EE363 Final Exam Solutions, June 2026

- This is a 3-hour in-class exam.
- Write your answers in the provided blue book. Make sure to clearly label each page with the problem number you are working on.
- If we ask you to "show" something is true, we are expecting a concise, intuitive derivation *not a formal proof.*
- You may use any books or paper notes.
- You may not use any electronic resources. For example, you cannot use a laptop, tablet or phone, run Julia code, search online, or consult a large language model.
- You may not discuss the exam or course material with others until everyone has taken the exam.
- If you have a question, please ask the staff. We have done our best to make the exam unambiguous. So unless there is a mistake, we are unlikely to say much.
- Please box or otherwise highlight your solution on the page.
- The problems on this exam are split into parts. Each part has a solution under 1/2 page. Thus, if you find yourself producing a complicated and time-consuming solution, stop and reconsider your assumptions.
- Good luck!

1. **Multiple choice questions (2 points each).** For each of the questions below, select exactly one correct answer. All matrices have compatible dimensions.

   a) Gauss–Newton LQR. Consider the nonlinear discrete-time system

   $$x_{t+1} = f(x_t, u_t)$$

   and the finite-horizon quadratic cost

   $$J = \sum_{\tau=0}^{N-1} \left(x_\tau^T Q x_\tau + u_\tau^T R u_\tau\right) + x_N^T Q_f x_N.$$

   Suppose Gauss–Newton LQR is applied by repeatedly linearizing the dynamics around the current nominal trajectory and solving the resulting finite-horizon LQR or tracking subproblem. Which of the following statements is **false**?

      i. A single Gauss–Newton LQR step solves a linear-quadratic approximation of the nonlinear optimal control problem, not necessarily the original nonlinear problem exactly.
      ii. The linearized subproblem can contain affine terms, and dynamic programming still applies to solve the resulting finite-horizon LQR or tracking problem.
      iii. For a fixed horizon, including reference-tracking terms changes the coefficients in the backward recursion, but it does not change the big-O computational complexity of the LQR backward pass.
      iv. If the goal is regulation rather than tracking a changing reference, Gauss–Newton LQR is guaranteed to converge to a time-invariant steady-state feedback law for every nonlinear system and every initial nominal trajectory.

   b) Observability with known inputs. Consider the discrete-time system

   $$x(t+1) = Ax(t) + Bu(t), \qquad y(t) = Cx(t) + Du(t),$$

   where the matrices are known and the input sequence $u(0), \ldots, u(T-1)$ is known. Define

   $$O_T = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{T-1} \end{bmatrix}.$$

   Select one statement that is always true.

      i. If O_T has nontrivial nullspace, choosing a sufficiently rich known input sequence can make $x(0)$ uniquely identifiable from $y(0), \ldots, y(T-1)$.
      ii. If O_T has full row rank, then $x(0)$ is uniquely identifiable from $y(0), \ldots, y(T-1)$.
      iii. If $z \in \text{null}(O_T)$, then the two initial states $x(0)$ and $x(0) + z$ produce the same output history $y(0), \ldots, y(T-1)$ for every choice of the known input sequence.
      iv. If the sensor noise is zero, the least-squares estimate using $O_T^\dagger$ always recovers the true $x(0)$, even when O_T has nontrivial nullspace.

   c) Internal stability and transfer-function stability. Consider a continuous-time realization

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

with transfer matrix

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

Call the realization internally stable if every eigenvalue of A has negative real part. Call G stable if every pole of G has negative real part. Select one statement that is always true.

i. If G is stable, then the realization is internally stable.
ii. If the realization is internally stable, then G is stable.
iii. Every eigenvalue of A must appear as a pole of G.
iv. If an eigenvalue of A does not appear as a pole of G, then the corresponding mode must be both uncontrollable and unobservable.

d) State feedback and uncontrollable modes. Consider the continuous-time system

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

with state feedback $u = Kx$. Suppose λ_u is an uncontrollable eigenvalue of (A, B), meaning that the PBH controllability test fails at $s = \lambda_u$. Select one statement that is always true.

i. No eigenvalues of A can be moved by state feedback.
ii. If $\operatorname{Re}\lambda_u < 0$, then arbitrary placement of all closed-loop eigenvalues is still possible.
iii. If $\operatorname{Re}\lambda_u > 0$, then sufficiently large feedback gains can always make the closed-loop system stable.
iv. The number λ_u is an eigenvalue of $A + BK$ for every feedback gain K.

e) Finite-horizon LQR tracking. Consider the finite-horizon tracking problem

$$x_{t+1} = Ax_t + Bu_t$$

with cost

$$J = \sum_{t=0}^{N-1} \big((x_t - r_t)^T Q(x_t - r_t) + u_t^T R u_t\big) + (x_N - r_N)^T Q_f (x_N - r_N),$$

where $Q, Q_f \succeq 0$, $R \succ 0$, and $r_0, \dots, r_N$ are known. Select one statement that is always true.

i. The optimal input can be written in the affine form

$$u_t = K_t x_t + g_t,$$

where the offset term g_t can be nonzero.
ii. Dynamic programming no longer applies because the cost is not centered at zero.
iii. If $x_t = r_t$, then the optimal input at time t must be $u_t = 0$.
iv. The problem is always equivalent to ordinary regulation of the error $e_t = x_t - r_t$ with no affine term in the feedback law.

f) A matrix and its elementwise absolute value. Consider the autonomous discrete-time systems

$$x(t+1) = Ax(t)$$

and

$$z(t+1) = |A|z(t),$$

where $A \in \mathbf{R}^{n\times n}$ and $|A|$ denotes the elementwise absolute value of A. A discrete-time autonomous system $x(t+1) = Fx(t)$ is called stable if $F^t \to 0$ as $t \to \infty$. Inequalities between matrices below are meant elementwise. Select one statement that is always true.

i. If $x(t+1) = Ax(t)$ is stable, then $z(t+1) = |A|z(t)$ is stable. The inequality $|A^t| \le |A|^t$ rules out any destabilizing effect of taking absolute values.

ii. If $z(t+1) = |A|z(t)$ is stable, then $x(t+1) = Ax(t)$ is stable. In fact, every matrix F satisfying $|F| \le |A|$ is stable whenever $|A|$ is stable.

iii. If A is diagonalizable over $\mathbf{C}$, then A and $|A|$ have the same spectral radius, so the two systems are stable or unstable together.

iv. Neither system's stability implies the other's, since changing signs of entries can arbitrarily change the eigenvalues.

g) The autonomous linear dynamical system $\dot{x} = Ax$ is a *constant norm* system if for every trajectory x, $\|x(t)\|$ is constant (as opposed to being a function of t). Which of the following statements is correct?

Hint: $\frac{\partial}{\partial t}\|x(t)\|^2 = 2x(t)^\top \dot{x}(t)$.

i. e^{tA} is an orthogonal matrix for all $t \ge 0$.

ii. $A + A^\top = 0$.

iii. $A = 0$.

iv. All of the above are valid sufficient conditions.

Solution.

a) Gauss–Newton LQR: the false statement is **iv**.

i. True. A Gauss–Newton LQR step solves a linear-quadratic approximation of the nonlinear problem. It need not solve the original nonlinear optimal control problem exactly.

ii. True. Linearizing around a nonzero nominal trajectory generally produces affine dynamics or affine tracking terms. Dynamic programming still applies.

iii. True. Tracking introduces additional linear and constant terms in the value function, but the backward pass has the same asymptotic order of computation.

iv. False. For a general nonlinear system, Gauss–Newton LQR is not guaranteed to converge, and a time-invariant steady-state feedback law need not exist.

b) Observability with known inputs: the correct answer is **iii**.

i. False. The known input contribution to the output history can be subtracted out. The ambiguity in $x(0)$ is determined by $\text{null}(O_T)$, not by how rich the known input sequence is.

ii. False. To uniquely identify $x(0)$, the map from $x(0)$ to the output history must have zero nullspace, i.e., O_T must have full column rank. Full row rank is not the right condition.

iii. True. If the same input sequence is applied from initial states $x(0)$ and $x(0) + z$, the difference between the two output histories is $O_T z$. If $z \in \text{null}(O_T)$, this difference is zero.

iv. False. If O_T has nontrivial nullspace, the output history does not determine the component of $x(0)$ in $\text{null}(O_T)$. The pseudoinverse returns one consistent initial state, not necessarily the true one.

c) Internal stability and transfer-function stability: the correct answer is **ii**.

i. False. A realization can have an unstable hidden mode that does not appear as a pole of G.

ii. True. Every pole of G is an eigenvalue of A, after any cancellations. If every eigenvalue of A has negative real part, then every pole of G has negative real part.

iii. False. Some eigenvalues of A can be hidden by uncontrollability or unobservability, and therefore need not appear as poles of G.

iv. False. A mode can fail to appear in G because it is uncontrollable or because it is unobservable. It need not be both.

d) State feedback and uncontrollable modes: the correct answer is **iv**.

i. False. Controllable modes can be moved by state feedback.

ii. False. Even if an uncontrollable eigenvalue is stable, it is still fixed. Therefore arbitrary placement of all closed-loop eigenvalues is not possible.

iii. False. An uncontrollable unstable eigenvalue cannot be moved by state feedback, no matter how large the gain is.

iv. True. Since the PBH test fails at $s = \lambda_u$, there is a nonzero left eigenvector w such that

$$w^T A = \lambda_u w^T, \qquad w^T B = 0.$$

Therefore

$$w^T(A + BK) = \lambda_u w^T + w^T BK = \lambda_u w^T$$

for every K, so λ_u remains a closed-loop eigenvalue.

e) Finite-horizon LQR tracking: the correct answer is **i**.

i. True. In a tracking problem, the value function generally has the form

$$V_t(z) = z^T P_t z + 2q_t^T z + s_t,$$

rather than just $z^T P_t z$. Minimizing the dynamic-programming expression then gives an affine feedback law $u_t = K_t x_t + g_t$.

ii. False. Dynamic programming still applies. The nonzero reference trajectory only changes the form of the value function by adding linear and constant terms.

iii. False. Even when $x_t = r_t$, a nonzero input may be needed to track future reference values.

iv. False. If the reference trajectory is not dynamically feasible, the error dynamics include an affine forcing term. In that case the optimal feedback generally needs a feedforward offset.

f) A matrix and its elementwise absolute value: the correct answer is **ii**.

Suppose $|F| \le |A|$ and $|A|$ is stable. Using $|FG| \le |F||G|$ and induction,

$$|F^t| \le |F|^t \le |A|^t, \qquad t = 0, 1, 2, \ldots.$$

Since $|A|$ is stable, $|A|^t \to 0$. Hence $|F^t| \to 0$, so $F^t \to 0$. Thus every F satisfying $|F| \le |A|$ is stable, and taking $F = A$ proves that stability of $|A|$ implies stability of A.

The converse is false. For example, let

$$A = \frac{2}{3}\begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}.$$

The eigenvalues of A are $\frac{2}{3}(1 \pm i)$, which have magnitude $2\sqrt{2}/3 < 1$, so $x(t+1) = Ax(t)$ is stable. However,

$$|A| = \frac{2}{3}\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix},$$

which has eigenvalues $4/3$ and 0. Hence $z(t+1) = |A|z(t)$ is not stable.

g) Constant norm systems. First we should figure out the conditions under which $\dot{x} = Ax$ is a constant norm system. This implies

$$\begin{aligned} 0 &= \frac{\partial}{\partial t}\|x(t)\|^2 \\ &= 2x(t)^\top \dot{x}(t) \\ &= 2x(t)^\top Ax(t) \end{aligned}$$

You might think this means $A = 0$, however recall that $2x(t)^\top Ax(t) = x(t)^\top (A + A^\top)x(t)$.

i. e^{tA} is an orthogonal matrix for all $t \geq 0$: This is a valid condition. It is equivalent to saying the velocity vector is always orthogonal to the position vector.

ii. $A + A^\top = 0$. This is a valid condition, derived above.

iii. $A = 0$. This is valid, but too restrictive.

iv. All of the above are are valid conditions, however some may be overly restrictive.

2. Lane keeping (10 points). A small robot moves forward at a constant speed of 1. Let $e(t)$ denote its lateral displacement from the center of a straight lane, and let $\theta(t)$ denote its heading angle relative to the lane direction. The yaw-rate command is $u(t)$. A simple nonlinear kinematic model is

$$\dot{e} = \sin\theta, \qquad \dot{\theta} = u, \qquad y = e.$$

We are interested in achieving centered, straight driving.

a) Find the equilibrium with $\theta \in [-\pi/2, \pi/2]$ corresponding to centered, straight driving.

b) Linearize the system around this equilibrium. Write the linearized system as

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

where $x = (e, \theta)^T$ denotes the small deviation from the equilibrium.

c) Suppose $u(t)$ is held constant on sampling intervals of length $\Delta > 0$. Let $x_d(k) = x(k\Delta)$, $u_d(k) = u(k\Delta)$, and $y_d(k) = y(k\Delta)$. Find the matrices A_d, B_d, C_d, and D_d, in terms of Δ, for the exact zero-order-hold discretization

$$x_d(k+1) = A_d x_d(k) + B_d u_d(k), \qquad y_d(k) = C_d x_d(k) + D_d u_d(k).$$

Use the exact zero-order-hold state-update formula

$$x_d(k+1) = e^{\Delta A} x_d(k) + \left(\int_0^{\Delta} e^{\tau A} B \, d\tau \right) u_d(k).$$

Hint: The correct matrix A from part (b) is nilpotent and

$$e^{\Delta A} = I + \Delta A + \frac{\Delta^2}{2!} A^2 + \frac{\Delta^3}{3!} A^3 + \cdots .$$

d) Is the sampled state-update system controllable for every $\Delta > 0$? Briefly justify your answer.

e) Now take $\Delta = 1$. For the sampled linearized system from part (c), find a discrete-time state-feedback gain

$$u_d(k) = K x_d(k), \qquad K = \begin{bmatrix} k_1 & k_2 \end{bmatrix},$$

that places both closed-loop eigenvalues at zero. For reference, the characteristic polynomial of a 2×2 matrix

$$M = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

is

$$\det(\lambda I - M) = \lambda^2 - (a+d)\lambda + (ad - bc).$$

f) For the gain found in part (e), let $A_{\text{cl}} = A_d + B_d K$ be the resulting closed-loop state-update matrix. Compute A_{cl}^2. What does its value imply for the sampled linearized closed-loop trajectories? Does the same conclusion hold for the original nonlinear system?

Solution.

a) Centered, straight driving corresponds to

$$e_e = 0, \qquad \theta_e = 0, \qquad u_e = 0.$$

At this point, $\dot{e} = \sin 0 = 0$ and $\dot{\theta} = 0$.

b) Let

$$f(e,\theta,u) = \begin{bmatrix} \sin\theta \\ u \end{bmatrix}.$$

The Jacobians at $(e_e, \theta_e, u_e) = (0,0,0)$ are

$$A = \left.\frac{\partial f}{\partial(e,\theta)}\right|_{(0,0,0)} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad B = \left.\frac{\partial f}{\partial u}\right|_{(0,0,0)} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Since $y = e$, we have

$$C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \qquad D = 0.$$

Thus

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u, \qquad y = \begin{bmatrix} 1 & 0 \end{bmatrix} x.$$

c) Since $A^2 = 0$,

$$e^{\tau A} = I + \tau A = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix}.$$

Therefore

$$A_d = e^{\Delta A} = \begin{bmatrix} 1 & \Delta \\ 0 & 1 \end{bmatrix},$$

and

$$B_d = \int_0^\Delta e^{\tau A} B \, d\tau = \int_0^\Delta \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} d\tau = \int_0^\Delta \begin{bmatrix} \tau \\ 1 \end{bmatrix} d\tau = \begin{bmatrix} \Delta^2/2 \\ \Delta \end{bmatrix}.$$

The output matrices are

$$C_d = C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \qquad D_d = D = 0.$$

d) The discrete-time controllability matrix is

$$\mathcal{C}_{\text{ctrb}} = \begin{bmatrix} B_d & A_d B_d \end{bmatrix} = \begin{bmatrix} \Delta^2/2 & 3\Delta^2/2 \\ \Delta & \Delta \end{bmatrix}.$$

Its determinant is

$$\det \mathcal{C}_{\text{ctrb}} = \frac{\Delta^2}{2}\Delta - \frac{3\Delta^2}{2}\Delta = -\Delta^3.$$

This is nonzero for every $\Delta > 0$, so the sampled state-update system is controllable for every positive sampling period.

e) With $\Delta = 1$,

$$A_d = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad B_d = \begin{bmatrix} 1/2 \\ 1 \end{bmatrix}.$$

For $K = \begin{bmatrix} k_1 & k_2 \end{bmatrix}$, the closed-loop matrix is

$$A_{\text{cl}} = A_d + B_d K = \begin{bmatrix} 1 + k_1/2 & 1 + k_2/2 \\ k_1 & 1 + k_2 \end{bmatrix}.$$

To put both eigenvalues at zero, the characteristic polynomial must be λ^2, so the trace and determinant must both be zero:

$$2 + \frac{k_1}{2} + k_2 = 0, \qquad 1 + k_2 - \frac{k_1}{2} = 0.$$

Solving gives

$$k_1 = -1, \qquad k_2 = -\frac{3}{2},$$

so

$$K = \begin{bmatrix} -1 & -3/2 \end{bmatrix}.$$

f) With the gain from part (e),

$$A_{\mathrm{cl}} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} \begin{bmatrix} -1 & -3/2 \end{bmatrix} = \begin{bmatrix} 1/2 & 1/4 \\ -1 & -1/2 \end{bmatrix}.$$

Then

$$A_{\mathrm{cl}}^2 = \begin{bmatrix} 1/2 & 1/4 \\ -1 & -1/2 \end{bmatrix}^2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Thus the sampled linearized closed-loop state reaches zero in at most two steps, from any initial sampled state.

No. This deadbeat conclusion is an exact statement about the sampled linearized system, not a global statement about the original nonlinear system. With u held constant over one sample interval of length $\Delta = 1$,

$$\theta(t) = \theta_k + u_k(t-k),$$

so

$$e_{k+1} = e_k + \int_0^1 \sin(\theta_k + u_k\tau)\, d\tau,$$

which is not equal in general to the linearized update

$$e_k + \theta_k + \frac{1}{2}u_k.$$

The linearized model is only a local approximation near $\theta = 0$, so the controller does not guarantee exact two-step convergence of the nonlinear robot for arbitrary initial conditions.

3. **Continuous-time oscillator observed through sampled data (10 points).** Consider the autonomous continuous-time system

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

where $x(t) \in \mathbb{R}^3$, $y(t) \in \mathbb{R}$, and

$$A = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}.$$

a) Write $x_1(t), x_2(t), x_3(t)$ in terms of $x_1(0), x_2(0), x_3(0)$.

b) Write $y(t)$ in terms of $x_1(0), x_2(0), x_3(0)$. Is every output trajectory bounded? Does every output trajectory converge? Briefly justify your answers.

c) Form the continuous-time observability matrix. Is the continuous-time system observable?

For the remaining parts, consider the sampled system

$$x_d(k+1) = A_d x_d(k), \qquad y_d(k) = C_d x_d(k),$$

with sampling period $\Delta > 0$, where

$$A_d = \begin{bmatrix} e^{-\Delta} & 0 & 0 \\ 0 & \cos\Delta & \sin\Delta \\ 0 & -\sin\Delta & \cos\Delta \end{bmatrix}, \qquad C_d = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}.$$

d) Consider determinant of the observability matrix for the sampled system

$$\det O_d = \sin\Delta \left(e^{-2\Delta} - 2e^{-\Delta}\cos\Delta + 1\right).$$

Is there a value of $\Delta > 0$ such that the sampled system is unobservable? If so, give such Δ and a corresponding nonzero initial condition that produces $y_d(k) = 0$ for all $k \geq 0$.

e) Explain why the answer in the previous part does not contradict the answer in part (c). In particular, can two different initial states give the same sampled output sequence but different continuous-time output signals between samples?

f) Suppose Δ is such that the sampled system is observable and we measure

$$\tilde{y}_d(k) = y_d(k) + v(k), \qquad k = 0, \ldots, N,$$

where $N \geq 2$. Write the least-squares estimate of $x(0)$ in terms of the measured output vector

$$\tilde{Y} = \begin{bmatrix} \tilde{y}_d(0) & \tilde{y}_d(1) & \cdots & \tilde{y}_d(N) \end{bmatrix}^T.$$

Solution. The state equations are

$$\dot{x}_1 = -x_1, \qquad \dot{x}_2 = x_3, \qquad \dot{x}_3 = -x_2.$$

a) The first state is

$$x_1(t) = e^{-t} x_1(0).$$

For the oscillator block,

$$\begin{bmatrix} x_2(t) \\ x_3(t) \end{bmatrix} = \begin{bmatrix} \cos t & \sin t \\ -\sin t & \cos t \end{bmatrix} \begin{bmatrix} x_2(0) \\ x_3(0) \end{bmatrix}.$$

Therefore

$$x_2(t) = x_2(0)\cos t + x_3(0)\sin t, \qquad x_3(t) = -x_2(0)\sin t + x_3(0)\cos t.$$

b) Since $y = x_1 + x_2$,

$$y(t) = e^{-t}x_1(0) + x_2(0)\cos t + x_3(0)\sin t.$$

Every output trajectory is bounded, since the first term decays and the last two terms are sinusoidal. Not every output trajectory converges. For example, if $x_1(0) = 0$, $x_2(0) = 1$, and $x_3(0) = 0$, then $y(t) = \cos t$, which does not converge. In fact, $y(t)$ converges if and only if

$$x_2(0) = x_3(0) = 0.$$

c) We have

$$C = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}, \qquad CA = \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}, \qquad CA^2 = \begin{bmatrix} 1 & -1 & 0 \end{bmatrix}.$$

Thus

$$O = \begin{bmatrix} 1 & 1 & 0 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}.$$

Its determinant is

$$\det O = 2,$$

so O has full rank. Therefore the continuous-time system is observable.

d) The sampled observability matrix using the first three samples is

$$O_d = \begin{bmatrix} C_d \\ C_d A_d \\ C_d A_d^2 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ e^{-\Delta} & \cos\Delta & \sin\Delta \\ e^{-2\Delta} & \cos(2\Delta) & \sin(2\Delta) \end{bmatrix}.$$

Its determinant is

$$\det O_d = \sin\Delta \left(e^{-2\Delta} - 2e^{-\Delta}\cos\Delta + 1\right).$$

The second factor is

$$e^{-2\Delta} - 2e^{-\Delta}\cos\Delta + 1 = |e^{-\Delta} - e^{i\Delta}|^2,$$

which is positive because $0 < e^{-\Delta} < 1$ and $|e^{i\Delta}| = 1$. Hence the sampled system is unobservable exactly when

$$\sin\Delta = 0,$$

i.e., when

$$\Delta = m\pi, \qquad m = 1, 2, 3, \ldots.$$

For $\Delta = \pi$, the nonzero initial condition

$$x(0) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

gives

$$y(t) = \sin t,$$

and therefore

$$y_d(k) = y(k\pi) = 0$$

for every $k \geq 0$.

e) There is no contradiction. Continuous-time observability uses the entire signal $y(t)$ over a time interval, while sampled observability uses only the values $y(0), y(\Delta), y(2\Delta), \ldots$. When $\Delta = \pi$, the initial states

$$x(0) = 0 \qquad \text{and} \qquad x(0) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

produce the same sampled output sequence, namely $y_d(k) = 0$ for all k. But their continuous-time outputs are different: the first gives $y(t) = 0$, while the second gives $y(t) = \sin t$. Thus sampling can hide a continuous-time oscillatory mode by observing it only at its zero crossings.

f) Define the sampled observability matrix over the available measurements by

$$O_N = \begin{bmatrix} C_d \\ C_d A_d \\ \vdots \\ C_d A_d^N \end{bmatrix}.$$

Since Δ is not one of the bad sampling periods and $N \geq 2$, the first three rows already have full column rank, so O_N has full column rank. The measurement model is

$$\tilde{Y} = O_N x(0) + V,$$

where $V = \begin{bmatrix} v(0) & \cdots & v(N) \end{bmatrix}^T$. The least-squares estimate is

$$\hat{x}(0) = (O_N^T O_N)^{-1} O_N^T \tilde{Y}.$$

Equivalently, $\hat{x}(0) = O_N^\dagger \tilde{Y}$.

4. Checkpoint control for a double integrator (10 points). Consider the discrete-time system

$$x(k+1) = Ax(k) + Bu(k), \qquad x(0) = 0,$$

where

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} 1/2 \\ 1 \end{bmatrix}.$$

The state is $x(k) = (p(k), v(k))^T$, where $p(k)$ is position and $v(k)$ is velocity. We choose the four scalar inputs $u(0), u(1), u(2), u(3)$ and require the vehicle to pass through the checkpoint

$$p(2) = 1$$

and return to rest at the origin at time 4:

$$x(4) = 0.$$

The input energy is

$$E = \sum_{k=0}^{3} u(k)^2.$$

a) Let

$$U = \begin{bmatrix} u(0) & u(1) & u(2) & u(3) \end{bmatrix}^T.$$

Find a matrix G and a vector h such that the checkpoint and terminal constraints are equivalent to

$$GU = h.$$

b) Are the constraints feasible? Briefly justify your answer.

c) Give an expression for the minimum-energy input $U^\star$ in terms of G and h. Is minimum-energy input unique?

d) Give the minimum energy $E^\star$ as an expression in terms of G and h.

e) Suppose the terminal time were changed from 4 to 5, while the checkpoint constraint $p(2) = 1$ is unchanged and the terminal constraint becomes $x(5) = 0$. Let $E_5^\star$ denote the minimum energy for this longer-horizon problem. What relationship must $E_5^\star$ have to $E^\star$? Briefly justify your answer.

Solution.

a) We have

$$x(2) = ABu(0) + Bu(1),$$

so

$$p(2) = \frac{3}{2}u(0) + \frac{1}{2}u(1).$$

Also,

$$x(4) = A^3Bu(0) + A^2Bu(1) + ABu(2) + Bu(3).$$

Since

$$A^rB = \begin{bmatrix} r + 1/2 \\ 1 \end{bmatrix},$$

the terminal constraint is

$$x(4) = \begin{bmatrix} 7/2 & 5/2 & 3/2 & 1/2 \\ 1 & 1 & 1 & 1 \end{bmatrix} U = 0.$$

Thus we can take

$$G = \begin{bmatrix} 3/2 & 1/2 & 0 & 0 \\ 7/2 & 5/2 & 3/2 & 1/2 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \qquad h = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

b) The constraints are feasible. For example, the first three columns of G form the matrix

$$\begin{bmatrix} 3/2 & 1/2 & 0 \\ 7/2 & 5/2 & 3/2 \\ 1 & 1 & 1 \end{bmatrix},$$

whose determinant is $1/2$. Hence G has full row rank. Since $G \in \mathbb{R}^{3\times 4}$ has rank 3, its range is all of $\mathbb{R}^3$, so $GU = h$ is feasible.

c) Since G has full row rank, GG^T is invertible. The minimum-energy input is the minimum-norm solution of

$$GU = h,$$

and is

$$U^\star = G^T(GG^T)^{-1}h.$$

This is feasible because

$$GU^\star = GG^T(GG^T)^{-1}h = h.$$

The minimum-energy input is unique. One way to see this is as follows. Let U be any other feasible input. Then

$$G(U - U^\star) = 0,$$

so $U - U^\star \in \text{null}(G)$. Also, $U^\star \in \text{range}(G^T)$, and

$$\text{range}(G^T) \perp \text{null}(G).$$

Therefore

$$\|U\|_2^2 = \|U^\star + (U - U^\star)\|_2^2 = \|U^\star\|_2^2 + \|U - U^\star\|_2^2.$$

Thus every feasible input other than $U^\star$ has strictly larger energy.

d) The minimum energy is

$$E^\star = \|U^\star\|_2^2.$$

Substituting

$$U^\star = G^T(GG^T)^{-1}h$$

gives

$$E^\star = h^T(GG^T)^{-1}GG^T(GG^T)^{-1}h.$$

Since GG^T is symmetric and invertible, this simplifies to

$$E^\star = h^T(GG^T)^{-1}h.$$

e) We must have

$$E_5^\star \leq E^\star.$$

The minimum energy cannot increase when the horizon is extended. To see this, take a minimum-energy feasible input for the four-step problem,

$$u^\star(0), u^\star(1), u^\star(2), u^\star(3),$$

and append

$$u(4) = 0.$$

Since the state is already $x(4) = 0$, the extra zero input keeps the state at zero, so $x(5) = 0$. The checkpoint constraint $p(2) = 1$ is unaffected. Therefore this gives a feasible input for the five-step problem with the same energy $E^\star$. Hence the minimum energy for the five-step problem is at most $E^\star$.

5. Interconnection, hidden modes, and minimal realization (10 points). Consider two continuous-time subsystems connected in cascade. The first subsystem is

$$\dot{x} = -x + u, \qquad w = x,$$

and the second subsystem is

$$\dot{z}_1 = 2z_1 + w, \qquad \dot{z}_2 = -3z_2 + w, \qquad y = z_2.$$

The input to the overall interconnected system is u, the output is y, and the overall state is

$$\xi = \begin{bmatrix} x & z_1 & z_2 \end{bmatrix}^T.$$

For this problem, call a realization internally stable if its state matrix is stable. Call a transfer function or transfer matrix stable if all of its poles are stable. We say that a mode is hidden if it is an eigenvalue of the state matrix of a realization, but it is not a pole of the transfer function from the specified input to the specified output.

a) Write the overall interconnected system in the form

$$\dot{\xi} = A_{\text{int}}\xi + B_{\text{int}}u, \qquad y = C_{\text{int}}\xi.$$

b) Compute the transfer function

$$G(s) = C_{\text{int}}(sI - A_{\text{int}})^{-1}B_{\text{int}}$$

from u to y.

c) For each eigenvalue of A_{int}, determine whether the associated mode is controllable and whether it is observable. Briefly justify your answer.

d) Which mode is hidden? Is the hidden mode stable or unstable?

e) Is the interconnected realization internally stable? Is $G(s)$ stable? Briefly explain.

f) Give a minimal realization of $G(s)$, and briefly justify why it is minimal.

g) Now suppose the output of the second subsystem is changed to

$$\tilde{y} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix},$$

while the input remains u. Would the mode at 2 still be hidden? Would the transfer matrix from u to $\tilde{y}$ be stable? Briefly explain.

Solution.

a) Since $w = x$, the interconnected equations are

$$\dot{x} = -x + u, \qquad \dot{z}_1 = x + 2z_1, \qquad \dot{z}_2 = x - 3z_2, \qquad y = z_2.$$

Therefore

$$A_{\text{int}} = \begin{bmatrix} -1 & 0 & 0 \\ 1 & 2 & 0 \\ 1 & 0 & -3 \end{bmatrix}, \qquad B_{\text{int}} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \qquad C_{\text{int}} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}.$$

b) One way to compute the transfer function is to use the cascade structure. From $\dot{x} = -x + u$, the transfer function from u to $w = x$ is

$$\frac{1}{s+1}.$$

From $\dot{z}_2 = -3z_2 + w$, the transfer function from w to $y = z_2$ is

$$\frac{1}{s+3}.$$

Thus

$$G(s) = \frac{1}{(s+1)(s+3)}.$$

Equivalently, $G(s) = C_{\text{int}}(sI - A_{\text{int}})^{-1}B_{\text{int}}$.

c) The matrix A_{int} is lower triangular, so its eigenvalues are

$$-1, \qquad 2, \qquad -3.$$

First, the controllability matrix is

$$\mathcal{C} = \begin{bmatrix} B_{\text{int}} & A_{\text{int}}B_{\text{int}} & A_{\text{int}}^2 B_{\text{int}} \end{bmatrix} = \begin{bmatrix} 1 & -1 & 1 \\ 0 & 1 & 1 \\ 0 & 1 & -4 \end{bmatrix}.$$

Its determinant is -5, so the realization is controllable. In particular, all three modes are controllable.

To analyze observability, note that

$$\mathcal{O} = \begin{bmatrix} C_{\text{int}} \\ C_{\text{int}}A_{\text{int}} \\ C_{\text{int}}A_{\text{int}}^2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & -3 \\ -4 & 0 & 9 \end{bmatrix}.$$

The nullspace of $\mathcal{O}$ is $\operatorname{span}\{e_2\}$. Since e_2 is the right eigenvector associated with the eigenvalue 2, the mode at 2 is unobservable. The modes at -1 and -3 are observable. Thus

eigenvalue	controllable?	observable?
-1	yes	yes
2	yes	no
-3	yes	yes

d) The hidden mode is the mode at eigenvalue 2. It is an eigenvalue of A_{int}, but it is not a pole of

$$G(s) = \frac{1}{(s+1)(s+3)}.$$

The hidden mode is unstable, since its real part is positive.

e) The interconnected realization is not internally stable, because A_{int} has the eigenvalue 2. The transfer function $G(s)$ is stable, because its only poles are -1 and -3, both of which have negative real part. This is possible because the unstable mode at 2 is hidden from the input-output map from u to y.

f) A two-state realization of $G(s)$ is obtained by keeping only the states x and z_2:

$$A_m = \begin{bmatrix} -1 & 0 \\ 1 & -3 \end{bmatrix}, \qquad B_m = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad C_m = \begin{bmatrix} 0 & 1 \end{bmatrix}, \qquad D_m = 0.$$

Indeed,

$$C_m(sI - A_m)^{-1}B_m = \frac{1}{(s+1)(s+3)}.$$

This realization is controllable and observable, since

$$\begin{bmatrix} B_m & A_m B_m \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}, \qquad \begin{bmatrix} C_m \\ C_m A_m \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & -3 \end{bmatrix}$$

are both nonsingular. Hence it is minimal.

g) With the changed output,

$$\tilde{C} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The transfer matrix from u to $\tilde{y}$ is

$$\tilde{G}(s) = \tilde{C}(sI - A_{\text{int}})^{-1}B_{\text{int}} = \begin{bmatrix} \dfrac{1}{(s+1)(s-2)} \\ \dfrac{1}{(s+1)(s+3)} \end{bmatrix}.$$

The mode at 2 is no longer hidden, since it appears as a pole in the first entry of $\tilde{G}(s)$. The transfer matrix from u to $\tilde{y}$ is not stable, because it has the pole $s = 2$.

6. Greedy one-step control versus pole placement (10 points). Consider the discrete-time system

$$x(k+1) = Ax(k) + Bu(k), \qquad A = \begin{bmatrix} 0 & 4 \\ 0 & 1/2 \end{bmatrix}, \qquad B = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

The input $u(k)$ is scalar. For this problem, call a feedback law *greedy one-step* if, at each time k, it chooses $u(k)$ to minimize the next-state norm

$$\|x(k+1)\|_2 = \|Ax(k) + Bu(k)\|_2.$$

a) Derive the greedy one-step feedback law in the form

$$u(k) = K_{\mathrm{g}} x(k).$$

b) Compute the greedy closed-loop matrix

$$F_{\mathrm{g}} = A + BK_{\mathrm{g}}.$$

c) Is the greedy closed-loop system $x(k+1) = F_{\mathrm{g}} x(k)$ stable? Briefly justify your answer.

d) Now ignore the greedy law. Is the pair (A, B) controllable? Briefly justify your answer.

e) Call a discrete-time feedback law *deadbeat* if all eigenvalues of the closed-loop matrix are zero. Find a state-feedback gain

$$u(k) = Kx(k), \qquad K = \begin{bmatrix} k_1 & k_2 \end{bmatrix},$$

that makes the closed-loop system deadbeat. For reference, the characteristic polynomial of a 2×2 matrix

$$M = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

is

$$\det(\lambda I - M) = \lambda^2 - (a + d)\lambda + (ad - bc).$$

f) Compare the closed-loop dynamics under the greedy gain and under the deadbeat gain. Which feedback law is preferable for long-term operation of the system?

Solution.

a) For a fixed state $x = x(k)$, the greedy law solves the scalar least-squares problem

$$\min_u \|Ax + Bu\|_2^2.$$

The optimality condition is

$$B^T(Ax + Bu) = 0.$$

Since $B^T B = 2$ and

$$B^T A = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 4 \\ 0 & 1/2 \end{bmatrix} = \begin{bmatrix} 0 & 9/2 \end{bmatrix},$$

we get

$$u = -\frac{1}{B^T B} B^T A x = -\frac{1}{2} \begin{bmatrix} 0 & 9/2 \end{bmatrix} x = \begin{bmatrix} 0 & -9/4 \end{bmatrix} x.$$

Thus

$$K_{\mathrm{g}} = \begin{bmatrix} 0 & -9/4 \end{bmatrix}.$$

b) The greedy closed-loop matrix is

$$F_{\mathrm{g}} = A + BK_{\mathrm{g}} = \begin{bmatrix} 0 & 4 \\ 0 & 1/2 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \begin{bmatrix} 0 & -9/4 \end{bmatrix} = \begin{bmatrix} 0 & 7/4 \\ 0 & -7/4 \end{bmatrix}.$$

c) The matrix F_{g} is upper triangular, so its eigenvalues are

$$0, \qquad -\frac{7}{4}.$$

Since one eigenvalue has magnitude $7/4 > 1$, the greedy closed-loop system is not stable.

d) The controllability matrix is

$$\mathcal{C} = \begin{bmatrix} B & AB \end{bmatrix} = \begin{bmatrix} 1 & 4 \\ 1 & 1/2 \end{bmatrix}.$$

Its determinant is

$$\det \mathcal{C} = \frac{1}{2} - 4 = -\frac{7}{2} \neq 0.$$

Hence (A, B) is controllable.

e) For $K = \begin{bmatrix} k_1 & k_2 \end{bmatrix}$,

$$A + BK = \begin{bmatrix} 0 & 4 \\ 0 & 1/2 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \begin{bmatrix} k_1 & k_2 \end{bmatrix} = \begin{bmatrix} k_1 & 4 + k_2 \\ k_1 & 1/2 + k_2 \end{bmatrix}.$$

To make both eigenvalues zero, the trace and determinant must both be zero. The trace is

$$k_1 + k_2 + \frac{1}{2},$$

and the determinant is

$$k_1 \left(\frac{1}{2} + k_2 \right) - k_1 (4 + k_2) = -\frac{7}{2} k_1.$$

Thus $k_1 = 0$, and then the zero-trace condition gives $k_2 = -1/2$. Therefore

$$K = \begin{bmatrix} 0 & -1/2 \end{bmatrix}.$$

With this gain,

$$F_{\mathrm{d}} = A + BK = \begin{bmatrix} 0 & 7/2 \\ 0 & 0 \end{bmatrix},$$

whose two eigenvalues are both zero.

f) The deadbeat feedback law is preferable for long-term operation. Under the deadbeat gain,

$$F_{\mathrm{d}} = \begin{bmatrix} 0 & 7/2 \\ 0 & 0 \end{bmatrix}, \qquad F_{\mathrm{d}}^2 = 0.$$

Therefore, for any initial condition,

$$x(k) = F_{\mathrm{d}}^k x(0) = 0, \qquad k \geq 2.$$

In contrast, the greedy gain gives

$$F_{\mathrm{g}} = \begin{bmatrix} 0 & 7/4 \\ 0 & -7/4 \end{bmatrix},$$

with eigenvalues 0 and $-7/4$. Since $|-7/4| > 1$, the greedy closed-loop dynamics are unstable. The reason is that the greedy law minimizes only the immediate quantity $\|x(k+1)\|_2$ at the current step; it does not directly choose the closed-loop eigenvalues or account for the behavior of the trajectory over many time steps.

7. Output tracking as a least-squares problem (10 points). Consider the discrete-time double-integrator system

$$x_{t+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x_t + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} u_t, \qquad y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} x_t, \qquad x_0 = 0.$$

We choose u_0, u_1, u_2 to make the position output track the desired values

$$y_1^{\text{des}} = 1, \qquad y_2^{\text{des}} = 2, \qquad y_3^{\text{des}} = 2.$$

For a fixed parameter $\rho > 0$, define

$$J = \sum_{t=1}^{3} (y_t - y_t^{\text{des}})^2 + \rho \sum_{t=0}^{2} u_t^2.$$

a) Let

$$U = \begin{bmatrix} u_0 \\ u_1 \\ u_2 \end{bmatrix}, \qquad Y = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}.$$

Find the matrix F such that

$$Y = FU.$$

b) Rewrite the cost as a regularized least squares problem and give the optimal input sequence $U^\star$ in terms of F, ρ, and $r = (1, 2, 2)^T$.

c) What happens as $\rho \to 0$? Is exact tracking possible in this example?

d) What happens as $\rho \to \infty$?

e) How is the batch least-squares solution related to the one that would be obtained by applying dynamic programming and the finite-horizon Riccati recursion? What is the main difference between the two approaches?

Solution.

a) Since $x_0 = 0$, we can propagate the dynamics directly. First,

$$x_1 = Bu_0 = \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} u_0, \qquad y_1 = \frac{1}{2} u_0.$$

Next,

$$x_2 = Ax_1 + Bu_1 = \begin{bmatrix} 3/2 \\ 1 \end{bmatrix} u_0 + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} u_1, \qquad y_2 = \frac{3}{2} u_0 + \frac{1}{2} u_1.$$

Finally,

$$x_3 = A^2 B u_0 + ABu_1 + Bu_2,$$

so

$$y_3 = \frac{5}{2} u_0 + \frac{3}{2} u_1 + \frac{1}{2} u_2.$$

Therefore

$$Y = FU, \qquad F = \begin{bmatrix} 1/2 & 0 & 0 \\ 3/2 & 1/2 & 0 \\ 5/2 & 3/2 & 1/2 \end{bmatrix}.$$

b) Let

$$r = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}.$$

Since $Y = FU$, the tracking-error vector is

$$Y - r = FU - r.$$

Also,

$$\sum_{t=0}^{2} u_t^2 = \|U\|_2^2.$$

Hence the cost can be written as the regularized least-squares objective

$$J = \|FU - r\|_2^2 + \rho\|U\|_2^2.$$

Expanding,

$$J = (FU - r)^T(FU - r) + \rho U^T U.$$

Setting the gradient with respect to U equal to zero gives

$$2F^T(FU - r) + 2\rho U = 0.$$

Thus

$$(F^T F + \rho I)U = F^T r.$$

Since $\rho > 0$, the matrix $F^T F + \rho I$ is positive definite. Therefore the optimal input sequence is unique and is

$$U^\star = (F^T F + \rho I)^{-1} F^T r.$$

c) As $\rho \to 0$, the input penalty disappears, and the problem approaches

$$\min_U \|FU - r\|_2^2.$$

In this example, F is lower triangular with nonzero diagonal entries, so F is invertible. Therefore exact tracking is possible, and

$$U^\star \to F^{-1} r.$$

Solving $FU = r$ gives

$$U = \begin{bmatrix} 2 \\ -2 \\ 0 \end{bmatrix}.$$

This input gives $Y = r$, so the tracking error is zero.

d) As $\rho \to \infty$, input effort dominates the cost. From

$$U^\star = (F^T F + \rho I)^{-1} F^T r,$$

we get

$$U^\star \to 0.$$

Consequently,

$$Y = FU^\star \to 0.$$

Thus, when input is extremely expensive, the optimal controller uses essentially no input and accepts the tracking error.

e) The solutions are the same. Both methods solve the same finite-horizon quadratic optimal-control problem, so they give the same optimal input sequence for the same initial condition and desired trajectory. The batch least-squares approach stacks all inputs into one vector and solves one linear algebra problem, while the dynamic-programming/Riccati approach works backward in time and gives a recursive feedback, or affine feedback for tracking, implementation.

8. Scalar deterministic and stochastic LQR (10 points). Consider the scalar discrete-time system

$$x_{t+1} = 2x_t + u_t.$$

We first consider the deterministic infinite-horizon LQR problem

$$J = \sum_{t=0}^{\infty} \left(x_t^2 + u_t^2\right),$$

where the initial state x_0 is given.

a) Assume the value function has the form $V(z) = Pz^2$. Write the Hamilton-Jacobi equation for P.

b) Derive the scalar algebraic Riccati equation

$$P = 1 + 4P - \frac{4P^2}{1+P},$$

and find its positive semidefinite solution.

c) Find the optimal feedback law

$$u_t = Kx_t.$$

Is the deterministic closed-loop system stable?

d) Under the optimal feedback law, find the deterministic stage cost

$$\ell_t^\star = x_t^2 + u_t^2$$

as a function of t and x_0. What are the optimal total cost $J^\star(x_0)$ and the optimal long-run average stage cost

$$J^\star_{\text{avg,det}} = \lim_{N\to\infty} \frac{1}{N} \sum_{t=0}^{N-1} \ell_t^\star?$$

e) Now suppose the system is instead

$$x_{t+1} = 2x_t + u_t + w_t,$$

where w_t are independent, identically distributed random variables with $\mathbf{E}w_t = 0$ and $\mathbf{E}w_t^2 = \sigma^2$. For the infinite-horizon stochastic average-cost problem

$$J_{\text{avg}} = \lim_{N\to\infty} \frac{1}{N} \mathbf{E} \sum_{t=0}^{N-1} \left(x_t^2 + u_t^2\right),$$

what is the optimal feedback gain? What is the optimal average stage cost? Briefly compare the deterministic and stochastic average stage costs.

Solution.

a) The Hamilton-Jacobi equation is

$$Pz^2 = \min_v \left\{z^2 + v^2 + P(2z+v)^2\right\}.$$

Equivalently,

$$Pz^2 = z^2 + \min_v \left\{v^2 + P(2z+v)^2\right\}$$

for every scalar z.

b) For a fixed z, the quantity to minimize over v is

$$v^2 + P(2z+v)^2.$$

Differentiating with respect to v gives

$$2v + 2P(2z+v) = 0,$$

so

$$v^\star = -\frac{2P}{1+P}z.$$

Substituting this expression into the Hamilton-Jacobi equation gives

$$P = 1 + 4P - \frac{4P^2}{1+P}.$$

Multiplying by $1+P$ gives

$$P + P^2 = 1 + 5P,$$

or

$$P^2 - 4P - 1 = 0.$$

Thus

$$P = 2 \pm \sqrt{5}.$$

The positive semidefinite solution is

$$P = 2 + \sqrt{5}.$$

c) From the minimizing input in part (b),

$$K = -\frac{2P}{1+P}.$$

With $P = 2+\sqrt{5}$, this simplifies to

$$K = -\frac{1+\sqrt{5}}{2}.$$

The closed-loop scalar is

$$2 + K = 2 - \frac{1+\sqrt{5}}{2} = \frac{3-\sqrt{5}}{2}.$$

Since

$$\left|\frac{3-\sqrt{5}}{2}\right| < 1,$$

the deterministic discrete-time closed-loop system is stable.

d) Let

$$a_{\text{cl}} = 2 + K = \frac{3-\sqrt{5}}{2}.$$

Under the optimal feedback law,

$$x_t = a_{\text{cl}}^t x_0, \qquad u_t = K a_{\text{cl}}^t x_0.$$

Therefore the optimal deterministic stage cost is

$$\begin{aligned}
\ell_t^\star &= x_t^2 + u_t^2 \\
&= (1+K^2)x_t^2 \\
&= (1+K^2)a_{\text{cl}}^{2t} x_0^2.
\end{aligned}$$

Since

$$K = -\frac{1+\sqrt{5}}{2},$$

we have

$$K^2 = \frac{3+\sqrt{5}}{2}, \qquad 1 + K^2 = \frac{5+\sqrt{5}}{2}.$$

Hence

$$\ell_t^\star = \frac{5+\sqrt{5}}{2}\left(\frac{3-\sqrt{5}}{2}\right)^{2t} x_0^2.$$

The optimal total cost is the value function evaluated at x_0:

$$J^\star(x_0) = V(x_0) = Px_0^2 = (2+\sqrt{5})x_0^2.$$

Since $|a_{\mathrm{cl}}| < 1$, the stage costs decay geometrically. Therefore

$$J^\star_{\mathrm{avg,det}} = \lim_{N\to\infty} \frac{1}{N}\sum_{t=0}^{N-1} \ell_t^\star = 0.$$

e) The optimal stochastic feedback gain is the same as in the deterministic problem:

$$K = -\frac{2P}{1+P} = -\frac{1+\sqrt{5}}{2}.$$

To see why the gain is unchanged, use the average-cost Bellman equation

$$\rho + Pz^2 = \min_v \mathbf{E}\left\{z^2 + v^2 + P(2z+v+w_t)^2\right\},$$

where ρ is the optimal average stage cost. Since $\mathbf{E}w_t = 0$ and $\mathbf{E}w_t^2 = \sigma^2$,

$$\mathbf{E}\left[P(2z+v+w_t)^2\right] = P(2z+v)^2 + P\sigma^2.$$

The extra term $P\sigma^2$ does not depend on v, so the minimizing feedback gain is unchanged. The same Riccati equation determines P, and the average-cost constant is

$$\rho = P\sigma^2 = (2+\sqrt{5})\sigma^2.$$

Thus the optimal stochastic average stage cost is

$$J^\star_{\mathrm{avg}} = (2+\sqrt{5})\sigma^2.$$

In the deterministic problem, the closed-loop state converges to zero and the optimal long-run average stage cost is

$$J^\star_{\mathrm{avg,det}} = 0.$$

In the stochastic problem, the feedback gain is the same, but the process noise continually injects energy into the system, so the optimal average stage cost is nonzero:

$$J^\star_{\mathrm{avg}} = (2+\sqrt{5})\sigma^2.$$