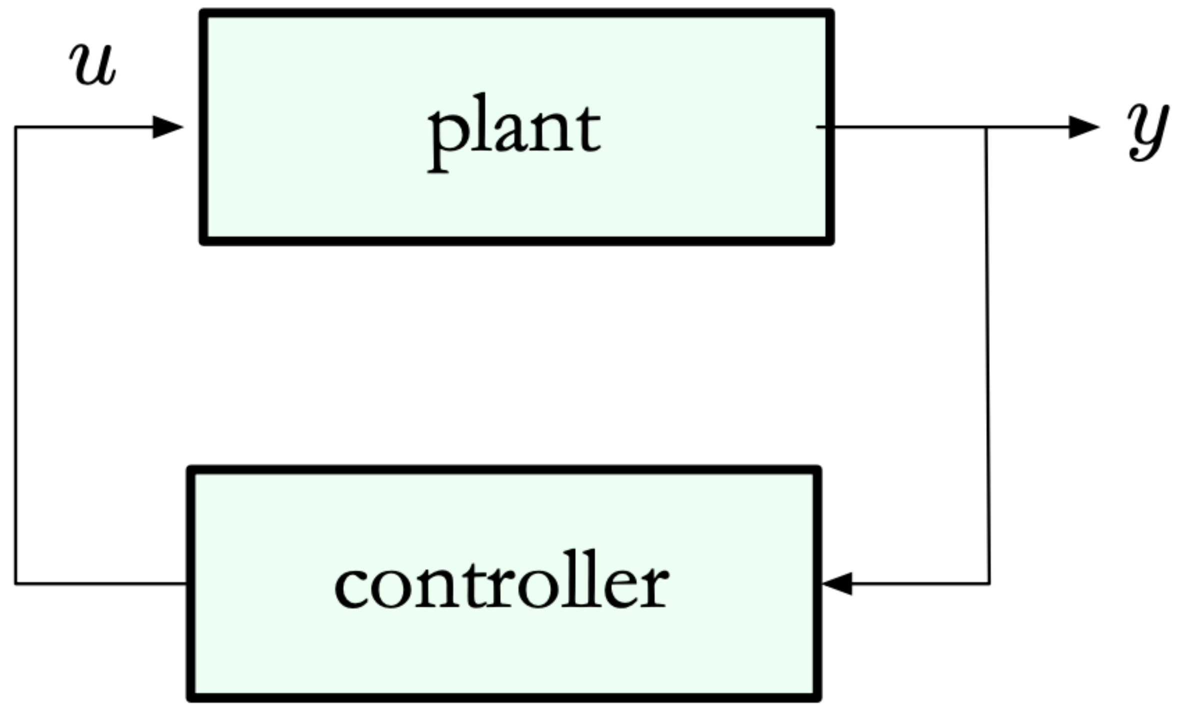

Full state feedback $u = -Kx$ is *not implementable*!!

In that case, an observer is used to estimate the state x:

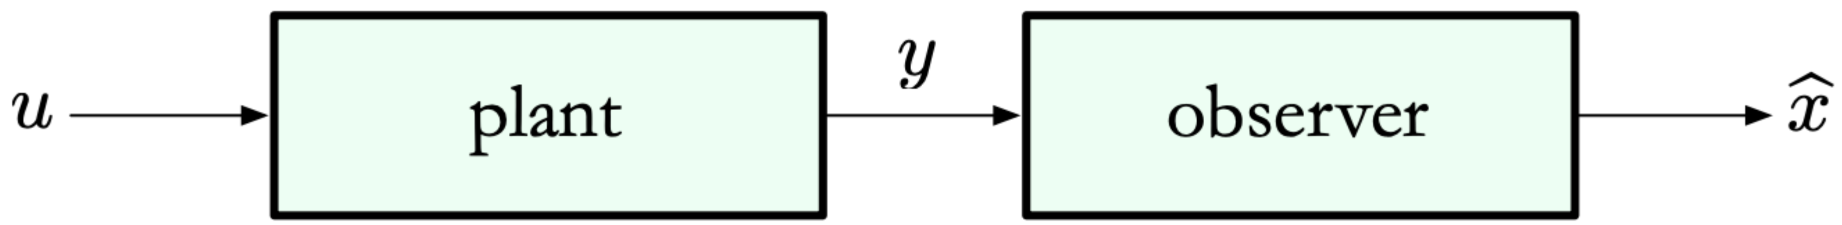

State Estimation Using an Observer

If the system is observable, the state estimate $\hat{x}$ is *asymptotically accurate*:

$$\|\hat{x}(t) - x(t)\| = \sqrt{\sum_{i=1}^{n} (\hat{x}_i(t) - x_i(t))^2} \xrightarrow{t\to\infty} 0$$

If we are successful, then we can try estimated state feedback:

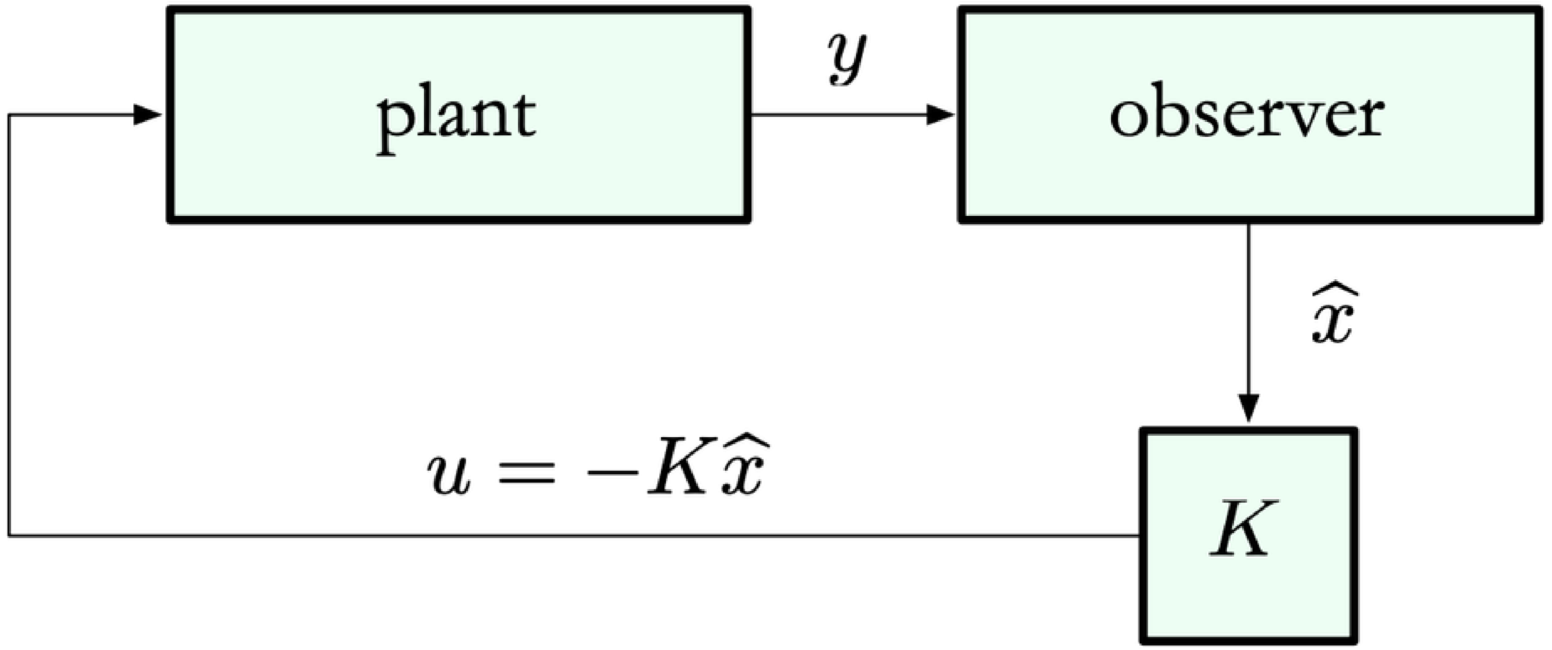

The Luenberger Observer

$$\text{System:} \quad \dot{x} = Ax$$
$$y = Cx$$
$$\text{Observer:} \quad \dot{\widehat{x}} = (A - LC)\widehat{x} + Ly.$$

What happens to state estimation error $e = x - \widehat{x}$ as $t \to \infty$?

$$\begin{aligned} \dot{e} &= \dot{x} - \dot{\widehat{x}} \\ &= Ax - [(A - LC)\widehat{x} + LCx] \\ &= (A - LC)x - (A - LC)\widehat{x} \\ &= (A - LC)e \end{aligned}$$

Does $e(t)$ converge to zero in some sense?

The Luenberger Observer

$$\begin{aligned} \text{System:} \quad & \dot{x} = Ax \\ & y = Cx \\ \text{Observer:} \quad & \dot{\widehat{x}} = (A - LC)\widehat{x} + Ly \\ \text{Error:} \quad & \dot{e} = (A - LC)e \end{aligned}$$

Recall our assumption that $A - LC$ is Hurwitz (all eigenvalues are in LHP). This implies that

$$\|x(t) - \widehat{x}(t)\|^2 = \|e(t)\|^2 = \sum_{i=1}^{n} |e_i(t)|^2 \xrightarrow{t \to \infty} 0$$

at an exponential rate, determined by the eigenvalues of $A - LC$.

For fast convergence, want eigenvalues of $A - LC$ far into LHP!!

Observability and Estimation Error

Fact: If the system

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

is observable, then we can arbitrarily assign eigenvalues of $A - LC$ by a suitable choice of the output injection matrix L.

This is similar to the fact that controllability implies arbitrary closed-loop pole placement by state feedback.

In fact, these two facts are closely related because CCF is dual to OCF.

Controllability–Observability Duality

Claim: The system

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

is observable if and only if the system

$$\dot{x} = A^T x + C^T u$$

is controllable.

Proof:

$$\mathcal{C}(A^T, C^T) = [C^T \,|\, A^T C^T \,|\, \ldots \,|\, (A^T)^{n-1} C^T]$$

$$= \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}^T = [\mathcal{O}(A, C)]^T$$

Thus, $\mathcal{O}(A, C)$ is nonsingular if and only if $\mathcal{C}(A^T, C^T)$ is.

Observer Pole Placement, O/C Duality Version

Given an observable pair (A, C):

1. For $F = A^T$, $G = C^T$, consider the system $\dot{x} = Fx + Gu$ (this system is controllable).
2. Use our earlier procedure to find K, such that

$$F - GK = A^T - C^T K$$

   has desired eigenvalues.
3. Then

$$\text{eig}(A^T - C^T K) = \text{eig}(A^T - C^T K)^T = \text{eig}(A - K^T C),$$

   so $L = K^T$ is the desired output injection matrix.

Final answer: use the observer

$$\begin{aligned}\dot{\hat{x}} &= (A - LC)\hat{x} + Ly \\ &= (A - K^T C)\hat{x} + K^T y.\end{aligned}$$

Recall: infinite-horizon Linear Quadratic Regulator (LQR)

Problem formulation: optimal control for integral-quadratic cost

$$\underset{u(t)}{\text{minimize}}\; J(u(t)) = \int_0^\infty x(t)^\top Q\, x(t) + u(t)^\top R\, u(t)\, dt$$

$$\text{subject to}\;\; \dot{x}(t) = Ax(t) + Bu(t)\,,\;\; x(0) = x_0$$

Feasible if $Q \succeq 0$, $R \succ 0$, (A, B) stabilizable, & $(A, Q^{1/2})$ detectable.

Solution: independent of the initial condition x_0, the linear state feedback

$$u^\star(t) = -\,K^\star x(t) = -\,\underbrace{R^{-1}B^\top P}_{=K^\star}\,x$$

where $P \succ 0$ solves the algebraic matrix Riccati equation

$$A^\top P + PA + Q = PBR^{-1}B^\top P$$

Equivalent problem formulation:

In hindsight, LQR can be interpreted as *optimal pole placement* for

$$\dot{x}(t) = (A - BK^\star)\, x(t)$$

trading off minimal state deviation and minimal control energy:

$$K^\star = \text{argmin}_K \int_0^\infty \underbrace{x(t)^\top Q\, x(t)}_{\text{state deviation}} + \underbrace{x(t)^\top K^\top R K\, x(t)}_{\text{control energy}}\, dt$$

Recall: dual notions of controllability/observability

The following statements are equivalent for **controller design**:

1. the system (A, B) is controllable
2. the controllability matrix

   $$W_c = [B \; AB \; \dots \; A^{n-1}B]$$

   has full rank n
3. the eigenvalues of $A - BK$ can be assigned via the matrix K
4. ...

The following statements are equivalent for **observer design**:

1. the system (A, C) is observable
2. the observability matrix

   $$W_o = \begin{bmatrix} C \\ CA \\ \dots \\ CA^{n-1} \end{bmatrix}$$

   has full rank n
3. the eigenvalues of $A - LC$ can be assigned via the matrix L
4. ...

idea: use duality $(A, B, K) \leftrightarrow (A^\top, C^\top, L^\top)$ to design *optimal observers*

Optimal design by duality $(A, B, K) \leftrightarrow (A^\top, C^\top, L^\top)$

LQ-optimal control for closed-loop dynamics: $\dot{x} = (A - BK)\,x$

$$\text{minimize}_K \int_0^\infty \underbrace{x(t)^\top Q\, x(t)}_{\text{state deviation}} + \underbrace{x(t)^\top K^\top R K\, x(t)}_{\text{control energy}}\, dt$$

$\Rightarrow K^\star = R^{-1} B^\top P$ where $P \succ 0$ solves $A^\top P + PA + Q = PBR^{-1}B^\top P$

LQ-optimal estimation for estimation error dynamics: $\dot{\epsilon} = (A - LC)\,\epsilon$

$$\text{minimize}_L \int_0^\infty \underbrace{\epsilon(t)^\top Q\, \epsilon(t)}_{\text{estimation error}} + \underbrace{\epsilon(t)^\top L R L^\top\, \epsilon(t)}_{\text{output correction}}\, dt$$

$\Rightarrow L^\star = PC^\top R^{-1}$ where $P \succ 0$ solves $AP + PA^\top + Q = PC^\top R^{-1} C P$

Role of Q and R in LQ observer design

Optimal observer for integral-quadratic cost

$$\text{minimize}_L \int_0^\infty \underbrace{\epsilon(t)^\top Q\,\epsilon(t)}_{\text{estimation error}} + \underbrace{\epsilon(t)^\top LRL^\top\,\epsilon(t)}_{\text{output correction}}\,dt$$

trades off **prediction and correction** $\dot{\hat{x}} = \underbrace{A\hat{x} + Bu}_{\text{prediction}} + \underbrace{L(y - C\hat{x})}_{\text{correction}}$

$\Rightarrow$ $R \succ 0$ quantifies correction through measurement:

R "large" $\implies$ L "small" $\implies$ trust prediction

R "small" $\implies$ L "large" $\implies$ trust measurement

$\Rightarrow$ $Q \succeq 0$ quantifies prediction error: Q "large" $\implies$ "smaller" error ϵ

Summary: LQ optimal estimation (LQE)

Problem formulation: optimal observer for integral-quadratic cost

$$\text{minimize}_L \int_0^\infty \underbrace{\epsilon(t)^\top Q\,\epsilon(t)}_{\text{estimation error}} + \underbrace{\epsilon(t)^\top LRL^\top \epsilon(t)}_{\text{output correction}}\,dt$$

$$\text{subject to}\quad \dot{\epsilon}(t) = (A - LC)\,\epsilon(t)$$

Feasible if $Q \succeq 0$, $R \succ 0$, (A, C) detectable, & $(A, Q^{1/2})$ stabilizable.

Solution: independent of the initial condition ϵ_0, the output feedback

$$L^\star = PC^\top R^{-1}$$

where $P \succ 0$ solves the algebraic matrix Riccati equation

$$AP + PA^\top + Q = PC^\top R^{-1} CP$$

Combining Full-State Feedback with an Observer

- So far, we have focused on autonomous systems ($u = 0$).
- What about nonzero inputs?

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

Combining Full-State Feedback with an Observer

- So far, we have focused on autonomous systems ($u = 0$).
- What about nonzero inputs?

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

— assume (A, B) is controllable and (A, C) is observable.

- Today, we will learn how to use an observer together with estimated state feedback to (approximately) place closed-loop poles.

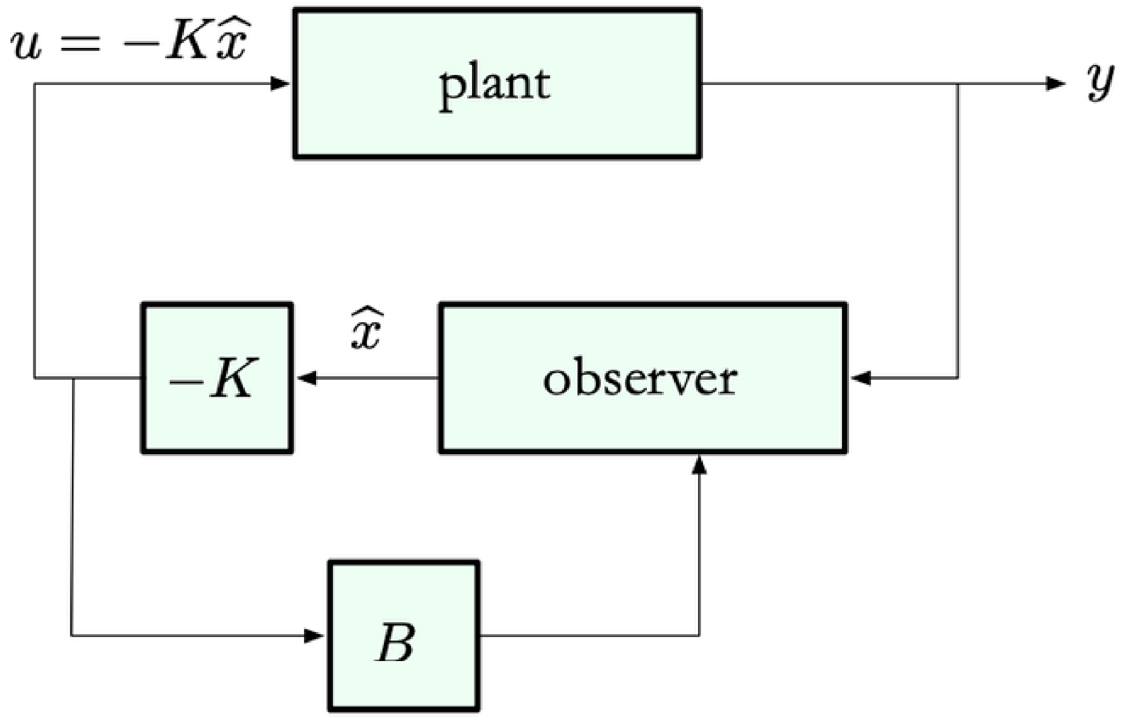

Combining Full-State Feedback with an Observer

- Consider

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

  where (A, B) is controllable and (A, C) is observable.
- We know how to find K, such that $A - BK$ has desired eigenvalues (controller poles).
- Since we do not have access to x, we must design an observer. But this time, we need a slight modification because of the Bu term.

Observer in the Presence of Control Input

- Let's see what goes wrong when we use the old approach:

$$\dot{\widehat{x}} = (A - LC)\widehat{x} + Ly$$

- For the estimation error $e = x - \widehat{x}$, we have

$$\begin{aligned} \dot{e} &= \dot{x} - \dot{\widehat{x}} \\ &= Ax + Bu - [(A - LC)\widehat{x} + LCx] \\ &= (A - LC)e + Bu \qquad \text{– not good} \end{aligned}$$

- Idea: since u is a signal we can access, let's use it as an input to the observer to cancel the Bu term from $\dot{x}$.
- Modified observer:

$$\begin{aligned} \dot{\widehat{x}} &= (A - LC)\widehat{x} + Ly + Bu \\ \dot{e} &= \dot{x} - \dot{\widehat{x}} \\ &= Ax + Bu - [(A - LC)\widehat{x} + LCx + Bu] \\ &= (A - LC)e \qquad \text{regardless of } u \end{aligned}$$

Observer and Controller

$$\begin{aligned} \text{System:} \quad & \dot{x} = Ax + Bu \\ & y = Cx \\ \text{Observer:} \quad & \dot{\widehat{x}} = (A - LC)\widehat{x} + Ly + Bu \\ \text{Error:} \quad & \dot{e} = (A - LC)e \end{aligned}$$

- By observability, we can arbitrarily assign $\text{eig}(A - LC)$; these should be farther into LHP than desired controller poles.

$$\text{Controller:} \quad u = -K\widehat{x} \quad (\text{estimated state feedback})$$

- By controllability, we can arbitrarily assign $\text{eig}(A - BK)$.

Observer and Controller

$$
\begin{aligned}
\text{System:} \quad & \dot{x} = Ax + Bu \\
& y = Cx \\
\text{Observer:} \quad & \dot{\widehat{x}} = (A - LC)\widehat{x} + Ly + Bu \\
\text{Controller:} \quad & u = -K\widehat{x}
\end{aligned}
$$

The overall observer-controller system is:

$$
\begin{aligned}
\dot{\widehat{x}} &= (A - LC)\widehat{x} + Ly + B\underbrace{(-K\widehat{x})}_{=u} \\
&= (A - LC - BK)\widehat{x} + Ly \\
u &= -K\widehat{x} \qquad \text{(dynamic output feedback)}
\end{aligned}
$$

— this is a dynamical system with input y and output u

Dynamic Output Feedback

$$\dot{x} = Ax + Bu$$
$$y = Cx$$
$$\dot{\widehat{x}} = (A - LC - BK)\widehat{x} + Ly$$
$$u = -K\widehat{x}$$

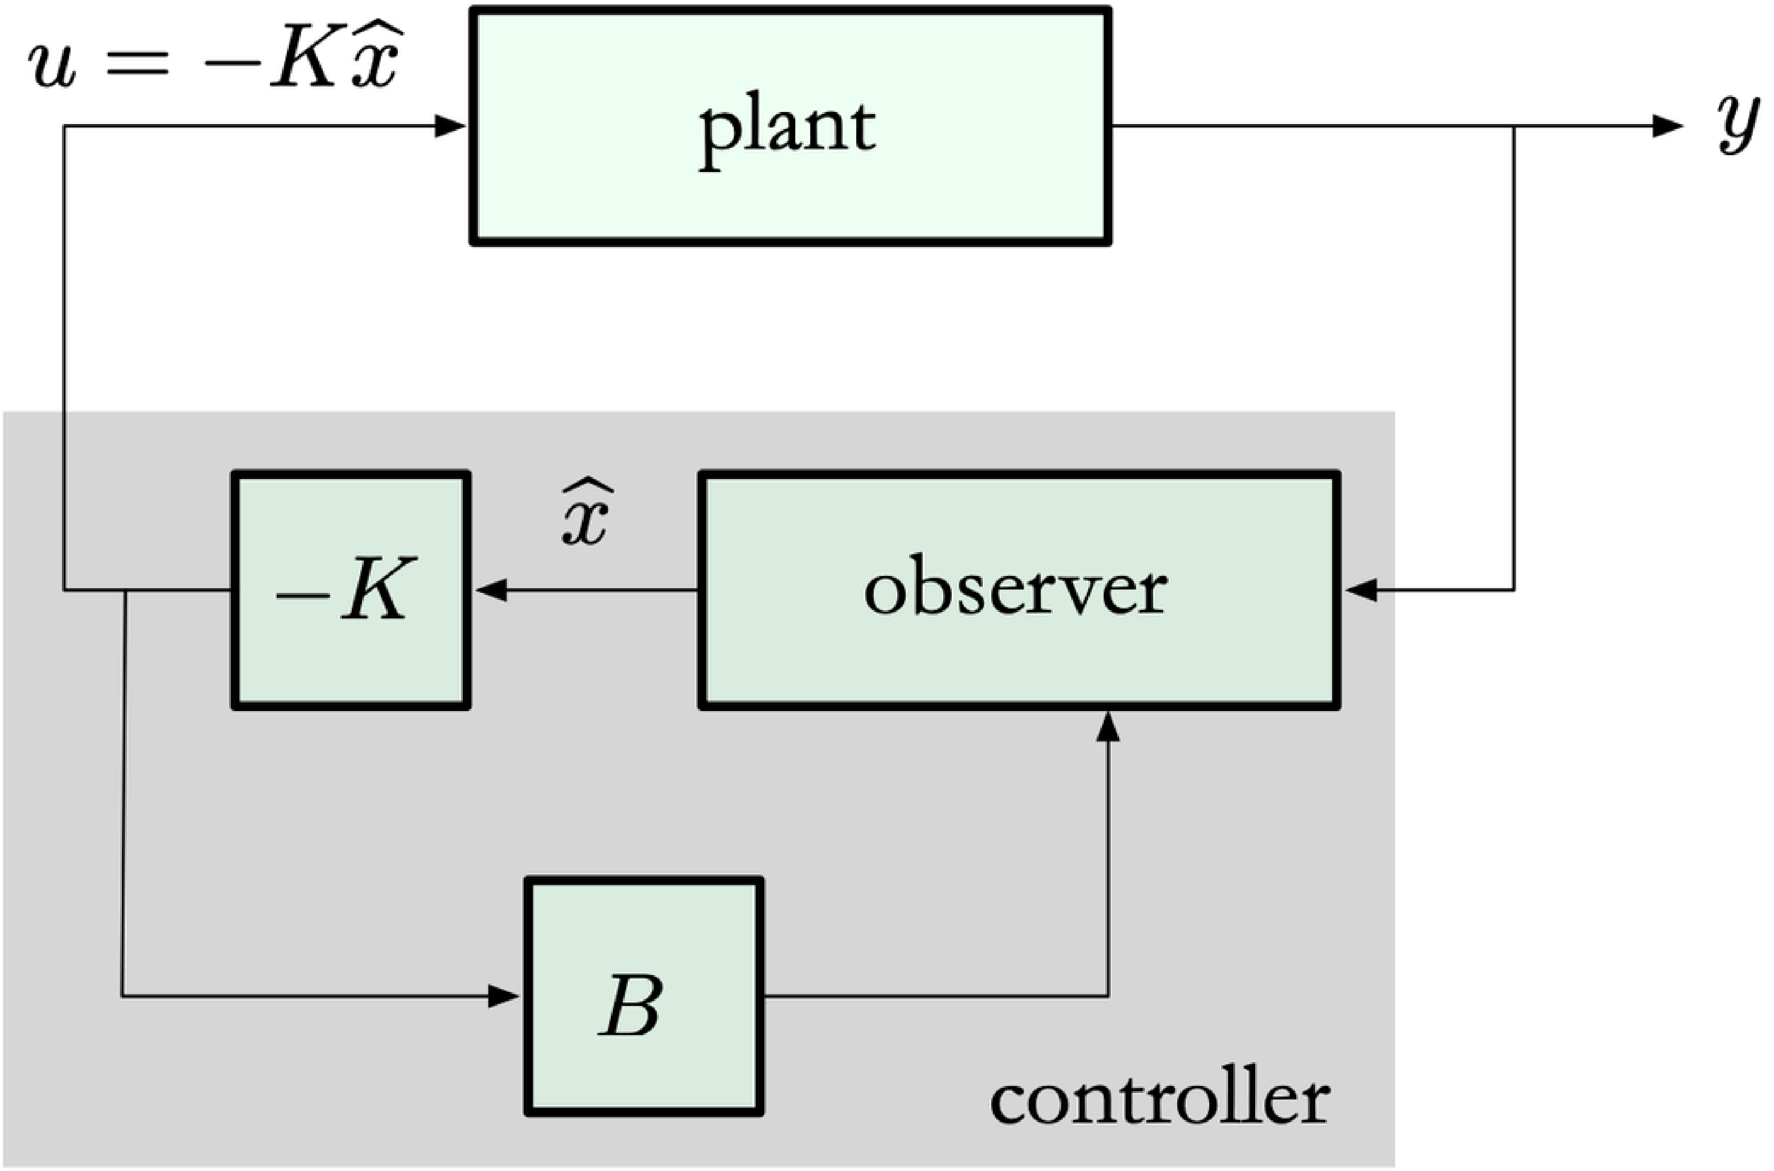

Dynamic Output Feedback: Does It Work?

Summarizing:

- When $y = x$, full state feedback $u = -Kx$ achieves desired pole placement.
- How do we know that $u = -K\widehat{x}$ achieves similar objectives?

Here is our overall closed-loop system:

$$\dot{x} = Ax - BK\widehat{x}$$
$$\dot{\widehat{x}} = (A - LC - BK)\widehat{x} + LCx$$

We can write it in block matrix form:

$$\begin{pmatrix} \dot{x} \\ \dot{\widehat{x}} \end{pmatrix} = \begin{pmatrix} A & -BK \\ LC & A - LC - BK \end{pmatrix} \begin{pmatrix} x \\ \widehat{x} \end{pmatrix}$$

How do we relate this to “nominal” behavior, $A - BK$?

Dynamic Output Feedback

$$\begin{pmatrix} \dot{x} \\ \dot{\widehat{x}} \end{pmatrix} = \begin{pmatrix} A & -BK \\ LC & A - LC - BK \end{pmatrix} \begin{pmatrix} x \\ \widehat{x} \end{pmatrix}$$

Let us transform to new coordinates:

$$\begin{pmatrix} x \\ \widehat{x} \end{pmatrix} \longmapsto \begin{pmatrix} x \\ e \end{pmatrix} = \begin{pmatrix} x \\ x - \widehat{x} \end{pmatrix} = \underbrace{\begin{pmatrix} I & 0 \\ I & -I \end{pmatrix}}_{T} \begin{pmatrix} x \\ \widehat{x} \end{pmatrix}$$

Two key observations:

- T is invertible, so the new representation is equivalent to the old one
- in the new coordinates, we have

$$\begin{aligned} \dot{x} &= Ax - BK\widehat{x} \\ &= (A - BK)x + BK(x - \widehat{x}) \\ &= (A - BK)x + BKe \\ \dot{e} &= (A - LC)e \end{aligned}$$

The Main Result: Separation Principle

So now we can write

$$\begin{pmatrix} \dot{x} \\ \dot{e} \end{pmatrix} = \underbrace{\begin{pmatrix} A - BK & BK \\ 0 & A - LC \end{pmatrix}}_{\text{upper triangular matrix}} \begin{pmatrix} x \\ e \end{pmatrix}$$

The closed-loop characteristic polynomial is

$$\det \begin{pmatrix} Is - A + BK & -BK \\ 0 & Is - A + LC \end{pmatrix}$$
$$= \det(Is - A + BK) \cdot \det(Is - A + LC)$$

Separation principle. The closed-loop eigenvalues are:

$$\{\text{controller poles (roots of } \det(Is - A + BK))\}$$
$$\cup \{\text{observer poles (roots of } \det(Is - A + LC))\}$$

— this holds only for linear systems!!

Separation Principle

Separation principle. The closed-loop eigenvalues are:

$$\{\text{controller poles (roots of } \det(Is - A + BK))\}$$
$$\cup \{\text{observer poles (roots of } \det(Is - A + LC))\}$$

— this holds only for linear systems!!

Moral of the story:

- If we choose observer poles to be several times faster than the controller poles (e.g., 2–5 times), then the controller poles will be dominant.
- Dynamic output feedback gives essentially the same performance as (nonimplementable) full-state feedback — provided observer poles are far enough into LHP.
- Remember: the system must be controllable and observable!!

Let’s go
a little bit out of the scope
of this class…

Stochastic model of plant

$$
\begin{aligned}
\dot{x}(t) &= Ax(t) + Bu(t) + w(t) \\
y(t) &= Cx(t) + v(t)
\end{aligned}
$$

Assumption: Gaussian, white, & uncorrelated $\left(E\left[w(t)v(\tau)^\top\right] = 0\right)$ noise:

$w(t)$: process noise with zero mean $E[w(t)] = 0$ & covariance:

$$E[w(t)w(\tau)^\top] = Q\,\delta(t-\tau)\,, \quad Q \succeq 0$$

$v(t)$: measurement noise with zero mean $E[v(t)] = 0$ & covariance:

$$E[v(t)v(\tau)^\top] = R\,\delta(t-\tau)\,, \quad R \succ 0$$

Kalman's question: What is the estimate $\widehat{x}(t)$ that minimizes the mean square error $E\left[(x(t)-\widehat{x}(t))^\top(x(t)-\widehat{x}(t))\right]$, given prior measurements? And how can we find an uncertainty estimate, e.g., a covariance matrix?

Kalman's solution

(see Wikipedia for controversial history & earlier/parallel authors: Bucy & Stratonovich)

Assume $Q \succeq 0$, $R \succ 0$, (A, C) detectable, $(A, Q^{1/2})$ stabilizable,
then the **mean-square optimal state estimator is a linear observer**

$$\boxed{\dot{\hat{x}} = A\hat{x} + Bu + L^\star(y - C\hat{x})}$$

where $\boxed{L^\star = PC^\top R^{-1}}$ is the **Kalman gain** & $P = P^\top \succ 0$ is the unique positive definite solution to the **algebraic matrix Riccati equation**

$$\boxed{0 = AP + PA^\top - PC^\top R^{-1} CP + Q},$$

& $P = \lim_{t \to \infty} E\left[\epsilon(t)\epsilon(t)^\top \middle| y(\tau),\, 0 \le \tau \le t\right]$ is **stationary error covariance**.

Properties of the Kalman Filter

- The plant (A, B, C, D) is not required to be stable.
- The Kalman filter results in **unbiased state estimate**: $E[\widehat{x}] = E[x]$.
  Further, the residual random process $y(t) - C\widehat{x}(t)$ is white,
  & therefore has no remaining dynamic information content.
  $\rightarrow$ "no need to work harder" (in case of Gaussian noise & LTI system)
- Kalman filter is **optimal** & minimizes stationary mean square error
  $$\lim_{t\to\infty} E\left[(x(t) - \widehat{x}(t))^\top (x(t) - \widehat{x}(t))\right] = \lim_{t\to\infty} E\left[\epsilon(t)^\top \epsilon(t)\right]$$
  given prior measurements $y(\tau)$ with $\tau \leq t$
- **stationary error covariance** $P = \lim_{t\to\infty} E\left[\epsilon(t)\epsilon(t)^\top \middle| y(\tau),\ 0 \leq \tau \leq t\right]$
- The Kalman filter is **recursive** & fully determined by P & $\widehat{x}(t)$.
- **duality** to LQR (all carries over) & **separation principle** for observer

LQG (linear-quadratic-Gaussian) control design

For a LTI control system with state and measurement noise

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) + w(t) \\ y(t) &= Cx(t) + v(t) \end{aligned}$$

we want to find the **output-feedback controller** that minimizes

$$E\left[\int_0^\infty x(t)^\top Q x(t) + u(t)^\top R(t) u(t) dt\right]$$

→ **solution:** combination of LQR state feedback & Kalman filter

separation principle

LQG design = LQR state feedback + time-invariant Kalman filter

LQR and Kalman gains can be designed separately

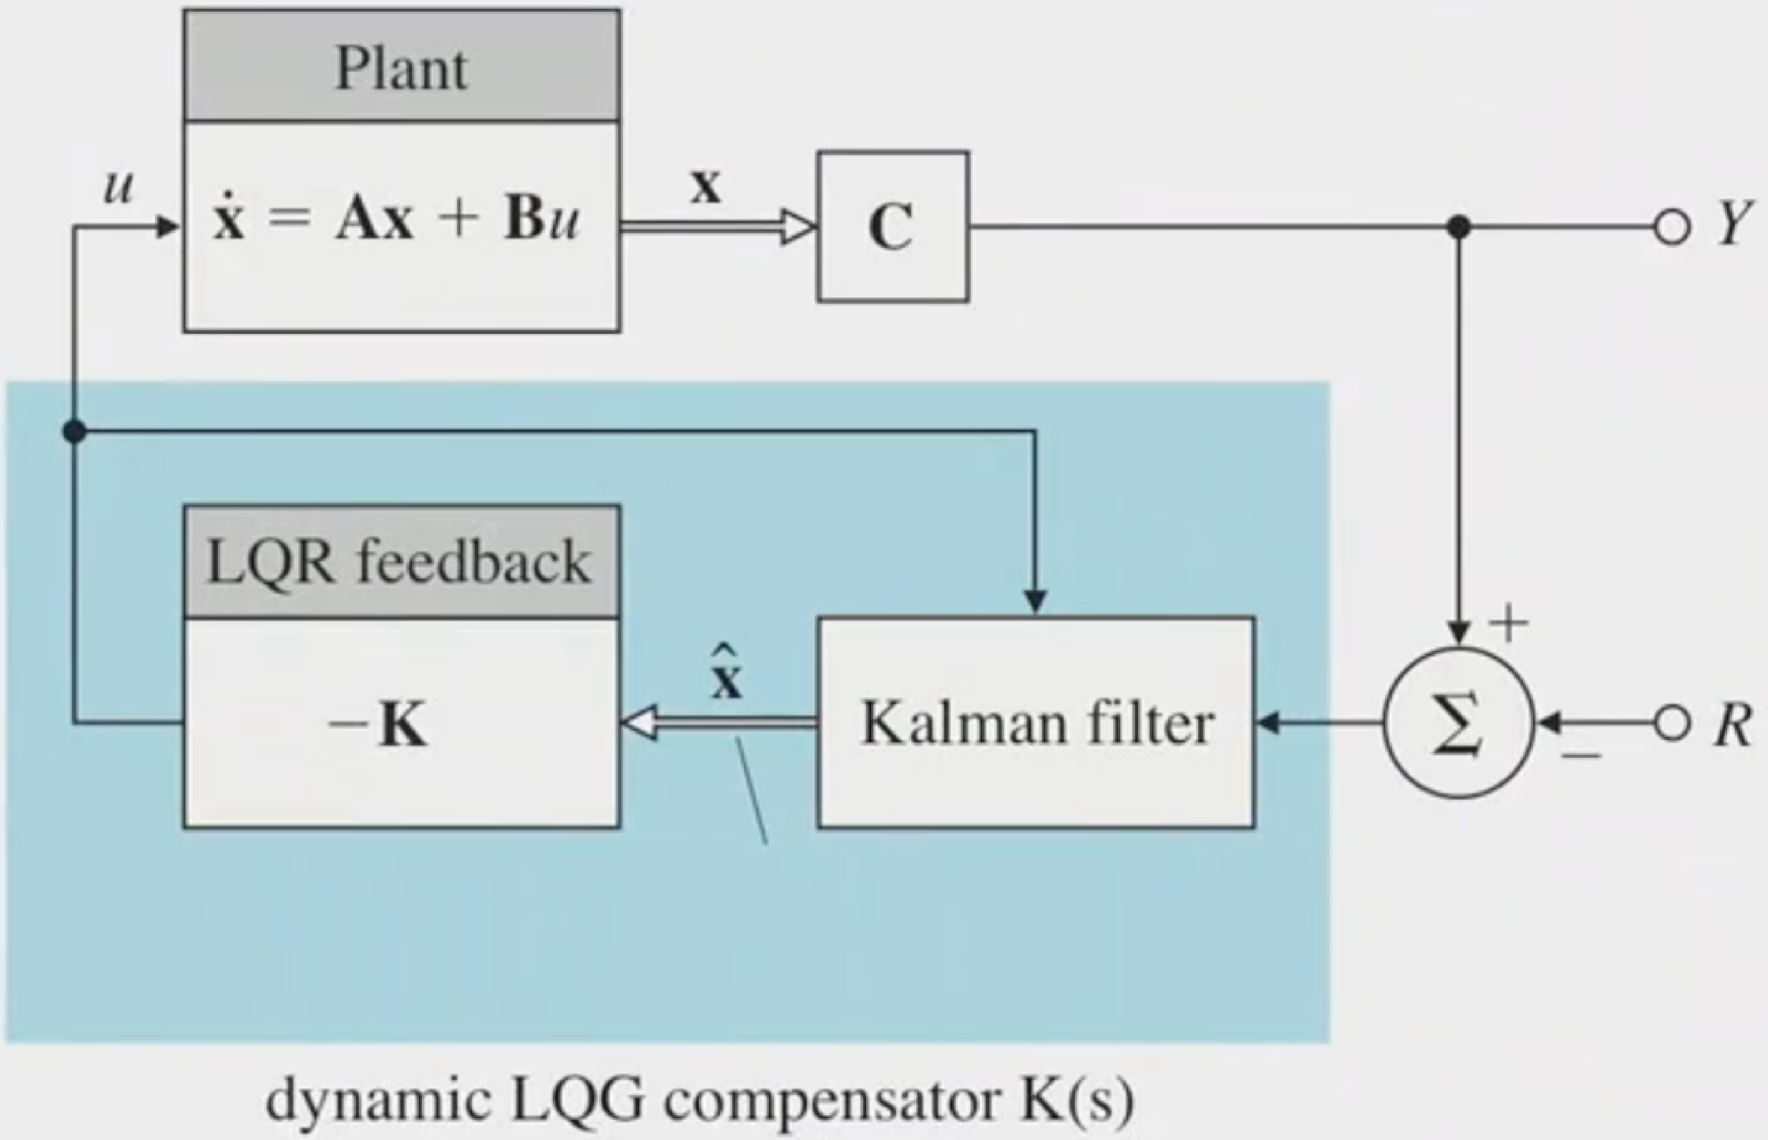

1 **state feedback:** $u = -K\hat{x}$ via LQR

2 **state estimation:** $x \approx \hat{x}$ via Kalman filter

LQG design is optimal but not robust in general

try the calculation at home

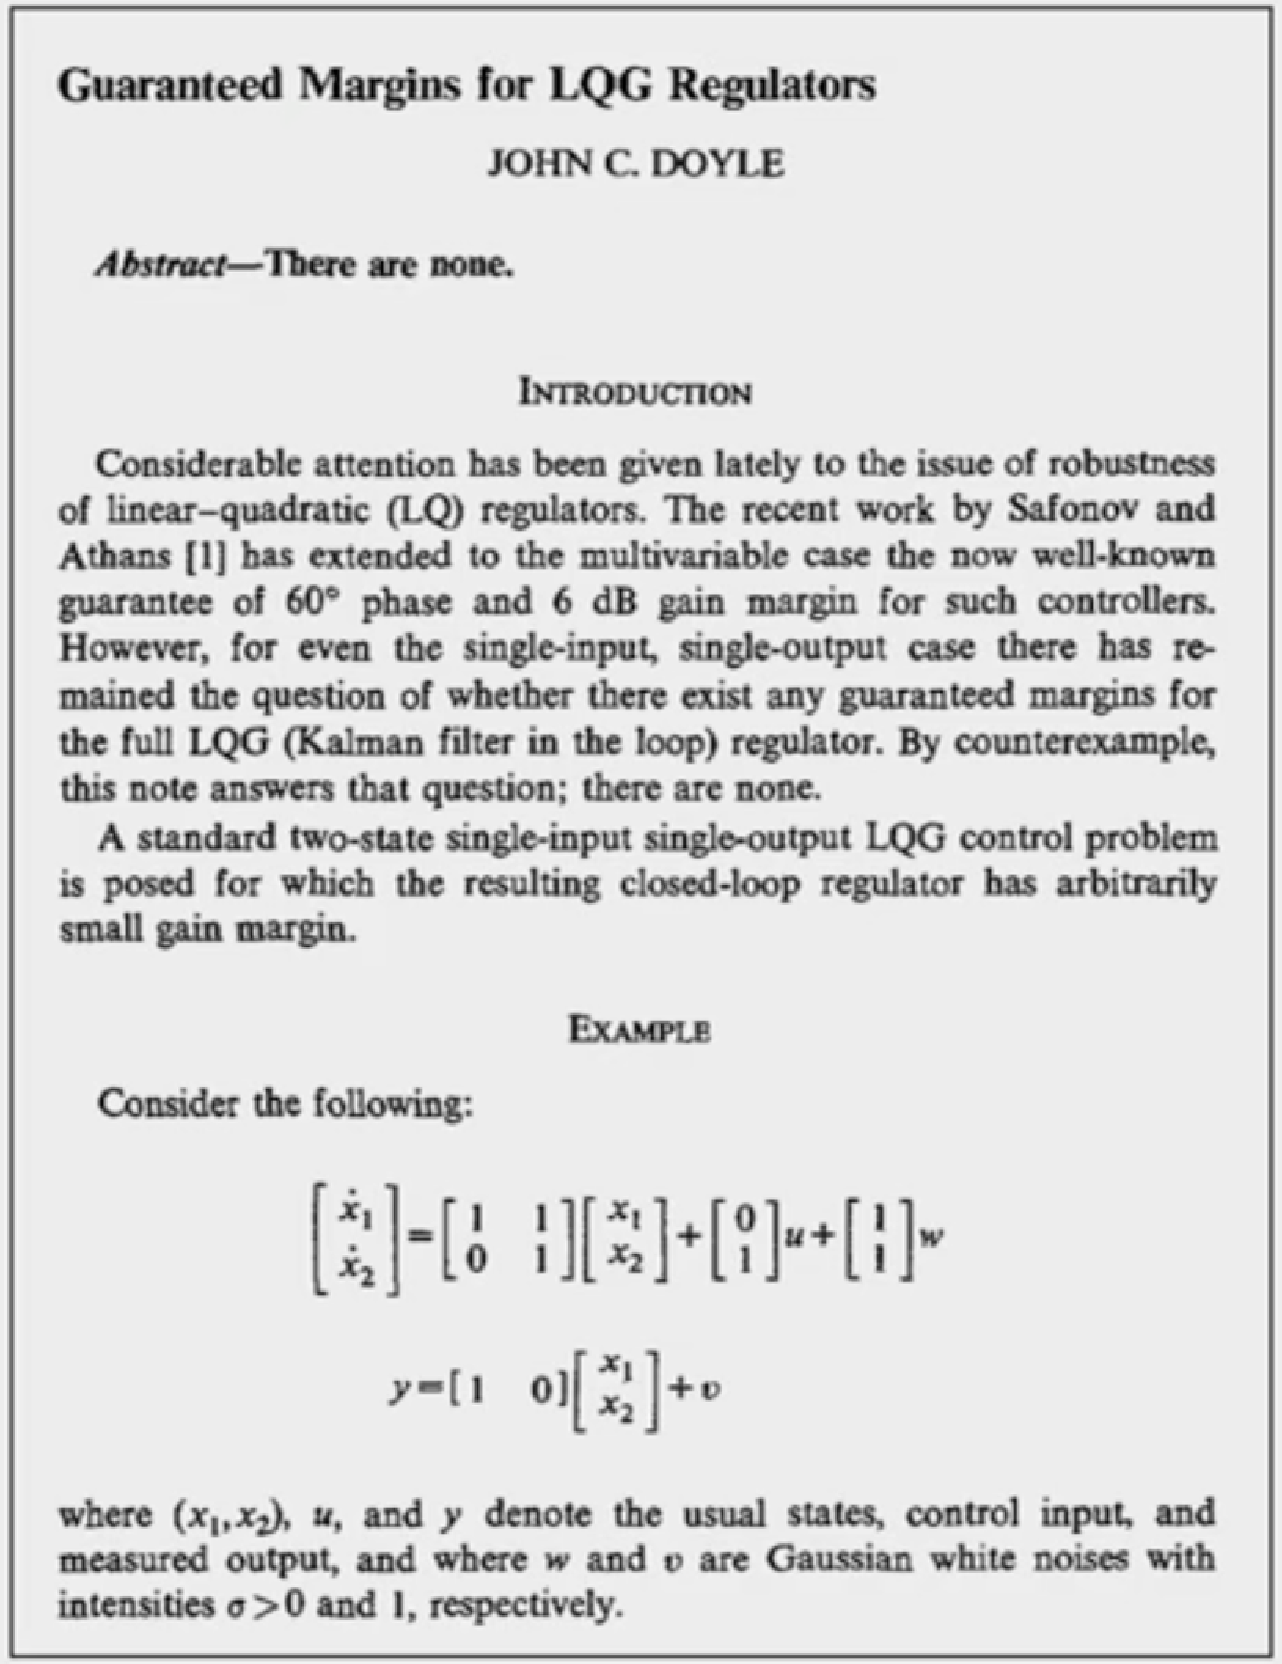

Guaranteed Margins for LQG Regulators

JOHN C. DOYLE

Abstract—There are none.

INTRODUCTION

Considerable attention has been given lately to the issue of robustness of linear–quadratic (LQ) regulators. The recent work by Safonov and Athans [1] has extended to the multivariable case the now well-known guarantee of 60° phase and 6 dB gain margin for such controllers. However, for even the single-input, single-output case there has remained the question of whether there exist any guaranteed margins for the full LQG (Kalman filter in the loop) regulator. By counterexample, this note answers that question; there are none.

A standard two-state single-input single-output LQG control problem is posed for which the resulting closed-loop regulator has arbitrarily small gain margin.

EXAMPLE

Consider the following:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u + \begin{bmatrix} 1 \\ 1 \end{bmatrix} w$$

$$y = [1 \quad 0] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + v$$

where (x_1, x_2), u, and y denote the usual states, control input, and measured output, and where w and v are Gaussian white noises with intensities $\sigma > 0$ and 1, respectively.

and that’s why we need robust control...

We need volunteers for a robotics experiment!

or how to get a little bonus to your final note...