INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 5

A Desktop Ball and Beam Whose Only Damping Comes From the Controller

Pages 107–130

Edited by John D. Hedengren

How to cite this chapter. "A Desktop Ball and Beam Whose Only Damping Comes From the Controller," in *Nonlinear Dynamics, Oscillations, and Stabilization*, John D. Hedengren, Ed., vol. 7 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 5, pp. 107–130.

The page numbers in this file are those of the printed volume.

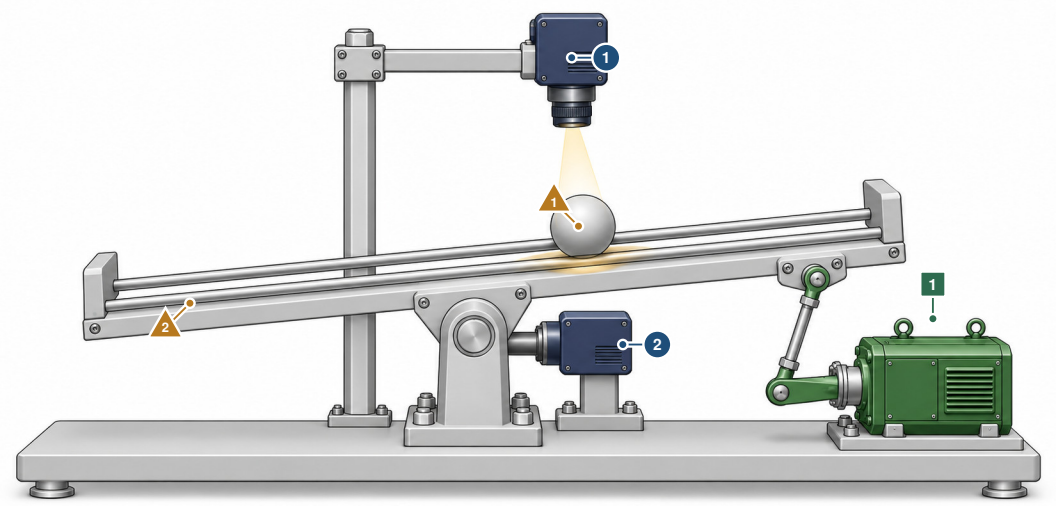
The app for this chapter
https://apmonitor.com/apps/ball-and-beam/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v07-nonlinear-stabilization/#ball-and-beam

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 5

A Desktop Ball and Beam Whose Only Damping Comes From the Controller



The controller moves 1 servo that tilts the beam through a push rod

The controller reads 1 camera above the beam, the ball's position 2 beam angle readback at the pivot

What disturbs it 1 a flick that sets the ball rolling 2 bent beam or unlevel table

The servo tilts the beam to bring the ball to its setpoint, and once the ball rolls almost nothing but the controller slows it. Illustration, not to scale.

IN PLAIN TERMS A Rolling Ball That Nothing Slows Down

The main idea. A ball rests on a long, thin beam that a small motor can tilt. The controller tilts the beam to bring the ball to a chosen spot, and once the ball rolls, almost nothing else can stop it.

The challenge. Suppose the controller only pushes harder the farther the ball is from its spot. Then the ball swings past again and again, wider each time. It also has to know how fast the ball is moving, and a noisy or late reading of the ball's place limits how well it can tell.

Why it matters. Spacecraft, disk-drive heads and machine slides move the same way, with little to slow them.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the history in Section 5.2, the fact on event cameras in Section 5.3, the two try-its and the fact on the inner loop in Section 5.6, and the key points
🕮	**Study**	45 min	the whole chapter with its three pause points and the twelve exercises at its end
⛰	**Practice**	15 min	the first page, the results (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the open challenge

The control question A ball on a tilting beam is two integrators with almost no friction, so a level beam leaves the ball rolling and a loop that knows only the position cannot hold it still. Where does the damping come from, and how much of it do the servo's lag, the sensor's noise, a camera's delay, a slow inner loop and an angle limit take away?

Learning objectives

- Derive the rolling factor of a ball on two rails and the acceleration a beam angle gives it.
- Predict the period and growth per swing of a proportional loop through a lag, and the damping a derivative adds.
- Explain why a derivative on a noisy reading needs a filter, and what the filter and a camera's delay cost.
- Decide from a stopping distance whether any controller can keep a flicked ball on the beam.
- Compare PD, PID and a four-state observer LQR on travel, integrated error and bend offset.

Prerequisites Sections 1.4.7 (lag limit, actuator headroom), 1.4.5 and 1.4.6 (LQR, Bryson's rule, observers) and 1.4.9 (growth through a lag), and Sections 1.6.1 and 1.6.2 (engine and velocity-form PID). Chapter 3 for a desktop kit with its own sample time and limits.

The app https://apmonitor.com/apps/ball-and-beam/

See also Volume 2, Chapter 4 for a sticking valve whose limit cycle no tuning removes, and Volume 2, Chapter 12 for profile feedforward on a machine-tool axis. Volume 3, Chapter 3 for the inner-loop speed of a cascade measured on a fermenter, and Volume 3, Chapter 10 for a position loop that commands an angle.

5.1 Introduction

Tilt a beam by 1° and a steel ball resting on its rails starts to roll at 0.112 m/s^2. Level the beam again and nothing slows the ball except a rolling resistance worth a twentieth of a degree of tilt. In the simulator of this chapter, a ball that has rolled for one second from the middle of a 1 m beam tilted 1° is moving at 10.6 cm/s when the beam is leveled, and it still rolls off the end 4.45 seconds later at 8.8 cm/s. Stopping it takes an opposite tilt held as long as the first one, and a level beam at the instant the speed reaches zero. A hand on a slider is too late for both, which is the first lesson the rig teaches.

The rig is built to correct a common belief, that a controller which pushes harder the farther the output is from its target will bring it there. On a plant with no damping of its own a proportional law is a spring. With 0.30 °/cm of tilt per centimeter of error the ball swings past its target with a period of 3.47 s, close to the 3.43 s of an undamped oscillator, and the servo's 60 ms lag feeds energy in at every swing, so each swing is 1.37 to 1.42 times the last until the ball leaves the beam at 16.41 s. Doubling the gain shortens the period to 2.48 s and loses the ball sooner, at 9.94 s. Only a term that knows the ball's speed adds damping, and the sensor that supplies the speed then limits how much of it the loop can carry.

The shape is general. Two integrators from a force or an angle to a position, with little natural damping, describe a spacecraft turning on its thrusters, a positioning stage on air bearings and a disk-drive head, and the double integrator is a standard test of how controllers behave under off-nominal conditions [1, 2]. A dead zone like the ball's breakaway is the stiction of a valve stem or the static friction of a machine slide, and integral action across one makes a loop hunt between two sides of its setpoint [3]. Volume 2, Chapter 4 models a sticking valve whose limit cycle no tuning removes. A position loop that asks an inner drive for an angle is a cascade, whose inner loop must be faster than the outer one it serves [4]. Each of these appears in the simulator as a case of its own.

The published work on the ball and beam is mostly about one design on one rig, a nonlinear stabilizing law [5–7] or a comparison of model-based and model-free controllers on a laboratory apparatus [8], and the virtual and remote laboratories built around it give access to a rig rather than a way to separate its effects [9, 10]. A student who sees the ball ring on hardware cannot tell how much of the ringing comes from the servo, the sensor, the sampling or the rail. This chapter separates those causes on a simulated desktop ball and beam, brings them in one by one, and measures what each one costs a loop. The rolling-ball model, the shared velocity-form PID, the trajectory feedforward [11],

the dither [12] and the linear quadratic regulator are all standard methods. The measurements are made on one model and one seed.

Section 5.2 follows the rig through the literature. Section 5.3 states the model, with every equation, constant and default, and with the ball's true state logged beside each sensor's reading. Section 5.4 analyzes the linear loop. It shows a P-only growth per swing that follows $e^{\pi\omega(\tau_s+T_s/2)}$ less a fixed loss to the rolling resistance, the damping that only the derivative supplies, and what the sensor, a delay and an inner drive cost. Section 5.5 describes the controllers as coded, together with a loop view whose original-law stability verdict comes from the eigenvalues of the sampled closed-loop matrix. Section 5.6 reports ten cases and a relay race as measured. Among the results are a camera that costs the loop 14.4° of phase at its 5.01 rad/s crossover, a cascade that goes unstable when the inner bandwidth falls below 1.53 times the outer crossover, and a ninth case that no controller can meet. Section 5.7 describes the sensors, servos and controllers of real rigs. Sections 5.8 and 5.9 state the limitations and summarize the results, and the exercises follow in Section 5.10.

5.2 The Ball and Beam in the Laboratory and the Literature

Rolling-ball rigs entered the control literature as laboratory apparatus that deserved analysis in their own right. A ball rolling inside a hoop that a motor turns, a close relative of the beam, was described in 1983 as a laboratory system of that kind [13]. The ball on a beam became a benchmark for nonlinear design a few years later, as the history segment of this section tells. Through the 1990s the plant stayed a proving ground for semi-global stabilization by output feedback [6] and for global asymptotic stabilization [7]. The simulator keeps the centripetal term in its model and leaves it out of its linear analysis, and case 4 below shows the term at work when a noisy derivative shakes the beam.

> ⧖ HISTORY **A Bead on a Wire Defeats Exact Linearization, 1989**
>
> A research memorandum of 1989 took the ball and beam, a common undergraduate control laboratory experiment, as the motivating example for approximate input-output linearization [14], the method that Section 1.4.3 introduces. Its ball slid without friction along the beam, like a bead on a wire, so the beam could turn freely without losing it. Differentiating the position brought in the input with a coefficient that is zero whenever the beam's angular velocity or the ball's position is zero, so the relative degree is not defined and exact input-output linearization fails. Full state linearization fails too. Neglecting the centrifugal term, the $r\dot{\alpha}^2$ of Equation (5.1), gave approximate designs that tracked a large sinusoid where a design on the Jacobian linearization went unstable. The work appeared as a journal paper in 1992 [5].

Laboratory practice followed a simpler path. A proportional-derivative law on the measured position is what most rigs run, and its stability on the ball and beam has been analyzed directly [15]. The difficulty in practice is the speed. A laboratory study that modeled the rig with its nonlinear and coupling effects tuned a linear quadratic regulator by a genetic algorithm, and because the position sensor was noisy it estimated the ball's speed with a state observer before comparing the result with PID control [8]. The original laws offer the same two routes to the speed, a filtered difference of the reading and an α-β estimator. Access then widened through virtual and remote laboratories that put a ball and beam behind a web page [9], one of them inside the Moodle course management system [10], and a desktop heat-transfer kit built around a microcontroller showed what a small, cheap, sampled rig offers as a benchmark for modeling and control [16].

Three remedies the cases teach come from other machines. Dead zones and the hunting that integral action makes across them were catalogued for machines with friction, along with dither as a compensation [3]. The theory of dither treats a fast periodic signal as a way to smooth a nonlinearity seen by a slower loop [12], and process practice adds short pulses to the command of a sticking control valve for the same purpose [17]. Motion systems plan their setpoint as a trajectory and feed forward the effort that trajectory needs, which leaves the feedback loop only the error [11]. Volume 2,

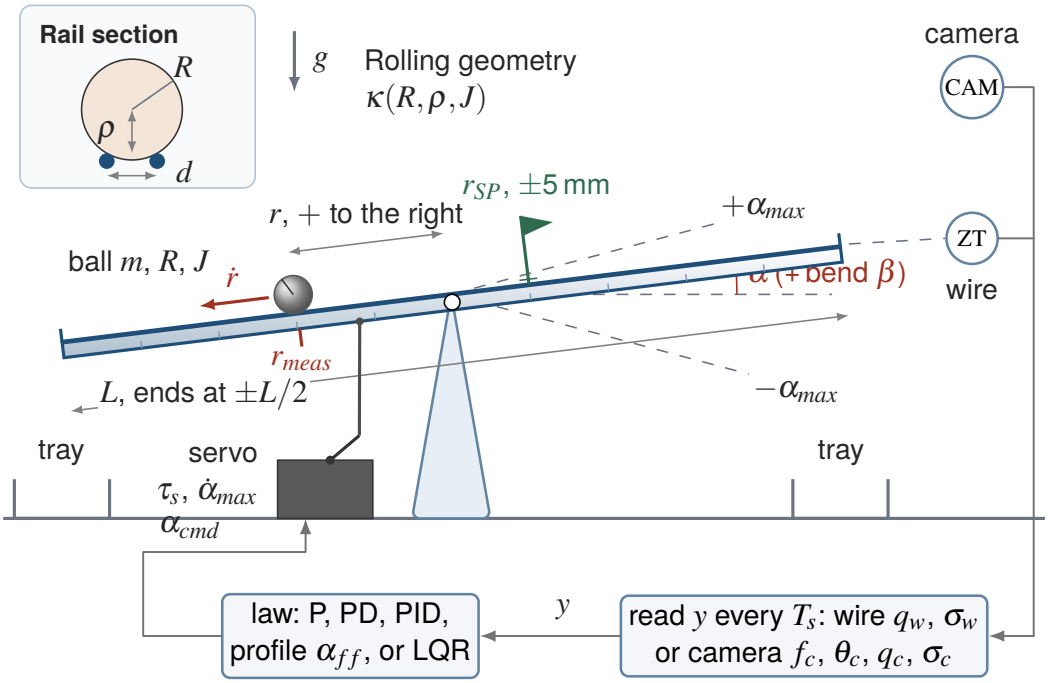


Figure 5.1 The simulated ball and beam, as the Rig view draws it, and the kit's chain from the reading to the servo, not to scale. The beam of length L is tilted by α (with an unknown bend β) within $\pm\alpha_{max}$ by a servo with lag τ_s and rate limit $\dot{\alpha}_{max}$ that receives the command α_{cmd}. The ball of mass m, radius R and inertia J sits at r with speed $\dot{r}$, and the section shows it touching rails d apart at the rolling radius ρ, which sets the rolling factor κ. The target flag marks r_{SP} with its ± 5 mm band, the short tick below the rail is the wire's quantized reading r_{meas}, and the trays catch a ball past $\pm L/2$. Every T_s the kit reads y from the wire (ZT) or the camera (CAM) and runs its law.

Chapter 12 applies a jerk-limited profile with its feedforward to a machine-tool axis. Case 6 brings the dead zone and the dither to the beam, and cases 5 and 11 the profile and its feedforward.

5.3 The Kit as Simulated

The simulated kit is a 1.0 m aluminum beam on a center pivot, tilted by a hobby servo on the desk through a horn and a push rod 12 cm left of the pivot, as shown in Figure 5.1. Two rails 12.7 mm apart carry a solid steel ball 25.4 mm across (66 g), or a hollow ping-pong ball 40 mm across (2.7 g), and a resistance wire between the rails reads the ball's position over ± 45 cm. A camera above the beam can replace the wire. The position r is measured along the beam from the pivot, positive to the right, and the beam angle α is positive when it raises the right end, so a positive angle rolls the ball toward $-r$. Past either end the ball drops into a tray. Every dimension, limit and threshold of the rig is a teaching value chosen for a desktop kit rather than taken from any product, which the model dialog and the Instructions page both state. The signals are sorted by role in Table 5.1.

5.3.1 The Ball on Two Rails

A sphere of mass m, radius R and moment of inertia J about its center rolls without slip on two rails a distance d apart. It touches them at the rolling radius ρ, so it turns through r/ρ while it travels r, and

Table 5.1 Signals of the Kit by Role

Symbol	Role	Nominal	Range
r	CV and state, ball position	per case	SP ±45 cm
$\dot{r}$	state, ball speed	0	—
α	state, beam angle	0	±15°
ω	state, beam rate (torque beam)	0	—
α_{cmd}	MV, beam angle command	per law	$\pm\alpha_{max}$
y	measured, wire or camera reading	r	±45 cm wire
$\hat{\alpha}$	measured, servo readback	α	0.01° steps
β	DV, bent beam or unlevel table	0	±2°
$\Delta\dot{r}$	DV, finger flick	20 cm/s	±50 cm/s, event
$\Delta\alpha$	DV, bump to the beam	2°	event
—	DV, sensor dropout	1 s	event
$m/\hat{m}$	DV, heavier ball (torque beam)	1	1 to 3
$\hat{\dot{r}}$	estimate, α-β speed	—	—
PM	output, phase margin of the tuning	—	—

Lagrange's equations for the ball on a beam turning at $\dot{\alpha}$ give Equation (5.1).

$$\left(1+\frac{J}{m\rho^2}\right)\ddot{r} = r\dot{\alpha}^2 - g\sin(\alpha+\beta) - g\sin\alpha_k\, s$$
$$\rho = \sqrt{R^2 - d^2/4}, \qquad \kappa = \frac{1}{1+J/(m\rho^2)} \tag{5.1}$$

where β is a tilt the controller does not know about, α_k is the rolling resistance written as an equivalent tilt, and $s = \pm 1$ is the rolling direction, frozen over each integration step, which keeps the derivative function smooth inside the step. Multiplying through by κ gives the acceleration the code integrates. For a solid ball $J = \frac{2}{5}mR^2$, so $\kappa = 5/7$ on a flat track. On rails one radius apart $\rho = 11.00$ mm for the steel ball and $\kappa = 0.652$, and the hollow ping-pong ball with $J = \frac{2}{3}mR^2$ has $\rho = 18.97$ mm and $\kappa = 0.574$ on the same rails. The rails make the ball spin faster for the same travel, so more of gravity's work goes into the spin. Near rest $\ddot{r} \approx -\kappa g\sin\alpha$, which is 0.558 m/s^2 at 5° and 1.656 m/s^2 at the 15° limit. The term $r\dot{\alpha}^2$ flings the ball outward when the beam turns fast. The rolling resistance, $\alpha_k = 0.05°$ or a coefficient of 0.00087, decelerates a rolling steel ball at 5.6 mm/s^2.

5.3.2 The Breakaway and BALL OFF

Between integration steps a latch decides whether the ball rolls. With the drive $a_d = \kappa\big(r\dot{\alpha}^2 - g\sin(\alpha+\beta)\big)$, a rolling ball stops when it is slower than the creep speed $v_c = 2$ mm/s, or its speed changes sign inside a step, while $|a_d| \le \kappa g\sin\alpha_0$, and a ball at rest starts only when $|a_d| > \kappa g\sin\alpha_0$, as shown in Equation (5.2).

$$\text{stuck} \;\Leftrightarrow\; |\dot{r}| < v_c \text{ and } |a_d| \le \kappa g\sin\alpha_0, \qquad \alpha_0 = 0.4° \tag{5.2}$$

A stuck ball has $\dot{r} = 0$ and $\ddot{r} = 0$. The breakaway α_0 is the dent, seam or speck of dirt on a real rail. It decides where a PD law leaves the ball, within about α_0/K_c of its setpoint, which is 1.33 cm at 0.30 °/cm. When $|r|$ passes $L/2 = 50$ cm the ball has left the beam. The instant is found by interpolating back inside the step, $t_{off} = t_k - \Delta t\,(|r_k| - L/2)/(|r_k| - |r_{k-1}|)$, the event is marked on the trends, the ball falls into the tray in 0.3 s on the drawing, and the run stops when it lands.

5.3.3 The Servo-Driven and the Torque-Driven Beam

The hobby servo is a stiff position loop of its own, seen through the lever as a first-order lag with a rate limit, as shown in Equation (5.3).

$$\dot{\alpha} = \text{clamp}\left(\frac{\alpha_{cmd} - \alpha}{\tau_s},\ \pm\dot{\alpha}_{max}\right), \qquad |\alpha_{cmd}| \le \alpha_{max} \tag{5.3}$$

where $\tau_s = 60\,\text{ms}$, $\dot{\alpha}_{max} = 150^\circ/\text{s}$ and $\alpha_{max} = 15^\circ$. The lever only rescales the servo's angle, so it is drawn and not modeled, and the ball's weight does not move the beam. The advanced torque-driven beam replaces the servo by a motor at the pivot, and the beam then feels the ball, as shown in Equation (5.4).

$$\begin{aligned} (J_b + mr^2)\,\ddot{\alpha} &= \tau - mgr\cos\alpha - 2mr\dot{r}\dot{\alpha} - b_b\dot{\alpha} \\ \tau &= \text{clamp}\big(K_p(\alpha_{cmd} - \alpha) - K_d\dot{\alpha} + \hat{m}gy\cos\alpha,\ \pm\tau_{max}\big) \end{aligned} \tag{5.4}$$

where $J_b = 0.042\,\text{kg}\,\text{m}^2$, $b_b = 0.01\,\text{N}\,\text{m}\,\text{s/rad}$, $\tau_{max} = 3\,\text{N}\,\text{m}$, and the inner proportional-derivative loop runs every millisecond with $K_p = J_b\omega_{in}^2$ and $K_d = \max(0,\ 2\zeta_{in}\omega_{in}J_b - b_b)$ from the bandwidth ω_{in} and damping ζ_{in} the reader sets. Its last term feeds forward the declared ball's weight on the measured position. The inner loop reads the beam angle and rate exactly, as an encoder and a tachometer fine enough to treat as exact would, and the heavier-ball disturbance is the mass $m - \hat{m}$ that the feedforward does not know.

5.3.4 What the Kit Reads

The wire reads the ball once every control interval $T_s = 10\,\text{ms}$ with no delay, and the camera catches a frame every $1/f_c$ on its own clock and reports it θ_c later, as shown in Equation (5.5).

$$\begin{aligned} y_k &= \text{clip}\Big(q_w\,\text{rnd}\Big(\frac{r(t_k) + \sigma_w n_k}{q_w}\Big),\ \pm x_w\Big) \\ y_{cam}(t) &= q_c\,\text{rnd}\Big(\frac{r(t_f) + \sigma_c n_f}{q_c}\Big), \quad t_f + \theta_c \le t \end{aligned} \tag{5.5}$$

where rnd rounds to the nearest integer, $q_w = 0.5\,\text{mm}$, $\sigma_w = 1\,\text{mm}$, $x_w = 45\,\text{cm}$, t_f is the newest frame on the grid k/f_c that has arrived, $f_c = 30\,\text{Hz}$, $\theta_c = 33\,\text{ms}$, $q_c = 2\,\text{mm}$ and $\sigma_c = 0.5\,\text{mm}$. A camera reading is therefore 33 to 66 ms old. The wire draws its unit normal deviates n_k from the engine's generator, seeded at 20261019, and the camera draws n_f from a stream of its own seeded from the same number, so switching sensors or changing one noise level leaves the other's noise as it was. The servo's angle readback $\hat{\alpha}$ is rounded to 0.01°. A sensor dropout freezes the reading for one second.

> ✦ DID YOU KNOW? **A Pencil Balanced by Event Cameras**
>
> Event cameras report each pixel's brightness change as it happens, not whole frames. In an early demonstration two of them watched a pencil standing on a small movable cart, its state was estimated with a latency below 1 ms, and a simple hand-tuned PID controller kept it upright [18]. The kit's camera reports a frame 33 to 66 ms old.

5.3.5 Integration, Starting Point and Replay

Table 5.2 lists the constants. A fourth-order Runge–Kutta step of 1 ms advances the four states, the controller is invoked every millisecond but executes its law only when a full T_s has passed, and the trends record at 50 Hz. Each step follows the order of Section 1.6.1, with the inner loop of the torque beam run beside the controller and the breakaway latch and the BALL OFF test applied after the integration. A case starts with the ball at rest at its starting position, on a level beam in automatic mode or at the manual angle in manual mode, and a run plays at 0.1 to 5 times real time. Reset and the reload button act as Section 1.6.1 describes. The application passes its 51 verification checks, among them $\kappa = 5/7$ flat and 0.652 on rails one radius apart, conservation of $\frac{1}{2}\dot{r}^2/\kappa + gr\sin\alpha$ to 10^{-9} with the beam still and no rolling resistance, a ball that never starts below the breakaway, the P-only period on an ideal servo within 2 % of $2\pi/\sqrt{K_c\kappa g}$, and, with the noise off, halving the step moving the BALL OFF instant of case 2 and the settling time of case 3 by less than 10 ms.

Table 5.2 Model Constants. H Is Hidden, A Adjustable in the Model Dialog

Symbol	Meaning	Value	Range	
L, x_w	beam length, wire range	1.0 m, ±0.45 m	fixed	H
g	gravity	9.81 m/s^2	fixed	H
m, R	steel ball, solid	66 g, 12.7 mm	fixed	H
m, R	ping-pong ball, hollow	2.7 g, 20 mm	fixed	H
d	rail spacing	12.7 mm	0 to 24	A
α_k	rolling resistance	0.05°	0 to 0.3	A
α_0	breakaway angle	0.4°	0 to 1.5	A
v_c	creep speed	2 mm/s	fixed	H
τ_s	servo lag	60 ms	5 to 300	A
$\dot{\alpha}_{max}$	servo rate limit	150 °/s	20 to 600	A
α_{max}	beam angle limit	±15°	1 to 20	A
J_b	beam inertia (torque beam)	0.042 kg m^2	0.01 to 0.2	A
b_b	pivot damping	0.01 N m s/rad	0 to 0.2	A
τ_{max}	motor torque limit	3 N m	0.5 to 10	A
σ_w, q_w	wire noise, quantization	1 mm, 0.5 mm	0–5, 0.1–5	A
f_c, θ_c	camera rate, latency	30 Hz, 33 ms	10–120, 0–200	A
q_c, σ_c	camera pixel, noise	2 mm, 0.5 mm	0.2–10, 0–5	A
κ	rolling factor, steel, ping-pong	0.652, 0.574	derived	
ρ	rolling radius, steel, ping-pong	11.00 mm, 18.97 mm	derived	

5.3.6 Assumptions

One ball rolling without slip on a rigid, straight beam. Slip, bouncing, the ball's wobble between the rails and the beam's flexing are left out. A real ball skids when the beam is jerked faster than friction can spin it up, which would weaken the $r\dot{\alpha}^2$ fling of case 4.

A stiff servo. The servo is a first-order lag with a rate limit and is not moved by the ball. A real hobby servo has a dead band, gear backlash and a speed that depends on its load, and its lever makes the beam angle a nonlinear function of the horn angle.

Rolling resistance and breakaway as a latch. The values of α_k and α_0 are teaching values for a ball that presses on a resistance wire, and a clean steel ball on steel rails rolls with less. They decide where a PD law stops the ball and the size of the hunting in case 6.

Ideal instruments apart from noise, rounding and delay. The wire's noise is white at the control interval, without contact bounce or drift, and the camera's frames are exact snapshots with a fixed latency, without blur, dropped frames or jitter.

Teaching thresholds. The ±5 mm band, the 1 s that makes an arrival, the challenge's 15 legs in 60 s and the rig's limits are teaching values. A real kit's angle limit, servo speed and lag come from its servo's data sheet and its lever geometry. The torque-driven beam has no mechanical stop beyond its command limit.

5.4 Two Integrators and a Lag

5.4.1 The Linear Plant

About rest on a level beam, with the rate limit, the angle limit, the breakaway, the rolling resistance, the quantization and the centripetal term left out, the servo-driven kit from the angle command to the position is Equation (5.6).

$$P(s) = \frac{R(s)}{A_{cmd}(s)} = \frac{-\kappa g}{s^2(\tau_s s + 1)} \tag{5.6}$$

The double integrator contributes −180° of phase at every frequency, the servo takes more, and the sign is the reverse action of the loop. A proportional law $\alpha_{cmd} = K_c(y - r_{SP})$ on an ideal servo

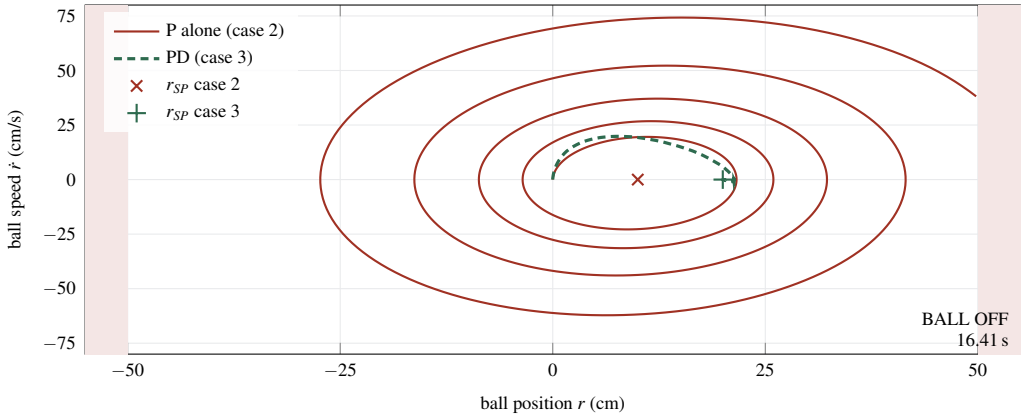


Figure 5.2 The ball's position against its speed from the engine, true values every 20 ms. Case 2 as shipped (solid line, P alone, 0.30 °/cm, setpoint 10 cm from 1 s, marked ×) spirals out at 1.37 to 1.42 a swing until the ball leaves the right end at 16.41 s. Case 3 as shipped (dashed line, PD, $\tau_D = 0.70$ s, setpoint 20 cm from 1 s, marked +) winds in from 1 s to 8 s and stops 2.5 mm past its setpoint. The shaded bands are past the ends of the beam.

makes the ball an undamped oscillator at $\omega = \sqrt{K_c \kappa g}$, a mass on a spring, with a period of 3.43 s at 0.30 °/cm (0.524 rad/m). The servo's lag and the hold of the sampled command act at that frequency as a delay of about $\tau_e = \tau_s + T_s/2 = 65$ ms, and a delay τ_e in $\ddot{r} = -\omega^2 r(t - \tau_e)$ moves the roots to $s \approx \omega^2 \tau_e/2 \pm j\omega$. Each swing therefore grows by $e^{\pi\omega\tau_e}$, which is 1.45 at 0.30 °/cm. The rolling resistance takes a fixed $4 \sin \alpha_k / K_c$ off each full swing, 6.7 mm at that gain, so the measured growth of 1.37 to 1.42 sits a little below the formula, and the sampled loop's least damping ratio is −0.058. The planning value for the kit was a gain of 0.1 rad/m. There the formula gives 1.18 a swing, but the rolling resistance takes 3.5 cm off every 7.9 s swing, and a 10 cm swing spirals in. The kit therefore ships with 0.30 °/cm, where the growth is visible inside its 20 s trend window.

🕮 **Pause.** Try Exercise 5.1 before reading on.

5.4.2 Derivative Is the Only Damping

A derivative term adds $K_c \tau_D \dot{r}$ to the angle, and on an ideal servo the loop becomes $s^2 + \kappa g K_c \tau_D s + \kappa g K_c$, with a damping ratio $\zeta = \omega\tau_D/2$ that the derivative time alone sets at a given gain. At 0.30 °/cm and $\tau_D = 0.70$ s that is 0.64. With the servo's lag the characteristic polynomial is $\tau_s s^3 + s^2 + \kappa g K_c \tau_D s + \kappa g K_c$, and the Routh test requires $\tau_D > \tau_s$ at every gain. This is the lag limit of Equation (1.9): the derivative time must be longer than the lag, and a proportional law alone fails it for any lag. The derivative is the phase lead the loop has, and the Loop view reads the result as a phase margin of 48.5° at 2.70 rad/s. The phase plane makes the two laws visible, as shown in Figure 5.2. Under P alone the orbit spirals out until it reaches the end of the beam, and under PD it winds into the target and stops there on the breakaway.

5.4.3 What the Sensor, the Delay and the Drive Cost

The derivative differences a reading that moves in 0.5 mm steps with 1 mm of noise. Over one 10 ms interval that noise is about 0.14 m/s of false speed, which a derivative time of 1 s at 0.30 °/cm turns into several degrees of command, so a filter is needed, and its lag is paid at the crossover. The camera is a delay. A frame is on average half a frame old when it is read, so the Loop view treats it as $\theta_c + 1/(2f_c) = 49.7$ ms, rounded to five control intervals, and a delay θ costs $\omega\theta$ of phase at the crossover and nothing in gain. A torque-driven beam with its inner loop closed is a second-order lag

Table 5.3 Controller Settings, Defaults and Ranges

Symbol	Setting	Default	Range
—	structure	PD	P, PD, PID, LQR4
K_c	proportional gain	0.30 °/cm	0 to 2
τ_I	integral time, PID only	1.0 s	0 to 30
τ_D	derivative time	0.70 s	0 to 2
τ_f	derivative filter	30 ms	0 to 300
T_s	control interval	10 ms	2 to 50
I_{max}	integral limit	±3°	0.2 to 15
a_p, v_p	profile acceleration, speed	0.8 m/s^2, 0.6 m/s	0.05–3, 0.05–1.5
A, P	relay marks, period of a leg	±30 cm, 6 s	5–45, 1–20
A_d, f_d	dither amplitude, frequency	±1.5°, 10 Hz	0–3, 1–50
$\omega_{\alpha\beta}$	α-β bandwidth	15 rad/s	2 to 60
ω_{in}, ζ_{in}	inner loop, torque beam	20 rad/s, 0.7	1–60, 0.2–1.5
r_s, α_s	LQR position, angle scales	5 cm, 3°	0.5–30, 0.5–15

where the servo was a first-order one, and its phase at the outer crossover comes out of the outer loop's margin.

5.4.4 A Limit No Tuning Overcomes

A ball moving at v toward an end stops in no less than $v^2/(2\kappa g \sin\alpha_{max})$, whatever law decides when to tilt. For the ping-pong ball on a beam limited to ±3° the largest deceleration is 0.295 m/s^2, so a ball flicked at 30 cm/s needs 15.3 cm to stop (15.0 cm with the rolling resistance helping), and case 9 gives it 10 cm. This is the actuator-headroom limit of Section 1.4.7 in the form it takes on a double integrator. It belongs to the plant and the actuator, and no controller removes it.

5.5 The Kit's Controller as Coded

The controller reads only the wire's or the camera's reading y, its age and a dropout flag, and the angle readback $\hat{\alpha}$, with the tachometer rate also supplied to the four-state law. It never reads the ball's true position or speed, which the CSV export logs as columns prefixed `truth_`. It executes once a whole T_s has passed on the 1 ms clock, writes α_{cmd} clamped to $\pm\alpha_{max}$, and holds it until the next execution. All laws are reverse acting, so a ball to the right of its reference asks for a positive angle. The settings are listed in Table 5.3.

5.5.1 The Shared PID

The PID is the shared velocity-form algorithm of Section 1.6.2, written in Equation (5.7) in the symbols of this loop.

$$\begin{aligned} e_k &= y_k - r_{ref,k}, \qquad P_k = K_c e_k \\ D_k &= K_c \tau_D \,(z_{f,k} - z_{f,k-1})/T_s \\ S_k &= S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ \hat{u}_k &= b_k + S_k + I_{k-1} \\ \alpha_{cmd} &= \mathrm{clamp}(\hat{u}_k + \delta_k,\ \pm\alpha_{max}) \\ I_k &= \mathrm{clamp}\big(I_{k-1} + \big[K_i e_k + (u_k - \hat{u}_k)/\tau_I\big]T_s,\ \pm I_{max}\big) \end{aligned} \tag{5.7}$$

where z_f is the measured variable through the exact lag $z_{f,k} = a\,z_{f,k-1} + (1-a)z_k$ with $a = \mathrm{e}^{-T_s/\tau_f}$ ($a = 0$ when $\tau_f = 0$), $z = y - r_{ref}$ with the profile on and $z = y$ without it, $K_i = K_c/\tau_I$ for the PID, and the integral is updated only while $K_i \neq 0$, $u_k = \mathrm{clamp}(\hat{u}_k, \pm\alpha_{max})$ drives the back-calculation, δ_k is the dither, and b_k is the bias that carries the feedforward and the hand-over. Here $\hat{u}_k$ and α_{cmd} play

the parts of v_k and u_k in Equations (1.13) to (1.15), and the tracking time is τ_I. Without the profile r_{ref} is the setpoint, held constant between moves, so a setpoint step gives no derivative kick. Beyond the shared algorithm, the integral alone is clamped to $\pm I_{max}$. With the α-β option the PID's own derivative is off and $K_c\tau_D(\hat{\dot{r}} - v_{ref})$ enters through b_k. The saturation flag compares the unclamped law with $\pm\alpha_{max}$ and is cleared every execution.

In manual mode every law tracks the manual angle, so automatic mode starts without a bump, as in Section 1.6.2. Tracking leaves the manual angle less P and D in the integral, and a P or PD law would then hold that bias for good, which on a plant that rests only on a level beam keeps the ball off its setpoint. On the first automatic execution the part of the integral the law may not keep (all of it for P and PD, whatever lies beyond $\pm I_{max}$ for the PID) moves into b_k and decays as e^{-t/τ_h} with $\tau_h = 1$ s. When the ball of case 1 is caught by hand at −11.35 cm and the loop is switched to automatic at 6 s, it is 0.32 cm from its setpoint at 30 s under PD and 0.41 cm under PID.

5.5.2 Profile, Feedforward and Dither

With the profile on, a new setpoint is reached on a trapezoid, accelerating at a_p to the speed v_p and braking at a_p to stop on the target, and the feedforward is the angle that gives the profile's acceleration, as shown in Equation (5.8).

$$\alpha_{ff} = -\arcsin\frac{a_{ref}}{\hat{\kappa} g} \tag{5.8}$$

where $\hat{\kappa}$ is the rolling factor of the ball the controller is told is loaded. The generator runs $\mathrm{round}(\tau_s/T_s) = 6$ intervals ahead of the reference the PID tracks, so the beam is at the feedforward's angle when the ball needs it, and the PID acts on the tracking error $y - r_{ref}$. At 0.8 m/s^2 the feedforward is 7.18°. The breakaway dither adds a square wave of $\pm A_d$ to the command, its sign taken from the half-period index $\lfloor 2 f_d t + 10^{-6} \rfloor$ so that a 10 Hz wave on the 10 ms interval is exactly five intervals high and five low. A first-order servo at 60 ms passes $1/\sqrt{1+(2\pi f_d \tau_s)^2} = 0.256$ of the fundamental at 10 Hz to the beam.

5.5.3 Speed Estimate and State Feedback

The α-β estimator updates on every fresh reading, as shown in Equation (5.9).

$$\begin{aligned} \bar{r} &= \hat{r}_{k-1} + T_s\hat{\dot{r}}_{k-1}, \qquad \hat{r}_k = \bar{r} + a(y_k - \bar{r}) \\ \hat{\dot{r}}_k &= \hat{\dot{r}}_{k-1} + \frac{b}{T_s}(y_k - \bar{r}) \end{aligned} \tag{5.9}$$

with $a = 2\omega_{\alpha\beta}T_s$ and $b = (\omega_{\alpha\beta}T_s)^2$, the small-$\omega T$ form of a double pole, so 0.30 and 0.0225 at 15 rad/s and 10 ms, which place the estimate's error poles at 10.8 and 24.9 rad/s. The comparison law is a discrete linear quadratic regulator on the declared servo-beam model $[r, \dot{r}, \alpha]$, discretized for a zero-order hold at T_s and designed with Bryson's weights $Q = \mathrm{diag}(1/r_s^2, 0, 0)$ and $R = 1/\alpha_s^2$ by the shared Riccati module, as Section 1.4.5 describes. Chapter 7 designs a discrete regulator with the same rule on a cart and pole. It acts as $\alpha_{cmd} = \alpha_{ff} - K[y - r_{ref},\ \hat{\dot{r}} - v_{ref},\ \hat{\alpha} - \alpha_{ff}]$, and at the defaults $K = [-1.0283\ \mathrm{rad/m},\ -0.6254\ \mathrm{rad\,s/m},\ 0.2166]$. The beam rate is not a state of a first-order servo, so the law has three gains, and it reads the readback and the estimate, never the truth.

For the original laws, the Loop view writes each as a discrete state-space exactly as it runs, closes it around Equation (5.6) sampled at T_s, decides stability from the eigenvalues of that sampled closed loop, and takes the phase margin where the open-loop magnitude last drops through unity.

5.5.4 Four-State Feedback on the Torque Rig

The additional LQR4 uses $x = [r, v, \alpha, \omega]^T$ and commands the angle $u = \alpha_{cmd}$ through the existing inner motor loop. It is distinct from the three-state servo comparison above. Its local model assumes

Table 5.4 Four-State Design Used Throughout the Comparison

Setting	Value
Decision interval T_s	10 ms
Q diagonal inverse-square scales	3 cm, 0.3 m/s, 10°, 2 rad/s
R inverse-square command scale	5°
Observer bandwidth ω_o	12 rad/s
Inner bandwidth, damping	20 rad/s, 0.7
Adjustable position scale	0.5 to 30 cm
Adjustable command scale	0.5 to 15°
Adjustable observer bandwidth	2 to 40 rad/s

a level beam, small angles, nominal weight cancellation and negligible friction:

$$\begin{aligned} \dot{r} &= v, \qquad \dot{v} = -\kappa_n g_n \alpha, \\ \dot{\alpha} &= \omega, \quad \dot{\omega} = \omega_{in}^2 (u - \alpha) - 2\zeta_{in}\omega_{in}\omega. \end{aligned} \tag{5.10}$$

Declared steel geometry gives $\kappa_n = 0.652173913$ and $g_n = 9.81\,\mathrm{m/s^2}$. The selected ball and controller settings ω_{in}, ζ_{in} are known, while plant-parameter edits remain mismatch. The nonlinear plant and unchanged inner loop remain Equations (5.1) and (5.4).

For $\dot{x} = Ax + Bu$, the shared Riccati routines compute $\Phi = \exp(AT_s)$ and $\Gamma = \int_0^{T_s} \exp(A\tau)B\,d\tau$. For regulation at a fixed equilibrium, the discrete regulator minimizes $\sum_{k=0}^{\infty}(\tilde{x}_k^T Q \tilde{x}_k + \tilde{u}_k^2 R)$ on this unclipped model [19], where $\tilde{x} = x - [r_{ref}, v_{ref}, \alpha_{ff}, 0]^T$ and $\tilde{u} = u - \alpha_{ff}$. The initial design is retained across every case: The gain is $K = [-2.73608860, -1.20491006, 0.60986450, 0.02793657]$, with units rad/m, rad s/m, 1 and s, respectively. The applied command is

$$u_k = \mathrm{clamp}\big(\alpha_{ff,k} - K(\hat{x}_k - x_{ref,k}) + b_{h,k},\ \pm\alpha_{max}\big). \tag{5.11}$$

The model and gain are cached separately for each engine and recomputed on explicit design-setting changes. Profile feedforward uses nominal κ_n, g_n in Equation (5.8), with the arcsine argument clipped to $[-1, 1]$, and a lead of $\mathrm{round}(2\zeta_{in}/(\omega_{in}T_s))$ intervals. There is no future disturbance information or error integral.

The Luenberger observer [20] of Section 1.4.6 first predicts $\bar{x}_k = \Phi\hat{x}_{k-1} + \Gamma u_{k-1}$, then replaces its angle and rate by the current encoder and tachometer readings. For a fresh position captured at $t_c = t_k - d_k$, its residual is $\varepsilon_k = y_k - \bar{r}(t_c)$, where d_k is the reported age and $\bar{r}(t_c)$ is interpolated from estimated history. With h the larger of T_s and the time since the last corrected capture,

$$\begin{aligned} p &= \exp(-\omega_o h), \quad a = 1 - p^2, \quad b = (1-p)^2, \\ \hat{r}_k &= \bar{r}_k + (a + d_k b/h)\varepsilon_k, \qquad \hat{v}_k = \bar{v}_k + (b/h)\varepsilon_k. \end{aligned} \tag{5.12}$$

History after t_c receives the correction transported to each stored time. Held/dropout frames do not correct. A 1 s interpolation window is capped at 1024 entries. Entry resets the calendar and initializes $[y, 0, \hat{\alpha}, \hat{\omega}]^T$ from instruments. Frames predating activation wait for history warm-up, while dropout entry holds until a real reading. True ball velocity is never an input.

Nominal wire feedback poles are $0.97016824 \pm 0.02700964i$ and $0.83613899 \pm 0.10361177i$. Current-update observer poles are $\exp(-\omega_o T_s) = 0.88692044$ twice and zero twice. These results describe the frictionless, unsaturated wire model, not delayed-camera or nonlinear stability. The original frequency-margin plots are withheld under LQR4.

LQR4 never changes the physical rig. On the servo it visibly holds the previous command with zero activity. It takes precedence over the stored original LQR switch, which is restored on return to P, PD or PID. Matched PD and PID runs therefore explicitly turn that switch off. Manual tracking, entry, recovery and departure seed the first sampled applied command, with transient bias b_h washing

Table 5.5 The Ten Cases and the Challenge

#	Configuration and events	Teaching point
1	Manual, beam at 1°, ball at rest at 0	it always overshoots
2	P 0.30 °/cm; setpoint 0 to 10 cm at 1 s	a spring, and the lag
3	PD, τ_D 0.70 s, τ_f 30 ms; to 20 cm at 1 s	derivative damps
4	PD, τ_D 1.00 s, $\tau_f = 0$; to 20 cm at 1 s	noise against phase
5	quick PD 1.00 °/cm, 0.40 s; relay ±30 cm every 6 s	profile and feedforward
6	PD, beam bent $\beta = 0.5°$	offset, hunting, dither
7	quick PD on the camera; to 20 cm at 1 s	phase the delay eats
8	quick PD, torque beam ω_{in} 20 rad/s; to 20 cm at 1 s	inner loop faster
9	ping-pong ball, camera, ±3°, at 40 cm; flick 30 cm/s at 1 s	infeasible by arithmetic
10	LQR4, torque beam; to 20 cm at 1 s, flick 20 cm/s at 8 s, bend 0.5° at 16 s	estimation, travel and offset
11	P 0.30 °/cm; relay ±35 cm every 4 s for 60 s; flick 20 cm/s at 20 s; $\beta = 0.5°$ from 35 s	the relay race

out over 1 s. Startup in AUTO instead computes feedback at $t = 0$. Invalid instruments, design or calendar reject an update, hold a finite bounded applied command, mark the hold and visibly rearm on valid data. This adds no error integral. Existing angle-command and applied-torque caps remain, without a mechanical angle stop.

5.6 Eleven Cases, Measured

The cases appear in Table 5.5. The numbers in this section come from headless runs of the engine the browser uses, on the default seed, with each case either unchanged or carrying the single change a sentence names entered before Run. Overshoot is measured on the true position past the new setpoint, and settling time from the setpoint change until the ball stays within ±5 mm. Each case produces the same run when it is the first one opened, when it follows another case from the menu, and after Reset.

5.6.1 By Hand and Proportional Alone

As shipped, case 1 starts the beam at 1° with the ball at rest in the middle, and the ball leaves the left end at 3.07 s at 32.6 cm/s. Leveling the beam at 1 s, when the ball is 5.3 cm left of center at 10.6 cm/s, only slows the escape to 5.45 s. A tilt of −1° from 1 to 2 s and then a level beam stops the ball 11.35 cm left of center. Case 2 moves the setpoint of the P-only loop to 10 cm at 1 s. The ball crosses it and comes back with periods of 3.46 to 3.48 s between swings on the same side, against 3.43 s for the ideal servo, swings 11.6, 16.0 and 22.3 cm past the target on the right, and leaves the beam at 16.41 s at 37.5 cm/s. The Loop view gives a phase margin of −6.8°. With a 5 ms servo the swings shrink from 10.0 to 9.4 cm over 20 s, slowed only by the rolling resistance. Halving the gain to 0.15 °/cm lengthens the period to 4.88 s and slows the growth to 1.16 a swing, rising to 1.25 as the swings widen and the fixed loss to the rolling resistance becomes a smaller share, so the ball lasts until 37.56 s. Doubling it to 0.60 °/cm gives 2.48 s and BALL OFF at 9.94 s. A higher gain makes the same lag a larger share of a shorter period.

> **TRY IT IN THE APP** **2 · Proportional Alone: Spring and Lag**
>
> Load the scenario and press Run, and watch the Phase view. After the setpoint moves to 10 cm, the orbit of position against speed spirals outward: each swing past the target is 1.37 to 1.42 times the last, about 3.47 s apart, until the ball leaves the beam at 16.41 s.

5.6.2 Derivative, and the Sensor That Limits It

With $\tau_D = 0.70$ s the 20 cm step of case 3 overshoots by 1.31 cm, and the ball rests there for a while, as listed in Table 5.6, and Figure 5.3 draws three of its columns. A 1.31 cm error asks for 0.39° at this gain, just inside the 0.4° breakaway, so the ball waits until the wire's noise tips the beam over it, and it is inside the band to stay 4.09 s after the step. Longer derivative times raise the margin

Table 5.6 Loop Margins and Step Responses on the 20 cm Steps. PM and ω_c From the Loop View, Overshoot and Settling From the Run

Case	Setting	PM (°)	ω_c (rad/s)	Overshoot (cm)	Settling (s)
3	τ_D 0.30 s	22.5	1.99	9.66	—[a]
	τ_D 0.50 s	38.0	2.28	4.47	—[a]
	τ_D 0.70 s (shipped)	48.5	2.70	1.31	4.09
	τ_D 1.00 s	55.6	3.49	0.03	2.62
4	$\tau_f = 0$ (shipped)	60.2	3.43	9.43	—[b]
	τ_f 30 ms	55.6	3.49	0.03	2.62
	τ_f 100 ms	43.2	3.51	0.00	3.42
	τ_f 200 ms	29.7	3.37	0.00	—[a]
7	wire	38.3	5.01	1.28	1.44
	camera (shipped)	24.0	5.01	5.75	1.69
	camera, 0.60 °/cm, 0.50 s	32.1	3.83	2.90	1.72
8	ω_{in} 20 rad/s (shipped)	34.2	5.20	1.23	1.31
	ω_{in} 15 rad/s	26.7	5.18	2.51	1.25
	ω_{in} 10 rad/s	11.4	5.08	7.79	3.97
	ω_{in} 6.7 rad/s	−9.1	4.76	17.1	—[c]

[a] Stops on the breakaway outside the band: 6.3 mm, 7.2 mm and 7.6 mm short. [b] Flung 66 mm to 94 mm beyond the setpoint from 8 s on. [c] Swings between 10.6 cm short and 14.7 cm past for good.

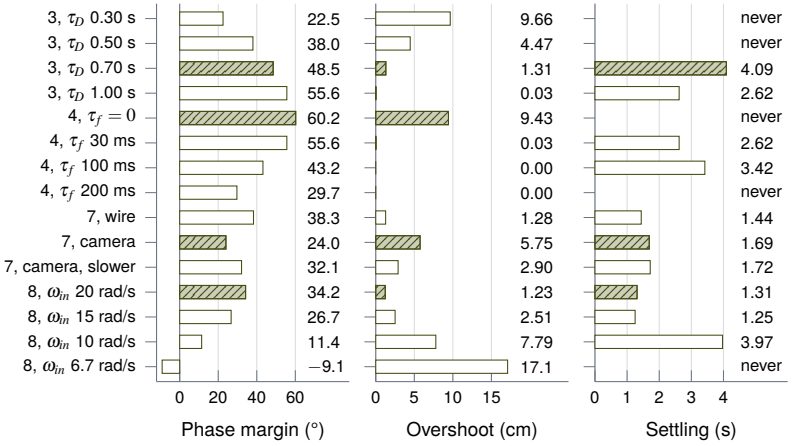


Figure 5.3 The 20 cm steps of Table 5.6: the phase margin from the Loop view (left), the overshoot (middle) and the settling time (right) of each run. Hatched bars are the setting each case ships with and white bars its variants, and "never" marks the five runs that never stay in the ±5 mm band. Notice that the unfiltered derivative of case 4 has the largest margin of all and still never settles, and that the inner loop of case 8 loses its margin while its crossover hardly moves.

and the crossover together, and at 1.00 s the overshoot is 0.3 mm. Shorter ones leave the ball on the breakaway outside the band.

Case 4 runs $\tau_D = 1.00$ s with the filter off. The Loop view rates it best, at 60.2°, because the margin describes small, slow motions. In the run the command jumps 7.43° RMS from one interval to the next, the servo spends 0.56 s of the first 15 s at its rate limit, and the beam shakes at 67 °/s RMS. That shaking flings the ball outward through $r\dot{\alpha}^2$ until a mean tilt of 2.22° holds it, against $r\langle\dot{\alpha}^2\rangle/g = 2.22°$, so the ball sits 66 to 94 mm beyond its setpoint. A 10 ms filter cuts the chatter to 3.79° and still leaves the ball 18 to 24 mm out. At 30 ms the chatter is 1.45° and the ball settles in 2.62 s for a margin of 55.6°. At 200 ms the chatter falls to 0.30° but the margin has fallen to 29.7° and the ball stops 7.6 mm short. The filter buys quiet with phase.

🕮 **Pause.** Try Exercise 5.2 before reading on.

Table 5.7 Matched Torque-Rig Runs, Seed 20261019. Each Group Lists PD, PID and LQR4 in That Order

	IAE (m s)			Logged travel (deg)			Last 5 s error (mm)			Motor cap time (s)		
Case	PD	PID	LQR4	PD	PID	LQR4	PD	PID	LQR4	PD	PID	LQR4
Step	0.09956	0.13432	0.12204	2386.6	2370.7	243.4	0.30	-0.33	0.84	0.017	0.017	0.016
Flick only	0.02671	0.04749	0.05169	2389.1	2377.4	221.4	0.16	-0.18	-0.59	0.000	0.000	0.000
New lesson	0.17922	0.17201	0.25993	2387.2	2388.9	255.1	-4.34	-0.14	-8.63	0.017	0.017	0.016
Constant bend	0.45983	0.14587	0.26433	1198.6	1204.5	210.0	-15.70	-0.50	-8.96	0.000	0.000	0.000
Mass ×3	0.35549	0.14704	0.28936	2373.1	2363.2	244.9	9.59	-0.23	6.86	0.017	0.017	0.016
Camera step	0.11724	0.15234	0.11383	816.0	875.4	229.3	0.44	0.21	0.37	0.023	0.029	0.016
Camera + dropout	0.42825	0.46312	0.42073	1050.8	1285.5	235.1	0.18	0.21	-0.02	0.076	0.091	0.040
Motor cap 0.5 N m	0.23071	0.22183	0.30677	2373.2	2364.4	275.6	-3.29	0.10	-8.67	7.738	7.813	0.780
Limited angle	0.04797	0.04797	0.04797	3.0	3.0	3.1	—	—	—	0.000	0.000	0.000
Torque race, base tuning	10.27726	0.67968	8.79947	2492.8	63.6	1118.3	-135.41	—	-120.35	0.338	0.022	0.339
Torque race, suggested tuning	10.79869	10.75737	10.99966	4011.0	4023.0	1440.9	-145.59	-141.30	-151.35	0.000	0.000	0.324

Requested windows are 30 s, or 60 s for race variants. All limited-angle runs exit the rail at 1.353 s and stop at 1.654 s. Base-tuning race PID exits at 1.969 s and stops at 2.269 s. Their IAE and travel use these actual short windows, and their final errors are undefined. The camera/dropout variant has the step at 1 s, flick at 8 s and a one-second hold from 8.05 s, without the bend. Mass mismatch triples physical ball mass. The motor-cap case keeps the new lesson events. Limited-angle runs retain the ping-pong ball, camera, 3 deg command cap and 30 cm/s flick. Torque races explicitly change the original physical rig and are not the graded servo challenge.

5.6.3 The Relay, the Bent Beam and the Camera

Case 5 sends the setpoint between −30 cm and +30 cm every 6 s under a quicker PD. A 60 cm step asks the proportional term for 60°, four times the limit, so every leg pins the beam at ±15°, for 3.93 s of the first 24 s in all, and the legs overshoot by 4.9 to 5.2 cm and settle in 2.08 to 2.10 s. With the profile and its feedforward the overshoot is 0.6 to 0.8 cm, settling takes 2.16 to 2.38 s, and the beam never reaches its limit. All four legs arrive within ±5 mm both ways, but only the profiled loop has angle to spare.

> **TRY IT IN THE APP** **5 · The Relay: Back and Forth on Purpose**
>
> Load the scenario and press Run. Every 60 cm leg pins the beam at ±15°, and the legs overshoot by 4.9 to 5.2 cm. Press Reset, turn Profile and feedforward on and run again: the overshoot falls to 0.6 to 0.8 cm, and the beam never reaches its limit.

In case 6 the beam is bent 0.5° and the PD stops the ball 15.8 mm short of its setpoint, where the bend and the proportional tilt leave less than the breakaway. Switched to PID with $\tau_I = 1$ s, the integral tilts the beam until the ball breaks loose, but by then it has tilted past what holding the ball needs, so the ball rolls through the setpoint and sticks on the other side. The stops alternate sides, 4.7 to 7.9 mm from the setpoint, every 2.2 to 4.0 s, and the HUNTING status, three alternating stops within 15 s, comes on at 7.18 s. The 10 Hz dither keeps the ball from sticking for long, and it then drifts slowly within ±4.0 mm without another hunting alarm. A dither of ±0.5° is too small, and HUNTING returns at 13.28 s. At 20 Hz the servo passes about half as much of the wave, and the ball again holds on alternate sides 5.1 to 9.8 mm from the setpoint, although its edges unstick the ball often enough that the detector stays quiet.

Case 7 puts the quick PD on the camera. At the 5.01 rad/s crossover the 50 ms delay is 0.25 rad, and the margin falls from 38.3° on the wire to 24.0° on the camera, a loss of 14.4°. The 20 cm step then overshoots by 5.75 cm instead of 1.28 cm. More derivative is no cure, because the camera's margin peaks at 24.7° at $\tau_D = 0.45$ s and falls again as the crossover moves up into the delay. A slower loop, 0.60 °/cm and 0.50 s, recovers 32.1°.

5.6.4 The Cascade and the Limit

On the torque-driven beam of case 8 the quick PD crosses over at 5.20 rad/s with 34.2° of margin when the inner loop has a bandwidth of 20 rad/s, and the step settles in 1.31 s. The outer loop is unstable for inner bandwidths below 7.93 rad/s, 1.53 times its crossover. At a third of the shipped

Table 5.8 The Relay Race Over 60 s, Default Seed

Configuration	Arrivals	Note
P 0.30 °/cm (shipped)	0	off 1.19 s
PID 1.30 °/cm, τ_I 3 s, τ_D 0.30 s, τ_f 40 ms, profile (suggested)	15	passes
suggested without the integral (PD)	12	legs 10, 12, 14
suggested without the profile	0	off 1.26 s
suggested with the filter off	0	7.90 s at limit
case 3's PD on the profile	0	—
LQR on the profile, r_s 2 cm, α_s 5°	15	—

bandwidth, 6.7 rad/s, the margin is −9.1° and the ball swings between 10.6 cm short of its setpoint and 14.7 cm past it for as long as the run lasts, with the command at its limit for 11.0 s of the 20 s. The rule of thumb that an inner loop should be three to five times faster than the outer one is the margin kept above that 1.53. Volume 3, Chapter 3 measures the same rule on a fermenter, and Volume 3, Chapter 10 closes a position loop that, like this one, commands an angle.

> **✦ DID YOU KNOW? Faster Than the Crossover, and Still Unstable**
>
> From the simulator: on the torque-driven beam of case 8, an inner bandwidth of 6.7 rad/s is still faster than the outer loop's crossover of 4.76 rad/s. Yet the phase margin is −9.1°, the ball overshoots by 17.1 cm and it never settles (Table 5.6).

Case 9 flicks the ping-pong ball toward the near end at 30 cm/s with the beam held to ±3°, and the ball leaves at 1.36 s. It does so in all 145 configurations of a verification sweep, P, PD and PID from 0.3 to 10 °/cm with derivative times of 0, 0.3 and 1 s, other integral times, the α-β speed, faster sampling, the state feedback at twenty weightings on both sensors and a bang-bang law on the wire, with BALL OFF between 1.27 and 1.41 s. The 15.3 cm stopping distance of Section 5.4 decides it before any controller is chosen. Given ±10°, the state feedback on the wire with $r_s = 1$ cm and $\alpha_s = 10°$ stops the ball at 47.9 cm.

🕮 **Pause.** Try Exercise 5.3 before reading on.

5.6.5 The Relay Race

The challenge carries the ball between marks at ±35 cm, the command alternating every 4 s for 60 s, through a 20 cm/s flick at 20 s and a 0.5° bend from 35 s. A leg counts as an arrival when the ball has been within ±5 mm of its mark for the last second before the next command, judged on the true position every millisecond, and BALL OFF scores zero. The configurations are compared in Table 5.8. The shipped P-only loop pins the beam at 15° on the first 70 cm step and loses the ball at 1.19 s. The suggested loop arrives on all 15 legs, each inside the band 2.48 to 2.61 s after its command with a worst error of 1.5 to 3.4 mm over its last second, and it scores 15 at the default seed and seeds 1 to 50, and at 21 seeds with the step halved, while the shipped loop scores 0 at all of them. Every part of the suggested setting is necessary except the step from the shipped 30 ms filter to 40 ms, which also scores 15 at eleven seeds. The race is judged as the case ships it, so changing the relay, the rig, a model constant or a disturbance voids an attempt. At the fastest speed, 5, the 60 s race plays in 12 s.

Without the integral the loop misses legs 10, 12 and 14, three of the legs to the left mark after the bend. Back at 0.30 °/cm, at $\tau_D = 0.70$ s or at $\tau_I = 1$ s, no leg arrives at eleven seeds, apart from one leg at one seed for the short integral time. Case 3's gentle PD fails on the profile because its tilt is inside the breakaway over most of the last centimeter. The P-only swing of case 2 just before the ball leaves the beam is shown in Figure 5.4.



Figure 5.4 Proportional control alone, case 2 on the Phase view, captured at $t = 16.30$ s, 0.11 s before the ball leaves the beam (the header that prints the time is cut). The picture is a composite of two parts of the screen: the Phase view above and the trends below. The badge row (1) reads the ball, the setpoint, the beam and the speed estimate, and calls the loop's phase margin unstable. The orbit (2) of position against speed has spiraled out from the 10 cm setpoint, and the swing panel (3) reads a last swing of −37.3 cm after one of −26.3 cm, a growth of $\times 1.42$ and a period of 3.48 s, with the loop at rest unstable. The ball is at 45.4 cm (the wire reads 44.9 cm) moving right at 49.4 cm/s, which the α-β estimate puts at 60.9 cm/s, and the beam is at 9.55° on its way to the commanded 10.46°. The 10 s trends show the swings growing in position about the dashed setpoint (4), in angle and in speed, where the estimate (5) runs beside the true speed.

5.6.6 Observer LQR Against PD and PID

Case 10, Four states, one gain vector, combines a 20 cm step at 1 s, a 20 cm/s flick at 8 s and a 0.5° bend from 16 s. Table 5.7 matches torque rig, inner settings, instruments, noise, seed, initial state and events. PD/PID keep the source case's tuning, or the named suggested tuning, with the old LQR switch off. LQR4 retains Table 5.4. Scorecard IAE uses raw meters. Logged command travel sums 50 Hz samples, not all 100 Hz moves or physical wear. Final error averages true $r - r_{SP}$ over the last 5 s, and motor-cap time counts requested torque beyond its limit.

In the new lesson LQR4 reduces logged travel from PD's 2387.247° to 255.059°, but increases IAE from 0.1792172 to 0.259 932 2 m s and leaves −8.629 mm mean error instead of −4.344 mm. PID leaves −0.139 mm with IAE 0.172 012 2 m s. At 24 s both balls are stationary, at 19.561 cm under PD and 19.137 cm under LQR4, while the latter's observer estimates 0.931 cm/s. An unknown input and friction can bias the nominal velocity estimate. Static feedback, estimation, chosen weights and the breakaway all affect the outcome, so this offset is not attributed to one cause. The paired trajectories

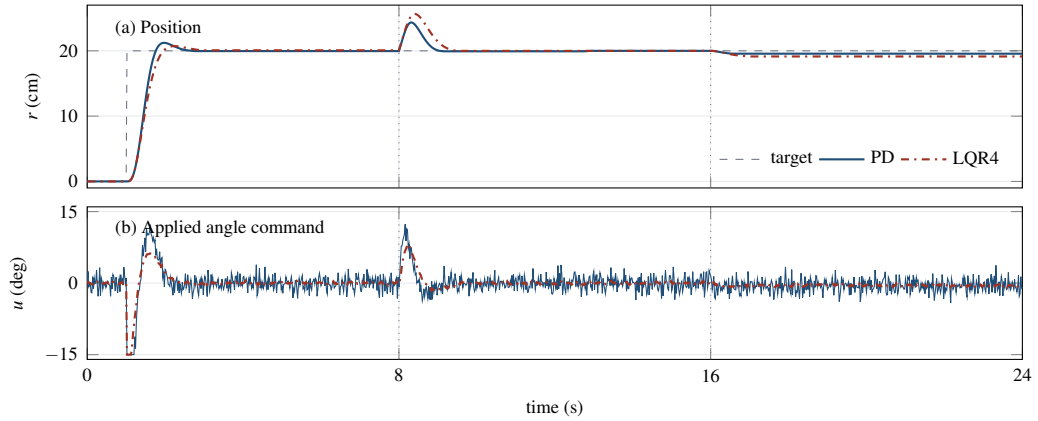


Figure 5.5 Common torque-rig lesson under PD (solid) and LQR4 (dash-dotted), seed 20261019, with the target as a thin dashed line. The step at 1 s, flick at 8 s and bend at 16 s give less logged command motion under LQR4 but a larger final position error. Every plotted point is a 50 Hz historian sample through 24 s. These traces have a shorter window than the 30 s table, and position is simulation truth.

are shown in Figure 5.5.

A fixed analytical diagnostic eliminates static angle feedback to give $K_c = -K_1/(1 + K_3) = 1.699\,576\,95\,\text{rad/m}$ and $\tau_D = K_2/K_1 = 0.440\,376\,84\,\text{s}$. With the original α-β estimator at 12 rad/s, it gives lesson IAE 0.346 513 0 m s, travel 495.657° and mean error −4.495 mm. This algebraic diagnostic still has a different estimator and is not an optimized or stronger PD comparator.

The step and push favor PD's IAE, and PID leaves smaller bend/mass errors. With a 0.5 N m motor cap LQR4 saturates less but has larger IAE and a true angle peak of 18.841° despite a 15° command cap. All three laws lose the limited-angle ball. Failed metrics end at the actual stop, including the 0.3 s fall, without final-window error or settling. Their shorter IAE window is not a win.

All 44 torque comparisons, including the analytical diagnostic, repeat at half the integration step. The lesson's travel/error tradeoff persists, but the camera/dropout ranking reverses. At 1 ms, PD/LQR4 IAE is 0.428 246 1 m s and 0.420 727 6 m s. At 0.5 ms, it is 0.421 755 8 m s and 0.441 870 6 m s. A robust camera advantage is not established. Torque races are distinct from the original graded servo race, which retains scores 0 shipped and 15 suggested. LQR4 on that unchanged servo holds, has zero activity and scores 0.

The eleven torque LQR4 runs have zero numerical failures. A browser measurement recorded when the law was added gives 7.792 μs per update at 1× and 3.002 μs at 5×, or 0.0779% and 0.1498% of one core, and the initial design costs 2.2 and 1.5 ms. These timings depend on the host and include timer quantization and JIT warm-up.

5.7 Technology in Practice

What Measures It

The one controlled variable is the ball's position along the beam, and no instrument on the kit measures its speed. The resistance wire makes the ball part of a voltage divider, read once every control interval. The camera above the beam is the other route. A frame camera delivers pictures at a fixed rate and must wait for a global exposure time, and motion blur is typical of it on fast motion [18]. The chapter's camera is a clean 33 ms delay on a 30 Hz frame clock, without blur, dropped frames or jitter, so the 14.4° it costs case 7 is a lower bound for a real one, as Section 5.8 states. The beam's own angle comes back from the servo's readback, rounded to 0.01°.

TECHNOLOGY Sensor: A Resistance Wire Between the Rails

On one laboratory ball and beam a high-resistance metal strip runs along the beam with a voltage across it, and the ball connects it to the other rod. The voltage sensed on that rod, from −5 V to +5 V over an effective length of 42 cm, gives the ball's position [8]. The same study found its signals noisy, most of all the resistive displacement sensor's, because of how they are read and wired. A numerical derivative of the noisy signals caused problems, so a filter or a state observer was required for the ball's speed [8]. The chapter's wire carries that noise as 1 mm of white noise with 0.5 mm rounding, and its dropout freezes the reading for a second.

What Moves It

The final element is a hobby servo that tilts the beam through a horn and a push rod. A research rig drives its beam instead with a direct-current motor through a gearbox of three simple gears and a lever arm, reads the gear's angle with a potentiometer over $\pm 180^\circ$, and relates the beam's angle to it through the lever's geometry, a sine relation that small angles make linear [8]. The chapter's lever only rescales the servo's angle, and its torque-driven beam, a motor at the pivot under an inner angle loop, is the kit's version of a motor-driven beam.

TECHNOLOGY Final Element: A Hobby Servo

A hobby servo is a position loop in a small case. A potentiometer inside senses the output shaft, the shaft position is compared with the desired one, and the amplified difference drives a small direct-current motor through gears that turn the shaft to remove the error [21]. Its command is a pulse 0.7 to 2.0 ms wide repeated every 20 ms, and the width sets the position of the output arm. One teaching laboratory removed the internal circuit so a microcontroller could close the loop itself, and on one small servo without end stops the potentiometer turns only about 210° before it slips on the shaft [22]. The chapter's servo is a 60 ms lag with a rate limit, and Section 5.3.6 names the dead band and backlash it leaves out.

Where the Controller Runs

One laboratory rig runs its laws on a desktop computer with analog-to-digital and digital-to-analog boards, and an interface board reads the sensors and applies the motor voltage [8]. A low-cost teaching rig puts the loop on a microcontroller board that drives a servo's motor through two outputs used as an H-bridge, with pulse-width modulation at 490 Hz, while an unmodified servo takes a new command pulse every 20 ms [22]. The chapter's kit runs its law every 10 ms on a 1 ms clock, and the inner loop of the torque-driven beam runs every millisecond. The velocity-form PID, with its profile, feedforward bias and dither, or the regulator that replaces it, is the one function that holds the chapter's structure. Figure 5.6 draws the loop on kit hardware.

What Is Modernizing It

First comes the event camera. Each pixel works on its own and reports a change of brightness as soon as it is detected, with a latency of about 10 µs on the laboratory bench and below a millisecond in the real world, and with less motion blur than a frame camera. New methods are needed to process that output [18]. It attacks the delay of case 7. The second is event-based control, in which the control signal changes asynchronously: an event-triggered controller recomputes when, for example, the tracking error exceeds a threshold, and a self-triggered one decides itself when to run next. The same survey states the principle behind both, that it makes no sense to change the control much faster than new information arrives or the actuators can move [18], the kit's servo and camera seen from the controller's side.

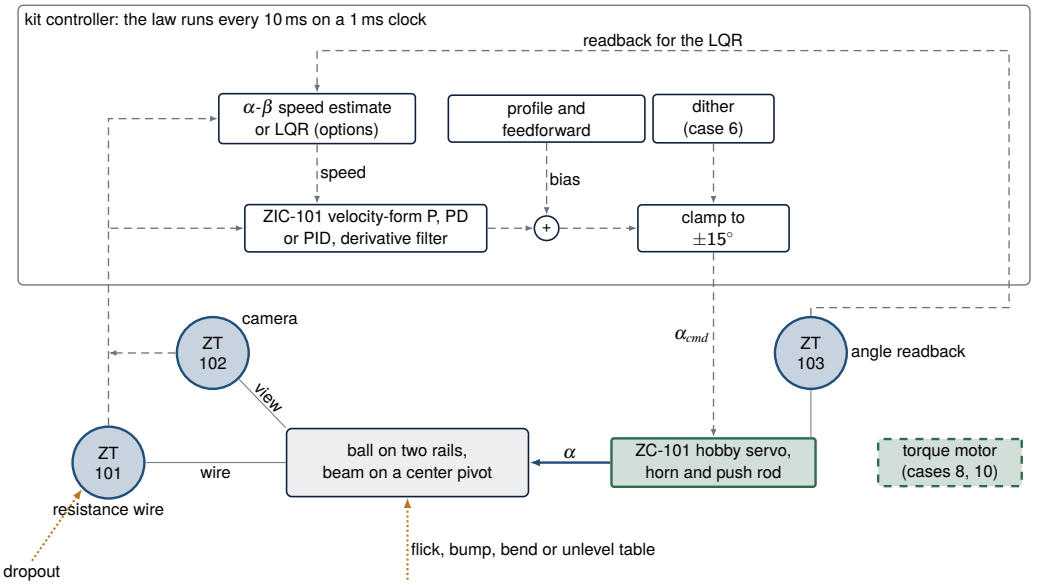


Figure 5.6 The position loop on kit hardware. The resistance wire ZT-101, or the camera ZT-102 that can replace it, reads the ball's position for the velocity-form controller ZIC-101, whose speed term comes from a filtered difference or from the α-β estimate. The profile's feedforward enters as a bias, the dither (case 6) rides on the command, and the clamped angle command goes to the hobby servo ZC-101, whose own position loop tilts the beam through a horn and a push rod; ZT-103 is its angle readback. The dashed torque motor replaces the servo on the torque-driven beam. Notice that nothing in the loop measures the ball's speed: every damping term has to be computed from the position reading.

> ⚠ **IN PRACTICE** **Filter the Derivative Before Tuning It**
>
> An ideal derivative has a high gain for high-frequency signals, so measurement noise makes large variations in the control signal. Practice replaces it with the derivative of a low-pass filtered signal, and the filter time constant is typically chosen as the derivative time divided by a factor in the range 5 to 20. Filtering the measured signal with a second-order filter instead also gives good high-frequency roll-off, which improves robustness [2]. Either way the filter is part of the tuning: in case 4 no filter shakes the beam, and a long one costs phase at the crossover.

5.8 Limitations

The kit is a teaching design whose sizes and limits make a consistent desktop rig without being fitted to any product. The ball rolls without slip on a rigid beam, so skidding, bouncing and beam flex are absent, and the servo is a lag with a rate limit that the ball's weight never moves. The wire's noise is white and Gaussian, and the camera is a clean delay without blur, dropped frames or jitter, so the costs measured in cases 4 and 7 are lower bounds for real sensors.

The linear analysis of the Loop view leaves out the rate limit, the angle limit, the breakaway and the centripetal term, which is why case 4's unfiltered loop reads as the best-damped one while its run is the worst. The camera enters that analysis as a whole number of control intervals. The torque-driven beam has no mechanical stop, and in case 8 detuned the beam reaches 16.44° while its command is held at 15°. The original servo state feedback is a comparison law, and no nonlinear law of the kind the literature built on this plant is implemented.

LQR4 has no integral or disturbance-bias state. Clipping, breakaway, unknown forces and delayed instruments are outside its nominal pole calculation. The observer can be biased, and delayed

initialization is approximate. Numerical hold bounds the applied command without ensuring ball survival. The camera/dropout refinement reversal precludes a robust IAE advantage.

OPEN CHALLENGE **Stabilizing With an Actuator That Runs Out**

A study of the ball and beam under a limited motor voltage judged the control of underactuated systems, with fewer controls than configuration variables, far from closed: design methods do not exist for many that matter in applications. When the linearized plant has unstable modes and the control is bounded, only initial states inside a controllability domain can be brought back, and the basin of attraction of any saturated feedback lies inside it. A law for the straight beam made the two coincide, one for a circular beam only approached it, both in simulation [23]. Progress would be such laws for more machines, tested on hardware. In case 9 the stopping distance decides before a controller is chosen.

5.9 Conclusion

A ball on a beam is two integrators with nothing to slow them, and this simulator models it as a desktop kit with rolling on rails, a breakaway, a lagging servo, a quantized wire or a delayed camera, and an optional torque-driven beam. Proportional control alone swings the ball past its target with a period the gain sets, 3.47 s at 0.30 °/cm, and the servo's lag makes each swing 1.37 to 1.42 times the last, close to $e^{\pi\omega\tau_e}$, until the ball leaves the beam at 16.41 s. Derivative action supplies the damping, 48.5° of margin at the shipped setting, and the wire's noise limits it, since the unfiltered derivative shakes the beam hard enough to fling the ball 66 mm to 94 mm off its setpoint. The camera's delay costs 14.4° of margin, the cascade fails below an inner bandwidth of 1.53 times the outer crossover, integral action hunts across the breakaway until a dither calms it, and no tuning keeps a flicked ping-pong ball on a beam limited to $\pm 3^\circ$. A PID on a profile with its feedforward turns the relay race from 0 arrivals into 15. The four-state torque observer LQR lowers logged command travel in its lesson while losing IAE and bend-offset performance to PID. The friction model is a latch and the instruments are idealized. A fuller model would add a servo with backlash and a dead band, a camera with blur and jitter, and a ball that can skid. The kit shows that a loop holding a ball still has to know how fast the ball is moving.

KEY POINTS

- A ball on a beam is two integrators with nothing to slow them, so a level beam leaves it rolling.
- Proportional control alone swings the ball past its target, and the servo's lag makes each swing 1.37 to 1.42 times the last until the ball leaves the beam.
- Derivative action supplies the damping, and the wire's noise limits how much of it the loop can carry.
- The camera's delay costs 14.4° of margin, the cascade fails below an inner bandwidth of 1.53 times the outer crossover, and no tuning keeps a flicked ping-pong ball on a beam limited to $\pm 3^\circ$.

Where else this applies. Volume 2, Chapter 4 meets the same dead zone as a sticking valve whose limit cycle no tuning removes, and Volume 3, Chapter 3 measures the inner-loop speed of a cascade on a fermenter.

References

[1] V. G. Rao and D. S. Bernstein, "Naive control of the double integrator," *IEEE Control Systems*, vol. 21, no. 5, pp. 86–97, 2001.

[2] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[3] B. Armstrong-Hélouvry, P. Dupont, and C. Canudas de Wit, "A survey of models, analysis tools and compensation methods for the control of machines with friction," *Automatica*, vol. 30, no. 7, pp. 1083–1138, 1994.

[4] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[5] J. Hauser, S. Sastry, and P. Kokotović, "Nonlinear control via approximate input-output linearization: the ball and beam example," *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 392–398, 1992.

[6] A. R. Teel, "Semi-global stabilization of the 'ball and beam' using 'output' feedback," in *1993 American Control Conference*, 1993, pp. 2577–2581.

[7] C. Barbu, R. Sepulchre, W. Lin, and P. V. Kokotović, "Global asymptotic stabilization of the ball-and-beam system," in *Proceedings of the 36th IEEE Conference on Decision and Control*, vol. 3, 1997, pp. 2351–2355.

[8] M. Keshmiri, A. F. Jahromi, A. Mohebbi, M. H. Amoozgar, and W.-F. Xie, "Modeling and control of ball and beam system using model based and non-model based control approaches," *International Journal on Smart Sensing and Intelligent Systems*, vol. 5, no. 1, pp. 14–35, 2012.

[9] J. L. Diez, M. Valles, and A. Valera, "Virtual and remote laboratory of the ball and beam system," in *2009 EAEEIE Annual Conference*, 2009, pp. 1–6.

[10] M. Guinaldo, L. de la Torre, R. Heradio, and S. Dormido, "A virtual and remote control laboratory in Moodle: The ball and beam system," *IFAC Proceedings Volumes*, vol. 46, no. 17, pp. 72–77, 2013.

[11] P. Lambrechts, M. Boerlage, and M. Steinbuch, "Trajectory planning and feedforward design for electromechanical motion systems," *Control Engineering Practice*, vol. 13, no. 2, pp. 145–157, 2005.

[12] G. Zames and N. Shneydor, "Dither in nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 5, pp. 660–667, 1976.

[13] P. E. Wellstead, "The ball and hoop system," *Automatica*, vol. 19, no. 4, pp. 401–406, 1983.

[14] J. Hauser, S. Sastry, and P. Kokotović, "Nonlinear control via approximate input-output linearization: The ball and beam example," Electronics Research Laboratory, University of California, Berkeley, Memorandum UCB/ERL M89/95, 1989. [Online]. Available: https://www2.eecs.berkeley.edu/Pubs/TechRpts/1989/1300.html

[15] W. Yu and F. Ortiz, "Stability analysis of PD regulation for ball and beam system," in *Proceedings of 2005 IEEE Conference on Control Applications (CCA 2005)*, 2005, pp. 517–522.

[16] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, "Benchmark temperature microcontroller for process dynamics and control," *Computers & Chemical Engineering*, vol. 135, p. 106736, 2020.

[17] T. Hägglund, "A friction compensator for pneumatic control valves," *Journal of Process Control*, vol. 12, no. 8, pp. 897–904, 2002.

[18] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, "Event-based vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 154–180, 2022.

[19] R. E. Kalman, "When is a linear control system optimal?" *Journal of Basic Engineering*, vol. 86, no. 1, pp. 51–60, 1964.

[20] D. G. Luenberger, "Observing the state of a linear system," *IEEE Transactions on Military Electronics*, vol. 8, no. 2, pp. 74–80, 1964.

[21] Electrical & Computer Engineering Department, University of Utah, "ECE 1250 lab 5: Servo (control system)," Laboratory handout. [Online]. Available: https://ctools.ece.utah.edu/ECE1250/ECE1250_Lab05_ServoNew.pdf

[22] C. Cohenour, "A low-cost control system experiment for engineering technology students," in *2017 ASEE Annual Conference & Exposition Proceedings*, Columbus, Ohio, 2017.

[23] Y. Aoustin and A. Formal'skii, "Ball on a beam: stabilization under saturated input control with large basin of attraction," *Multibody System Dynamics*, vol. 21, no. 1, pp. 71–89, 2009.

5.10 Exercises

5.1 ○ Concept check · 2 min

Case 2 runs a proportional law alone at 0.30 °/cm. After the setpoint step each swing past the target is 1.37 to 1.42 times the one before it on the same side, until the ball leaves the beam. What feeds the growth?

(a) The rolling resistance, which pushes the ball the way it is already rolling

(b) The wire's noise, which the proportional term passes straight to the beam

(c) The servo's lag and the held command, which apply each tilt about 65 ms late

(d) The centripetal term $r\dot{\alpha}^2$, which flings the ball outward on every swing

5.2 Concept check · 2 min

In case 4 the derivative filter is off. The Loop view gives this tuning a phase margin of 60.2°, more than any filtered setting, yet the ball comes to sit 66 to 94 mm beyond its setpoint. Why does the margin mislead?

(a) The margin is computed for the camera, while the case runs on the wire
(b) A setpoint step kicks the unfiltered derivative, and the integral winds up
(c) The margin leaves out the servo's lag, which the raw derivative excites
(d) The analysis leaves out the rate limit and the $r\dot{\alpha}^2$ term that the shaking beam excites

5.3 Concept check · 1 min

In case 9 a ping-pong ball resting 10 cm from the end of a beam limited to $\pm 3°$ is flicked toward the end at 30 cm/s. Every one of 145 configurations, from P to state feedback and a bang-bang law, loses it. What decides the outcome?

(a) The stopping distance at the angle limit is longer than the 10 cm the ball has
(b) The camera's 50 ms delay leaves too little phase margin for any stabilizing gain
(c) The light ball's small rolling factor makes the loop gain too low to react in time
(d) The integral term winds up against the angle limit while the ball rolls toward the end

5.4 Theory and calculation · 5 min

The steel ball has radius $R = 12.7\,\text{mm}$ and $J = \frac{2}{5}mR^2$, and it rolls on rails $d = 12.7\,\text{mm}$ apart. Use $g = 9.81\,\text{m/s}^2$.

(a) Find the rolling radius ρ and the rolling factor κ, and compare κ with its value on a flat track. *Checkpoint:* 11.00 mm and 0.652, against $5/7 = 0.714$ flat.

(b) Find the acceleration near rest at a beam angle of 1° and at the 15° limit. *Checkpoint:* $0.112\,\text{m/s}^2$ and $1.656\,\text{m/s}^2$.

(c) The hollow ping-pong ball has $R = 20\,\text{mm}$ and $J = \frac{2}{3}mR^2$. Find ρ, κ and the acceleration at 3° on the same rails. *Checkpoint:* 18.97 mm, 0.574 and $0.295\,\text{m/s}^2$.

5.5 Theory and calculation · 9 min

Use the steel ball's $\kappa = 0.652$, $g = 9.81\,\text{m/s}^2$, the servo lag $\tau_s = 60\,\text{ms}$, the control interval $T_s = 10\,\text{ms}$ and the rolling resistance $\alpha_k = 0.05°$.

(a) A proportional law with $K_c = 0.30\,°/\text{cm}$ (0.524 rad/m) runs on an ideal servo. Find the natural frequency and the period. *Checkpoint:* 1.830 rad/s and 3.43 s.

(b) Treat the lag and the hold as a delay $\tau_e = \tau_s + T_s/2$. Find the growth per swing $e^{\pi\omega\tau_e}$ and the fixed loss $4\sin\alpha_k/K_c$ per full swing. Repeat both at 0.15 °/cm, and estimate the ratio of a swing to the one before it when that one was 11.4 cm. *Checkpoint:* 1.45 and 6.7 mm. At 0.15 °/cm, 1.30 and 13.3 mm, a ratio of about 1.19.

(c) Add a derivative time. Find the damping ratio on an ideal servo at $\tau_D = 0.70\,\text{s}$ and the derivative time that gives $\zeta = 1$. Then include the servo's lag and state the condition on τ_D for stability. *Checkpoint:* 0.64 and 1.09 s, and $\tau_D > \tau_s = 60\,\text{ms}$ at any gain.

5.6 Theory and calculation · 6 min

In case 9 the ping-pong ball ($\kappa = 0.574$) rests 10 cm from the end of a beam limited to $\pm 3°$, and a flick sends it toward the end at 30 cm/s. Ignore every delay and assume the beam is already at its limit.

(a) Find the shortest stopping distance, without and with the rolling resistance $\alpha_k = 0.05°$ helping. *Checkpoint:* 15.3 cm and 15.0 cm.

(b) Find the smallest angle limit that stops the ball within 10 cm, and the largest flick that 3° can absorb in 10 cm. *Checkpoint:* 4.58° and 24.3 cm/s.

(c) Would the steel ball ($\kappa = 0.652$) survive the same flick at 3°? *Checkpoint:* No: it needs 13.4 cm.

5.7 App challenge · 8 min

Load scenario 2, Proportional alone: spring and lag, and keep the Phase view. Use the default seed.

(a) Run as shipped until BALL OFF stops the run. Read the On the beam tile and the swing panel: the last swing, the swing before it on the same side, the growth per swing and the period. *Checkpoint:* Off at 16.41 s, −37.3 cm after −26.3 cm, ×1.42 and 3.48 s.

(b) Press Reset, set Proportional gain to 0.15 °/cm, run, and pause between 12 and 13 s. Read the growth and the period. *Checkpoint:* ×1.17 and 4.88 s.

(c) Run on at speed 5 until BALL OFF. Read the On the beam tile and the growth per swing. *Checkpoint:* Off at 37.56 s, growth ×1.25.

Explain why halving the gain lengthens the period by about $\sqrt{2}$, and why the growth in (c) is larger than in (b).

5.8 App challenge · 7 min

Load scenario 4, The sensor: a derivative of noise. Use the default seed and read the phase margin on the Loop view in each part.

(a) Run as shipped and pause at about 15 s. Read the Beam chatter, Servo at its rate limit and Last overshoot tiles. *Checkpoint:* 7.43° within 0.05°, 0.56 s within 0.03 s, 9.4 cm, PM 60°.

(b) Press Reset, set Derivative filter to 30 ms, and run to about 15 s. Read the same tiles and Since the last move. *Checkpoint:* 1.45°, 0.00 s, 0.0 cm, settled in 2.62 s, PM 56°.

(c) Press Reset, set Derivative filter to 200 ms, and repeat (b). *Checkpoint:* 0.30°, still settling, PM 30°.

State what the filter buys and what it costs.

5.9 App challenge · 7 min

Load the challenge scenario, Challenge, the relay race. Use the default seed and speed 5, at which each 60 s run takes 12 s.

(a) Run as shipped. Read the Arrivals tile. *Checkpoint:* 0, ball off within the first 1.5 s.

(b) Press Reveal suggested tuning, then Reset, and run the whole race. Read Arrivals and Challenge target. *Checkpoint:* 15 of 15, met.

(c) Press Reset, set Controller structure to PD, keep everything else, and run again. Read Arrivals, and find on the ball trend which legs miss. *Checkpoint:* 12 of 15. The misses are legs to the left mark after the bend at 35 s.

5.10 Application area · 4 min

A kit builder considers replacing the 30 Hz camera with 33 ms of latency by a 120 Hz camera with 10 ms of latency, under the quick PD of case 7, which crosses over at 5.01 rad/s.

(a) Find the equivalent delay $\theta_c + 1/(2f_c)$ of each camera and the phase it costs at the crossover. *Checkpoint:* 49.7 ms and 14.3°, then 14.2 ms and 4.1°.

(b) The wire gives this loop 38.3°. Estimate the margin on the faster camera, and name two costs of a real camera that the model leaves out. *Checkpoint:* About 34°. Blur, dropped frames, jitter, lighting and price.

5.11 Application area · 4 min

Case 6 calms the hunting across the 0.4° breakaway with a square wave of $\pm 1.5°$ at 10 Hz added to the command. The servo is a first-order lag of 60 ms, and a square wave of amplitude A has a fundamental of amplitude $4A/\pi$.

(a) Find the amplitude of the fundamental at the beam for $\pm 1.5°$ at 10 Hz, $\pm 1.5°$ at 20 Hz and $\pm 0.5°$ at 10 Hz. *Checkpoint:* 0.49°, 0.25° and 0.16°.

(b) Compare each with the breakaway, and name one cost of dither on a real valve or machine slide. *Checkpoint:* Only the first exceeds 0.4°. Wear and vibration.

5.12 Review the agent · 5 min

An AI assistant reviews case 8, where the quick PD on the ball commands the angle loop of the torque-driven beam, and writes: “The outer loop crosses over at 5.20 rad/s. Any inner angle loop faster than that is fast enough, so the inner bandwidth can be cut from 20 to 6.7 rad/s to save motor effort, with the outer tuning left as it is.” Find the flaw, and say how you would check the claim in the app.