INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 7

Inverted Pendulum Simulator for Balancing a Pole Without Losing the Cart

Pages 158–187

Edited by John D. Hedengren

How to cite this chapter. "Inverted Pendulum Simulator for Balancing a Pole Without Losing the Cart," 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. 7, pp. 158–187.

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

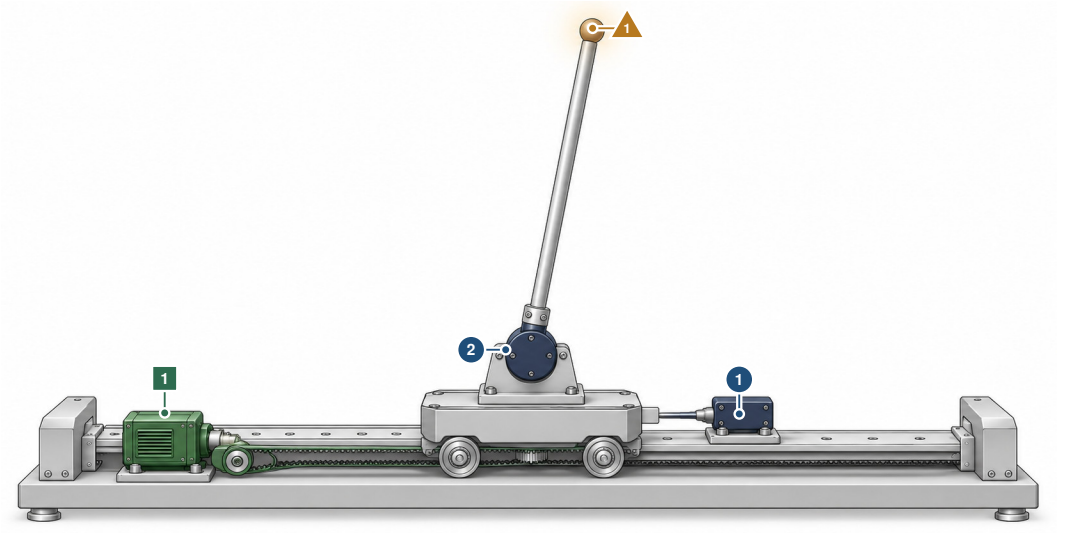
The app for this chapter
https://apmonitor.com/apps/inverted-pendulum/

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

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

CHAPTER 7

Inverted Pendulum Simulator for Balancing a Pole Without Losing the Cart



The controller moves 1 motor and belt that push the cart

The controller reads 1 where the cart is on the rail 2 the pole's angle at its hinge

What disturbs it 1 a tap on the pole

Notice that the only way to catch the leaning pole is to drive the cart under it, while the cart must stop short of the bumpers at both ends of the rail. Illustration, not to scale.

IN PLAIN TERMS **Keeping a Pole Up Without Losing the Cart**

The main idea. A motor pushes a small cart along a short rail, with a pole standing on it on a free hinge. The controller keeps the pole upright by moving the cart under it, like a hand under a stick balanced on a fingertip.

The challenge. Left alone, the pole falls over in less than a second. A controller that watches only the pole keeps it up while the cart drifts into the end of the rail.

Why it matters. Balancing robots, walking machines and rockets standing on their thrust all share this problem.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on a launch vehicle in Section 7.1, the history in Section 7.2, the try-its and the fact on a longer pole in Section 7.6, and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 7.6), Technology in Practice (Section 7.8), Limitations (Section 7.9) and the open challenge

The control question One motor pushes the cart, and the pole responds only to its acceleration, so stopping the fall and keeping the cart on a finite rail share one input. Which loops hold the pole without losing the cart, and which gains, delays and pole changes defeat them?

Learning objectives

- Derive the unstable root, the wrong-way zero and the angle-gain floor $(M+m)g$ from the linearized cart and pole.
- Explain why rail friction makes an angle loop alone run the cart into a bumper, and why every cart move starts the wrong way.
- Use the characteristic polynomial of the cascade to predict which coefficient a longer or a heavier pole attacks, and which fix restores it.
- Relate the stable window of the angle gain to the measurement delay and sample interval.
- Compare the regulator with a predictive controller on the same model, and explain what the terminal cost does for a short horizon.

Prerequisites Sections 1.4.4 and 1.4.5 (state space, LQR, Bryson's rule), 1.4.7 (unstable pole, RHP zero, delay window), 1.6.1 (engine) and 1.6.3 (predictive control). Chapter 6 for a gain floor and ceiling on a plant with one unstable pole.

The app https://apmonitor.com/apps/inverted-pendulum/

See also Volume 3, Chapter 10 and Volume 3, Chapter 9 for a position loop that commands an angle and for a right-half-plane zero that bounds the bandwidth. Volume 9, Chapter 4 for an offset-free predictive controller with a disturbance estimate, and Volume 9, Chapter 2 for predictive control of a delayed integrating plant with a terminal constraint.

7.1 Introduction

Balancing is a control task that engineering keeps meeting in new machines. A two-wheeled robot carries its load above the axle [1], a walking pattern generator treats a biped as a pendulum standing on its support foot [2], a rocket in the first seconds of launch has to be held upright on its thrust before the airstream can steady it [3], and most of the corrections a person makes while balancing a stick on a fingertip are faster than the measured neural delay [4]. The inverted pendulum on a cart is the reduced model behind all of them, and its literature as a control benchmark reaches back to 1960 [5].

> ✦ DID YOU KNOW? **A Cart and Pole Standing In for a Rocket**
>
> In 2011 a space flight center used a laboratory cart and inverted pendulum to simulate a launch vehicle during liftoff. A fluid slosh tank, built to ride on the pendulum, was to test a controller's robustness to slosh that it does not know [6]. Volume 3, Chapter 10 lands a rocket with the same structure, a position loop that commands an attitude.

The pole draws the eye, and the natural conclusion is that a loop which keeps it upright has done the work. It has not. The cart has a position of its own that the pole angle does not reveal, and a loop closed on the angle alone will hold the pole steady over a cart that is creeping toward the end of its travel. In the simulator of this chapter, once that loop has caught a 2° start it keeps the pole within 1.53° of upright while rail friction accelerates the cart into a bumper at 4.11 s. The omission is as old as the apparatus, since two of the first textbook treatments stabilized the angle and left the cart free to drift [3].

The difficulty belongs to a class. One actuator serves two outputs. The pole is open-loop unstable at a rate set by gravity and geometry, and the transfer function from force to cart position carries a right-half-plane zero, so every move of the cart begins in the wrong direction. The unstable pole puts a floor under the angle gain, the lags of the drive and the sensors put a ceiling above it, and the zero limits how fast the cart may be steered [7, 8]. Delay narrows the window between floor and ceiling, which is why the time delays of the sensory and motor pathways limit human balance [9]. Balancing

robots, legged machines and vehicles standing on thrust share that shape, a fast loop against the fall nested inside a slower loop that places the base.

Pendulum rigs are a staple of control teaching [3], and assessments of control education keep returning to laboratory work and to the gap between theory and practice [10, 11]. However, a hardware rig is protected against exactly the failures that teach the most, such as a pole dropped by hand, a cart driven into its end stop and a loop that rocks after a change of pole. Each of those is a limit of the class rather than a mistake in the gains.

This chapter explains which configurations of a cart and pole loop settle and why, using a simulator described closely enough to be rebuilt from the text. The model, the cascade, the regulator and the energy swing-up are all established methods. Section 7.2 traces the history of the apparatus. Section 7.3 gives the nonlinear model, with rail impacts and a deck stop, and the chain of quantizers, delay and filtered differences through which the controller sees it. Section 7.4 derives three results. A closed-form characteristic polynomial of the cascade shows the cart loop raising the angle-gain floor from 11.3 to 17.4 N/rad and separates the two ways a new pole defeats an old tuning. A friction root makes an angle loop alone run away. An exact sampled-data linearization of the implemented loop gives a stable window that agrees with 39 of 40 nonlinear runs. Section 7.5 describes the five control structures, among them a discrete Riccati solution computed in the browser and a model predictive controller placed beside the regulator. That controller shares the regulator's model, weights and Riccati solution, and a completed square shows that the two coincide until a drive, rail or lean limit binds. A second setting removes its terminal weight, which leaves a finite-horizon law whose stability depends on the horizon, and Section 7.5.6 computes where it becomes stable. Section 7.6 runs the scenario sequence, in which each case changes one quantity. It measures what the limits are worth on this plant, as the difference between the regulator and the predictive controller (Section 7.6.8), and how long a horizon the nonlinear rig needs without the terminal weight (Section 7.6.9). Section 7.7 explains why a quantized loop replays exactly and yet depends on the last bit of its arithmetic. Section 7.8 describes the sensors, drives and controllers of real rigs, Sections 7.9 and 7.10 state the limitations and the conclusions, and the exercises follow in Section 7.11.

7.2 Balancing a Pole on a Moving Base

The first known solution was a bachelor's thesis written at M.I.T. in 1960 under the apt title "The Mechanical Seal" [3]. A broom stood on a platform driven through electromechanical integrators on a plotting table, and without position feedback the platform drifted. Its author estimated that a misalignment of one second of arc would carry the base to the end of its travel in about 100 s, and he added a position loop whose polarity was chosen to be positive, so that a base displaced to one side moves slightly further out, the broom tips inward, and the broom falls back toward the center of the table [3]. Work at Stanford on systems of several independent pendula followed in 1963, and its analysis arrived at the same positive feedback around the cart position [3].

By the end of the 1960s the single pendulum had entered the standard textbooks. Two of them stabilized the angle alone and left the cart position unregulated, one with a gain condition that matches the floor $(M+m)g$ derived in Section 7.4.2, while a third set the cart loop as an exercise whose answer is that the cart corrects a position error by first backing up [3]. Both observations reappear below, the first as the runaway of case 4 and the second as the wrong-way start of every cart move.

Swinging the pendulum up from hanging and then stabilizing it was demonstrated in 1976 [12]. Swing-up laws were developed further in the decades that followed [13], and the energy formulation, which drives the energy of the pendulum toward its upright value with a pivot acceleration signed by $\dot{\theta}\cos\theta$, gave the method the form used in this application [14]. The cart and pole also became a test problem for learning control, in which adaptive elements learned to balance it from a failure signal alone [15].

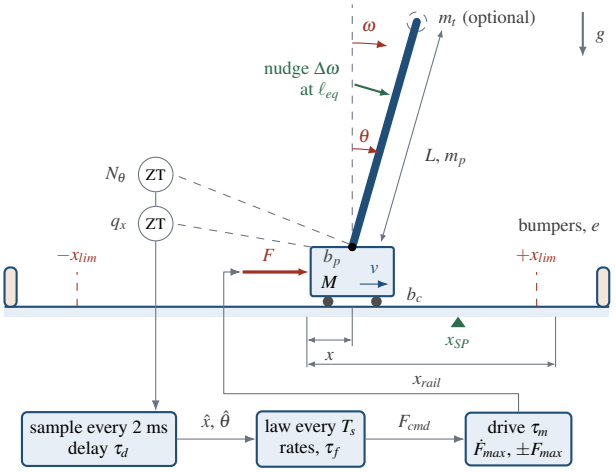


Figure 7.1 Cart and pole on a rail with bumpers, and the chain from the encoders to the drive. The cart of mass M on viscous friction b_c is pushed by the developed force F. The rod of length L and mass m_p, with an optional tip mass m_t, turns on a pivot with damping b_p, and its angle θ and rate ω are positive toward $+x$. A nudge changes ω by $\Delta\omega$ at the center of percussion, a distance ℓ_{eq} from the pivot. The cart center meets the bumpers at $\pm x_{rail}$, and they return it with restitution e, while the dashed lines at $\pm x_{lim}$ are the challenge limit. The encoders ZT, of resolution q_x and N_θ counts per turn, are sampled every 2 ms and delayed by τ_d, the law runs every T_s with rates filtered by τ_f, and the drive follows F_{cmd} with lag τ_m, slew limit $\dot{F}_{max}$ and limit $\pm F_{max}$.

> ⧖ HISTORY **Swinging Up by Energy, 1992 to 1996**
>
> Inverted pendulums had been classic tools in control laboratories since the 1950s, first to show how feedback stabilizes an unstable system [14]. A swing-up design published in 1992 noted that the optimal control found by the maximum principle is a function of time, applied as feedforward, so it is not robust and the transfer to upright is not assured [13]. An energy view, presented at a world congress in San Francisco in July 1996, treated the pendulum's energy as an integrator with varying gain, which loses control only when the pendulum is horizontal or reverses its swing. One ratio, the pivot's largest acceleration over gravity, sets the behavior: above four thirds one swing is enough, and below 0.577 it takes three swings or more. Energy control leaves the upright point a saddle, so another strategy must catch the pendulum [14].

The same structure moved into machines and into physiology. A two-wheeled mobile robot was designed and controlled as an inverted pendulum on a moving base [1], biped walking patterns are generated from a pendulum whose pivot is the support foot [2], and models of human balance use inverted pendula whose stabilization is limited by neural delay [4, 9]. Experiments on a cart and pole also found that a regulator designed with a viscous friction model left small oscillations that stick-slip friction explained [16], which bears on the friction model adopted below.

7.3 Cart and Pole Model

The apparatus in the application is a cart on a horizontal rail carrying a uniform rod on a free pivot, pushed by one force actuator, as shown in Figure 7.1. It is the upright cart of Section 1.4.1 with a force in place of a commanded acceleration, and the same cart with the pendulum hanging is the crane of Chapter 12. The cart position x is measured from the center of the rail and the pole angle θ from upright, both positive toward $+x$, so a pole leaning toward $+x$ has $\theta > 0$. The six states are the cart position and velocity v, the pole angle and rate ω, the drive force F as delivered so far, and a latch for a pole lying on the deck stop. Table 7.1 sorts the signals by role.

Table 7.1 Classification of the Signals

Symbol	Role	Nominal	Range
θ	CV, pole angle	0°	SP 0, band ±2°
x	CV, cart position	0 m	SP ±0.6, band ±0.05
F_{cmd}	MV, drive command	0 N	±20
x, v	states, cart	0	$\lvert x\rvert \le 1.0$ m
θ, ω	states, pole	0	full circle
F	state, developed force	0 N	±20
—	state, deck latch	0	0 or ±1
$\hat{x}$, $\hat{\theta}$	measurements	—	steps q_x, $2\pi/N_\theta$
$\hat{v}$, $\hat{\omega}$	rate estimates	0	filtered
$\Delta\omega$	DV, nudge	60 °/s	0 to 300
L, m_p	DV, rod	0.5 m, 0.15 kg	0.3–1.0, 0.05–1.0
$\theta(0)$	DV, initial tilt	0°	per scenario
θ_{lean}	output, lean asked	0°	±6°
E	output, pole energy	0 J	−0.736 J hanging
λ	output, unstable root	5.58 s^{-1}	—

7.3.1 Equations of Motion

Lagrange's equations for the cart and the pole, with the first moment $m\ell$ of the pole and its inertia J about the pivot, are given in Equation (7.1).

$$\begin{aligned}(M+m)\ddot{x} + m\ell\cos\theta\,\ddot{\theta} - m\ell\sin\theta\,\dot{\theta}^2 &= F - b_c\dot{x}\\ m\ell\cos\theta\,\ddot{x} + J\ddot{\theta} - m\ell g\sin\theta &= -b_p\dot{\theta}\end{aligned} \tag{7.1}$$

where $m = m_p + m_t$, $m\ell = m_pL/2 + m_tL$ and $J = m_pL^2/3 + m_tL^2$. The first row is a force balance in newtons and the second a moment balance about the pivot in newton meters. For the default rod $m\ell = 0.0375\,\mathrm{kg\,m}$ and $J = 0.0125\,\mathrm{kg\,m^2}$. Nothing is linearized. The two accelerations are solved from the 2×2 system at every evaluation of the right-hand side, and its determinant $\Delta(\theta) = (M+m)J - (m\ell\cos\theta)^2$ is never smaller than MJ, because $J \ge (m\ell)^2/m$ for any rigid pole, so the solve cannot fail anywhere on the circle that the swing-up visits. With no force and no friction the total energy $\frac{1}{2}(M+m)v^2 + m\ell v\omega\cos\theta + \frac{1}{2}J\omega^2 + m\ell g\cos\theta$ is conserved, and the verification suite of the application confirms it to one part in 10^6 over 10 s.

7.3.2 Drive, Rail and Deck Stop

The drive is a force source. The developed force follows the command through a first-order lag with a slew limit, as shown in Equation (7.2).

$$\dot{F} = \mathrm{sat}_{\dot{F}_{max}}\left(\frac{F_{cmd} - F}{\tau_m}\right), \qquad \lvert F_{cmd}\rvert \le F_{max} \tag{7.2}$$

where $\tau_m = 30\,\mathrm{ms}$ lumps the current loop, the motor and the belt, $\dot{F}_{max} = 1000\,\mathrm{N/s}$ and $F_{max} = 20\,\mathrm{N}$, which gives the bare cart $20\,\mathrm{m/s^2}$. The slew limit binds only when the command moves by more than $\dot{F}_{max}\tau_m = 30\,\mathrm{N}$, which within ±20 N requires a command swinging from near one limit toward the other, and the lag cannot carry F past a clipped command.

The cart meets rubber bumpers when its center reaches $x = \pm x_{rail} = \pm 1.0\,\mathrm{m}$. Contact is handled as a hard event between integration steps and never inside the derivative evaluation. A cart found beyond a stop is returned to it, and if it was moving into the stop its velocity is reversed with restitution $e = 0.4$, as shown in Equation (7.3).

$$\Delta v = -(1+e)\,v, \qquad \Delta\omega = -\frac{m\ell\cos\theta}{J}\,\Delta v \tag{7.3}$$

The second relation follows from applying the impulse to the cart alone through the mass matrix of Equation (7.1), so the pivot carries part of it into the pole, and a cart that strikes the right bumper throws its pole toward the right. For the default rod $m\ell/J = 3\,\mathrm{m}^{-1}$, so an arrival at 0.55 m/s turns the pole by 132 °/s. Below a resting-contact speed $v_{rest} = 0.02\,\mathrm{m/s}$ the restitution is zero, which also holds a cart that the drive keeps pressing into the stop, and with $e = 1$ an impact conserves kinetic energy exactly.

A pole that crosses $\pm 90^\circ$ outward lands on a stop on the cart deck, its rate is set to zero, and it then rides with the cart as one body of mass $M + m$. The stop is removed in the swing-up structure, whose pole must pass through the bottom. A nudge adds $\Delta\omega$ to the pole rate and leaves the cart velocity unchanged, which is what a tap at the center of percussion produces, a distance $\ell_{eq} = J/m\ell$ from the pivot ($2L/3 = 0.333\,\mathrm{m}$ for the bare rod). The pole energy relative to upright rest, Equation (7.4), is zero upright and $-2m\ell g = -0.736\,\mathrm{J}$ hanging.

$$E = \tfrac{1}{2} J\omega^2 + m\ell g\,(\cos\theta - 1) \tag{7.4}$$

7.3.3 What the Controller Receives

No control law reads the true state. Every $\Delta t_c = 2\,\mathrm{ms}$ the controller clock reads the two encoders, rounds the cart position to $q_x = 0.02\,\mathrm{mm}$ and the angle to $2\pi/N_\theta$ with $N_\theta = 4096$ counts per turn (0.088°), and pushes the pair into a delay line. When the law executes at t_k, once every sample interval T_s, it takes the sample from $t_k - \tau_d$, as shown in Equation (7.5), and forms the rates as filtered first differences, as shown in Equation (7.6).

$$\begin{aligned} \hat{x}_k &= q_x\,\mathrm{rnd}\left(\frac{x(t_k - \tau_d)}{q_x}\right), \\ \hat{\theta}_k &= \frac{2\pi}{N_\theta}\,\mathrm{rnd}\left(\frac{N_\theta\,\theta(t_k - \tau_d)}{2\pi}\right) \end{aligned} \tag{7.5}$$

$$\hat{v}_k = \hat{v}_{k-1} + a\left(\frac{\hat{x}_k - \hat{x}_{k-1}}{T_s} - \hat{v}_{k-1}\right) \tag{7.6}$$

where rnd rounds to the nearest integer, $a = 1 - \mathrm{e}^{-T_s/\tau_f}$, and $\hat{\omega}_k$ is formed from $\hat{\theta}$ in the same way. The filter time constant is $\tau_f = 15\,\mathrm{ms}$, so $a = 0.487$ at the default $T_s = 10\,\mathrm{ms}$, and one angle count differenced over one sample moves $\hat{\omega}$ by 0.075 rad/s. The delay is a multiple of the 2 ms clock between 0 and 120 ms, and the law holds its force between executions, so on average it sees the pole $\tau_d + T_s/2$ in the past. The encoders carry no random noise, and the application draws no random numbers, so the seed of Section 1.6.1 changes nothing in this app.

7.3.4 Parameters and Numerics

Table 7.2 lists every plant constant with the class the application gives it. Adjustable constants appear as sliders in the model dialog, the rod length and mass are disturbance sliders on the main panel, and hidden constants are fixed and shown only through their effect. The engine of Section 1.6.1 integrates the equations by the fourth-order Runge-Kutta method on a 0.5 ms step, the controller clock runs every 2 ms and the law every T_s, and the log samples at 50 Hz, with speeds from 0.1 to 2 seconds of simulated time per second of real time. The contact and deck corrections follow each integration step, before the outputs and the log sample. Halving the step to 0.25 ms moves the peak angle and the cart travel of case 1 by less than 0.5 %, which the verification suite checks.

7.3.5 Assumptions

The pole is a rigid uniform rod with viscous pivot damping. A real pole bends, and its first bending mode, far above the loop bandwidth for a short aluminum rod, approaches it for a long and light one.

Cart friction is viscous only. Stick-slip friction makes a balanced cart hunt in a small limit cycle about its target, which is what experiments on a cart and pole found when a regulator was designed

Table 7.2 Plant Parameters. Class a Is Adjustable in the Model Dialog, s Is a Disturbance Slider, h Is Hidden and Fixed.

Symbol	Meaning	Value	Range	
M	cart mass	1.0 kg	0.4–3	a
b_c	cart friction	1.0 N s/m	0–8	a
L	rod length	0.5 m	0.3–1.0	s
m_p	rod mass	0.15 kg	0.05–1.0	s
m_t	tip mass	0 kg	0–0.5	a
b_p	pivot damping	0.002 N m s/rad	0–0.05	a
g	gravity	9.81 m/s^2	1.6–24.8	a
F_{max}	drive force limit	20 N	5–40	a
τ_m	drive lag	30 ms	5–100	a
$\dot{F}_{max}$	drive slew limit	1000 N/s	100–5000	a
x_{rail}	bumper position	1.0 m	fixed	h
x_{lim}	challenge line	0.8 m	fixed	h
e	restitution	0.4	0–1	a
v_{rest}	resting-contact speed	0.02 m/s	fixed	h
q_x	cart encoder step	0.02 mm	0.001–1	a
N_θ	pole encoder counts	4096/turn	256–65536	a
$m\ell$, J	first moment, inertia	0.0375, 0.0125	derived	
ℓ_{eq}	$J/m\ell$	0.333 m	derived	
λ	unstable root	5.58 s^{-1}	derived	
z	wrong-way zero	5.35 rad/s	derived	

without it [16]. It would also alter the runaway of Section 7.4.2, because a lean too small to overcome breakaway would not move the cart.

The drive is a force source with one lag and one slew limit. Back electromotive force, current limits inside the amplifier, belt stretch and gear backlash are lumped into the lag, and a belt drive with backlash is harder to balance than this one.

The bumpers are instantaneous impacts with a fixed restitution, and the cart never lifts or rocks on its wheels. A real impact lasts milliseconds and can unseat the pole.

The encoders are exact apart from their resolution, with no noise and no missed counts, and the delay is a pure transport delay. A camera would add noise and dropped frames on top of the delay of case 8.

The regulator is designed with the true pole and redesigned at once after every change, which no engineer achieves in practice. Gravity, the masses and the frictions are constant during a run and the rail is level, so no constant force acts on the cart.

7.4 Linear Analysis of the Balancing Problem

Section 1.4 sets out the tools used here on a simple pendulum. Three models are kept apart below. The nonlinear equations of Section 7.3 are the simulation model, and the linearizations of this section serve analysis only. The controllers carry models of their own, the four-state linearization with friction on which the regulator is designed and the cart equation that the swing-up inverts, both evaluated with the true parameters rather than identified from data.

7.4.1 The Unstable Root and the Wrong-Way Zero

Linearized about upright rest with F as the input and no friction, Equation (7.1) becomes Equation (7.7).

$$\Delta_0\,\ddot{\theta} = (M+m)\,m\ell g\,\theta - m\ell F, \quad \Delta_0\,\ddot{x} = JF - (m\ell)^2 g\,\theta \tag{7.7}$$

where $\Delta_0 = (M+m)J - (m\ell)^2 = 0.01297\,\mathrm{kg^2\,m^2}$. The tilt grows at $\lambda_0 = \sqrt{(M+m)m\ell g/\Delta_0} = 5.71\,\mathrm{s^{-1}}$, and with the cart friction and the pivot damping the root is $5.58\,\mathrm{s^{-1}}$, one e-fold every 179 ms and a doubling every 124 ms. The remaining open-loop roots are 0 for the cart position, $-0.87\,\mathrm{s^{-1}}$ and $-5.86\,\mathrm{s^{-1}}$. Holding a tilted pole still requires accelerating the whole assembly at g per radian of tilt, so the force needed per radian is $(M+m)g = 11.3\,\mathrm{N/rad}$, a figure that recurs below as a gain floor.

The transfer function from force to cart position has the numerator of the pole equation, as shown in Equation (7.8).

$$\frac{x(s)}{F(s)} = \frac{Js^2 + b_p s - m\ell g}{D(s)}, \qquad z \approx \sqrt{g/\ell_{eq}} \tag{7.8}$$

where $D(s)$ is the open-loop characteristic polynomial. Its positive root z is 5.42 rad/s without pivot damping and 5.35 rad/s with it. It is the rate at which the pole would fall if the cart were clamped, since ℓ_{eq} is the length of the equivalent simple pendulum, and it lies only 4 % below the unstable root. The sensitivity function of any stable loop must equal one at a right-half-plane zero of the plant and vanish at a right-half-plane pole, so a maximum modulus argument bounds its peak below by $(z+\lambda)/(\lambda - z) = 46$ for a controller that measures only the cart position [7]. This is the bound of Section 1.4.7, which is infinite on the running example because a drive that sets the acceleration puts the zero on the pole. The force drive here separates them by only 4 %, so the bound is finite but far too large for a usable loop. Every structure in the application therefore reads the angle, and the cart position enters through it.

7.4.2 The Angle Loop Alone

With $F = K_{p\theta}\theta + K_{d\theta}\omega$ and the cart left out of the law, the linearized loop has the characteristic polynomial of Equation (7.9) when the small pivot damping is neglected.

$$s\big[\Delta_0 s^3 + (b_c J + m\ell K_{d\theta})s^2 + m\ell\,(K_{p\theta} - (M+m)g)\,s - b_c m\ell g\big] = 0 \tag{7.9}$$

Without friction the bracket reduces to $s[\Delta_0 s^2 + m\ell K_{d\theta} s + m\ell(K_{p\theta} - (M+m)g)]$, so the pole is held whenever $K_{p\theta} > (M+m)g = 11.3\,\mathrm{N/rad}$ and $K_{d\theta} > 0$, a floor that depends on the total mass and not on the length or the inertia of the pole. Chapter 6 derives the same kind of floor for a levitated ball. The cart is left with a double root at the origin and drifts. With friction the constant term is negative, so by Descartes' rule of signs one root is real and positive whatever the gains. For small friction it is close to $b_c g/(K_{p\theta} - (M+m)g)$, because a cart moving at speed v needs the force $b_c v$, which the loop can supply only by letting the pole lean forward by $b_c v/(K_{p\theta} - (M+m)g)$, and that lean accelerates the cart at g per radian. The estimate is $0.52\,\mathrm{s^{-1}}$ at the defaults, the exact root is $0.46\,\mathrm{s^{-1}}$, and it is $0.47\,\mathrm{s^{-1}}$ with the drive lag, the rate filter and the hold included. Friction, which damps most mechanical systems, destabilizes this one.

Pause. Try Exercise 7.1 before reading on.

7.4.3 The Cascade as Four-State Feedback

The cascade of Section 7.5 multiplies out, inside its lean clamp, to $F = K_{p\theta}\theta + K_{d\theta}\omega + k_x(x - x_{SP}) + k_v v$ with $k_x = K_{p\theta}K_{px} = 3.67\,\mathrm{N/m}$ and $k_v = K_{p\theta}K_{dx} = 5.24\,\mathrm{N\,s/m}$ at the defaults. With the frictionless model, that law closes into the quartic of Equation (7.10).

$$\frac{\Delta_0}{m\ell}s^4 + (K_{d\theta} - \ell_{eq}k_v)\,s^3 + (K_{p\theta} - (M+m)g - \ell_{eq}k_x)\,s^2 + gk_v s + gk_x = 0 \tag{7.10}$$

At the defaults its coefficients are 0.346, 3.25, 17.5, 51.4 and 36.0 in SI units, and its roots are $-0.97\,\mathrm{s^{-1}}$, $-2.04 \pm 4.52i\,\mathrm{s^{-1}}$ and $-4.37\,\mathrm{s^{-1}}$. Three consequences follow. The two lowest coefficients

Table 7.3 Dominant Modes of the Sampled-Data Linearization Against the Measured Behavior

Case, configuration	Mode (s^{-1})	Measured
1, cascade	-1.56	settles
1, regulator	$-1.96 \pm 1.61i$	settles
3, 75 N/rad	$+2.52 \pm 12.06i$	$\pm 16.7°$ at 2.10 Hz
3, 55 N/rad	$+0.55 \pm 10.80i$	$\pm 9.3°$ at 1.75 Hz
3, 40 N/rad	$-1.38 \pm 0.75i$	$\pm 0.08°$ residual
3, $\tau_m = 10$ ms	$+1.02 \pm 12.95i$	$\pm 8.8°$
4, angle loop alone	$+0.471$	speed grows at 0.47 s^{-1}
4, weak cart loop	$+0.024 \pm 0.531i$	growing swing, 11.9 s period
5, cascade	$+0.50 \pm 4.61i$	$\pm 8.3°$ at 0.70 Hz
5, 40 and 7 N s/rad	$-0.061 \pm 5.76i$	$\pm 1.00°$ from 30 to 60 s
5, soft cart loop	$-0.96 \pm 4.49i$	settles
6, cascade	$+0.022 \pm 3.92i$	$\pm 9.8°$ at 0.60 Hz
6, 40 N/rad	$-0.42 \pm 6.34i$	settles
8, cascade	$+1.05 \pm 8.04i$	falls at 6.84 s
8, 20 N/rad	$-0.26 \pm 2.92i$	holds
8, 15 N/rad	$+0.47 \pm 2.30i$	falls at 4.01 s

are positive only if k_x and k_v are, so the loop must push a cart that is right of its target further to the right, the polarity chosen deliberately in 1960 and the signature of the zero at z. The outer loop consumes the angle loop's damping in proportion to $\ell_{eq} k_v$ and its stiffness in proportion to $\ell_{eq} k_x$, so the Routh-Hurwitz conditions put the floor of the cascade at 17.4 N/rad rather than 11.3 N/rad. Because k_v grows with $K_{p\theta}$, the same conditions give the cascade a ceiling even without lags, at 73.2 N/rad. A changed pole also attacks a particular coefficient. A 1.0 m rod doubles ℓ_{eq} to 0.667 m and cuts the s^3 coefficient from 3.25 to 1.51, while a 1.0 kg rod raises $(M+m)g$ to 19.6 N/rad and cuts the s^2 coefficient from 17.5 to 9.16. More angle gain restores the second but not the first, because k_v grows with $K_{p\theta}$ in a cascade, while a softer cart loop restores the first.

7.4.4 The Sampled-Data Loop

Equations (7.9) and (7.10) leave out the drive lag, the rate filter, the hold and the delay. The linearized plant under the implemented law, without the quantizers, the clamps and the limits, is a linear periodic system, so its small-signal modes can be computed exactly. The plant with friction and the drive lag is discretized with a zero-order hold over one 2 ms clock period, the delay line and the filtered differences of Equation (7.6) are written as states, and the product over one sample interval is a transition matrix whose eigenvalues μ, mapped by $s = \ln(\mu)/T_s$, are the closed-loop modes. With the quantizers removed from the simulator and a 0.001 °/s tap that keeps the response linear, the model reproduces the growth and the frequency of the rocking mode at 55 N/rad (0.547 against 0.545 s^{-1}, at 1.718 Hz in both), of the heavy rod of case 6 (0.022 s^{-1} at 0.624 Hz in both) and of the stiffened loop on the long rod of case 5 (–0.061 s^{-1} at 0.916 Hz in both). Table 7.3 sets its predictions beside the measured outcomes.

At the default 10 ms sample interval and no delay the sampled cascade is stable from 18.6 to 50.7 N/rad, so the lags cut the lag-free ceiling of 73.2 N/rad by almost a third, while the lags and the friction raise the floor from 17.4 to 18.6 N/rad. At the 20 ms sample interval of case 8, the floor stays between 18.3 and 18.6 N/rad for every delay, 65 % above $(M+m)g$, while the ceiling falls from 47.8 N/rad with no delay to 38.6, 31.7, 26.3 and 21.6 N/rad at delays of 20, 40, 60 and 80 ms, as shown in Figure 7.2. The window closes between 94 and 96 ms, an average sensing lag of about 105 ms, or 0.59 of the e-folding time. Proportional feedback on the unstable mode alone keeps a window up to the bound $p\,\tau_d < 1$ of Section 1.4.7, with p the unstable root, which is a delay of one full e-folding time. The cascade, with its lags and its cart loop, loses the window well before that. Forty nonlinear runs with the 60 °/s tap agree with the window at 39 points. The exception is 20 N/rad

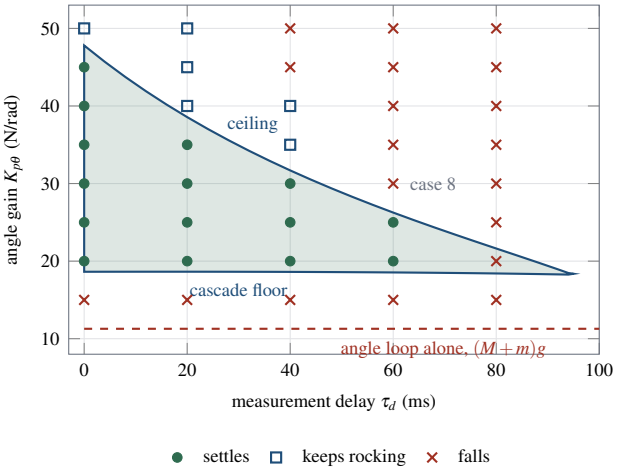


Figure 7.2 Stable window of the angle gain against the measurement delay for the default cascade at the 20 ms sample interval of case 8, from the sampled-data linearization. Markers are nonlinear runs with the 60 °/s tap, classified from 35 to 40 s.

at 80 ms, which is linearly stable with a pair at $-0.33 \pm 8.26i\,\mathrm{s}^{-1}$ and holds a 5 °/s tap, but which the full tap drives onto the lean clamp for 1.75 s and into the bumper at 4.77 s. A linear window is a necessary condition for recovering a disturbance and not a sufficient one.

7.5 Implemented Control Structures

7.5.1 The Sampled Law and the Drive Command

The law executes once every sample interval T_s, 10 ms by default and adjustable from 2 to 40 ms in steps of 2 ms, on the measurements of Equations (7.5) and (7.6), and it holds its force between executions. Every law computes a force toward $+x$. The command is clipped at $\pm F_{max}$, and the saturation flag compares the unclipped command with that limit and is cleared on every scan, so a law asking for 35 N reads as saturated while the developed force is still rising. A watchdog forces the command to zero once the measured angle reaches 90°, except in the swing-up. None of the five laws carries integral action, so the velocity-form PID of Section 1.6.2 is not used. There is no integrator to wind up, and the only limits are the command clamp, the drive slew limit of Equation (7.2) and, in the cascade, the lean clamp. The settings are collected in Table 7.4.

7.5.2 Angle Loop and Cascade

The angle loop pushes the cart toward the side to which the pole leans, as shown in Equation (7.11), with the defaults $K_{p\theta} = 30\,\mathrm{N/rad}$ and $K_{d\theta} = 5\,\mathrm{N\,s/rad}$.

$$F = K_{p\theta}\,\hat{\theta} + K_{d\theta}\,\hat{\omega} \tag{7.11}$$

The cascade wraps a slow cart loop around it that asks the pole for a lean toward the target, as shown in Equation (7.12).

$$\begin{aligned} \theta_{lean} &= \mathrm{sat}_{\theta_{lean,max}}\big(K_{px}(x_{SP} - \hat{x}) - K_{dx}\hat{v}\big) \\ F &= K_{p\theta}(\hat{\theta} - \theta_{lean}) + K_{d\theta}\,\hat{\omega} \end{aligned} \tag{7.12}$$

with $K_{px} = 7^\circ$ per meter, $K_{dx} = 10^\circ$ per m/s and $\theta_{lean,max} = 6^\circ$. The angle loop's only way to lean the pole toward the target is to pull the cart away from it first. This is the right-half-plane zero of Section 1.4.7, which asks for a fast loop on the angle inside a slow loop on the cart. The rocket of Volume 3, Chapter 10 has the same structure, a position loop that commands an attitude, and the

Table 7.4 Controller Settings

Symbol	Meaning	Default	Range
T_s	sample interval	10 ms	2–40
τ_d	measurement delay	0 ms	0–120
τ_f	rate filter	15 ms	0–50
$K_{p\theta}$	angle gain	30 N/rad	0–150
$K_{d\theta}$	angle damping	5 N s/rad	0–15
K_{px}	cart gain	7° per m	0–30
K_{dx}	cart damping	10° per m/s	0–30
$\theta_{lean,max}$	lean clamp	6°	1–20
x_s	regulator cart scale	0.5 m	0.05–2
θ_s	regulator angle scale	20°	1–45
F_s	regulator force scale	5 N	1–40
T_p	MPC prediction horizon	1 s	0.3–2
x_{max}	MPC rail limit	0.8 m	0.3–0.95
θ_{max}	MPC lean limit	15°	5–45
P	MPC terminal weight	LQR	LQR, Off
k_E	energy gain	100 m/s^2 per J	0–200
a_{sw}	pump acceleration limit	4 m/s^2	0.5–10
θ_{cap}	capture angle	20°	5–45
ω_{cap}	capture rate	180 °/s	30–600
$\Delta\theta_{rel}$	release margin	15°	5–45
ω_c, ζ_c	pump centering	1.2 rad/s, 0.9	fixed
F_{push}	manual push	2 N	0.5–20

same kind of zero. Written as $F = -K\tilde{x}$ with $\tilde{x} = (\hat{x} - x_{SP}, \hat{v}, \hat{\theta}, \hat{\omega})$, the cascade inside its clamp is $K = [-3.67, -5.24, -30, -5]$ in N/m, N s/m, N/rad and N s/rad, all four entries negative.

7.5.3 Discrete Linear Quadratic Regulator

The regulator applies $F = -K\tilde{x}$ with K from the discrete algebraic Riccati equation, the sampled counterpart of Equation (1.6), on the model linearized at upright. The design model is the four-state linearization with friction and without the drive lag, discretized with a zero-order hold at T_s through one matrix exponential of the augmented matrix $\left[\begin{smallmatrix} A & B \\ 0 & 0 \end{smallmatrix}\right] T_s$, evaluated by scaling and squaring with a 14-term Taylor series. The weights follow Bryson's rule [17] of Section 1.4.5, one over the square of the acceptable excursion, as shown in Equation (7.13).

$$Q = \operatorname{diag}\left(\frac{1}{x_s^2}, 0, \frac{1}{\theta_s^2}, 0\right), \qquad R = \frac{1}{F_s^2} \tag{7.13}$$

with $x_s = 0.5\,\text{m}$, $\theta_s = 20°$ and $F_s = 5\,\text{N}$. The recursion $P \leftarrow Q + \Phi^T P \Phi - \Phi^T P \Gamma (R + \Gamma^T P \Gamma)^{-1} \Gamma^T P \Phi$ runs from $P = Q$ until the relative change is below 10^{-11}, which takes 697 sweeps at the defaults, and gives $K = [-9.26, -9.44, -48.60, -8.28]$. The single input makes the inverse a scalar division. The gain is recomputed whenever the rod length or mass, the tip mass, the cart mass, either friction, gravity, the sample interval or a weight changes, as if the engineer redesigned after every change with a perfect model of the new pole. A change of any of these constants is therefore a redesign for a different rig, and not a test of the regulator's robustness, and the same holds for the predictive controller and the swing-up catch, which receive this design, and for the pump force of Equation (7.15), which uses the same constants. The regulator's design poles are $-1.75 \pm 1.51i\,\text{s}^{-1}$ and $-6.53 \pm 3.21i\,\text{s}^{-1}$, and with the drive lag, the rate filter and the hold the dominant pair moves to $-1.96 \pm 1.61i\,\text{s}^{-1}$. The angle entry becomes −61.1 N/rad for the 1.0 m rod and −63.8 N/rad for the 1.0 kg rod.

7.5.4 Energy Swing-Up and Handover

The swing-up pumps the energy of Equation (7.4), evaluated on the measurements, toward zero. Following the energy-control law [14], the pump asks for a cart acceleration, as shown in Equation (7.14).

$$a = \mathrm{sat}_{a_{sw}}\big(k_E E\,\mathrm{sgn}(\hat{\omega}\cos\hat{\theta})\big) - \omega_c^2(\hat{x} - x_{SP}) - 2\zeta_c\omega_c\hat{v} \tag{7.14}$$

with $k_E = 100\,\mathrm{m/s^2}$ per J, $a_{sw} = 4\,\mathrm{m/s^2}$ and a centering spring of $\omega_c = 1.2\,\mathrm{rad/s}$ and $\zeta_c = 0.9$ that keeps the pumping cart off the bumpers. At rest hanging, where $\hat{\omega}\cos\hat{\theta}$ is zero, the sign is taken as positive so that the swing starts itself. The acceleration becomes a force through the cart equation solved at the measured state, as shown in Equation (7.15).

$$F = \frac{\Delta(\hat{\theta})\,a + m\ell\cos\hat{\theta}\,(m\ell g\sin\hat{\theta} - b_p\hat{\omega})}{J} + b_c\hat{v} - m\ell\hat{\omega}^2\sin\hat{\theta} \tag{7.15}$$

This cancels the reaction of the swinging pole on the cart and the cart friction that would otherwise absorb the pumped energy. The regulator takes over when $|\hat{\theta}| < \theta_{cap} = 20^\circ$ and $|\hat{\omega}| < \omega_{cap} = 180^\circ/\mathrm{s}$, and the pump resumes only past $\theta_{cap} + \Delta\theta_{rel} = 35^\circ$, so the handover cannot chatter.

7.5.5 Model Predictive Control on the Regulator's Model

The fifth structure is a model predictive controller (MPC), an option that no case before case 10 selects. At each decision it plans the force over a horizon of $T_p = 1\,\mathrm{s}$, applies the first move and plans again from the next measurement, the receding-horizon law of Section 1.6.3, whose constrained form, with its stability and optimality conditions, is reviewed in [18]. Its model is that of the linear quadratic regulator (LQR) of Section 7.5.3. The matrices Φ, Γ, K and P and the weights Q and R of Equation (7.13) are handed to it from the same design call, so a changed pole reaches both controllers at once and both are given the present plant. Nothing is identified. This departs from the general rule of Section 1.6.3, under which a predictive controller never reads the parameters of the run, and it makes a changed pole a redesign rather than a model error. The MPC reads what the regulator reads and nothing else, the state $\tilde{x}$ formed from the delayed, quantized readings and filtered rates of Equations (7.5) and (7.6), the cart target and the drive limit F_{max}. The plant with its drive lag, the instruments with their quantizers, delay line and filters, and the controller's model, the four-state linearization held at T_s without the drive lag, therefore remain three separate things.

The decisions are offsets from the regulator's move, $F_k = -K\tilde{x}_k + c_k$, so the free variables perturb a fixed state feedback, a structure used in [19]. With the force carried as a fifth state, $\xi_k = (\tilde{x}_k, F_{k-1})$, the plan predicts the loop that the regulator closes, as shown in Equation (7.16). Chapter 8 plans departures from its regulator's force in the same way.

$$\xi_{k+1} = \begin{bmatrix} \Phi - \Gamma K & 0 \\ -K & 0 \end{bmatrix} \xi_k + \begin{bmatrix} \Gamma \\ 1 \end{bmatrix} c_k \tag{7.16}$$

where k counts intervals of T_s from the present state $\tilde{x}_0$ and $N = T_p/T_s$ steps, 100 at the defaults, span the horizon. The prediction thus starts on the stable closed loop rather than on the root at $+5.58\,\mathrm{s^{-1}}$, and the predicted force is an output that a limit can act on. The objective is the regulator's cost over the horizon with P as the terminal weight, plus the prices of the limits described below, as shown in Equation (7.17).

$$J = \sum_{k=0}^{N-1}\left(\tilde{x}_k^T Q\tilde{x}_k + RF_k^2\right) + \tilde{x}_N^T P\tilde{x}_N + \sum_i\left(\rho_i\varepsilon_i^2 + \sigma_i\varepsilon_i\right) \tag{7.17}$$

where $\varepsilon_i \geq 0$ are the slacks of the three limits and ρ_i and σ_i their quadratic and linear prices. The state term at $k = 0$ is fixed by the measurement and is omitted in the implementation. With $W = R + \Gamma^T P \Gamma$ and the gain $K = W^{-1}\Gamma^T P \Phi$, every state and force satisfies Equation (7.18).

$$\tilde{x}^T Q \tilde{x} + RF^2 + (\Phi\tilde{x} + \Gamma F)^T P (\Phi\tilde{x} + \Gamma F) - \tilde{x}^T P \tilde{x} = W(F + K\tilde{x})^2 \quad (7.18)$$

Expanded, the left side is $\tilde{x}^T(Q + \Phi^T P \Phi - P)\tilde{x} + 2F\Gamma^T P \Phi \tilde{x} + WF^2$. The Riccati equation turns its first term into $W(K\tilde{x})^2$ and the gain turns its second into $2WFK\tilde{x}$. Summed along any predicted trajectory, the P terms telescope, and the cost without slacks becomes Equation (7.19) for every horizon and every sequence of offsets.

$$J_0 = \tilde{x}_0^T P \tilde{x}_0 + W \sum_{k=0}^{N-1} c_k^2 \quad (7.19)$$

Three consequences follow. With no limit binding the best offsets are zero for any horizon and any blocking, so the first move is the regulator's, and the regulator is this MPC with its limits removed. The Hessian in the offsets is diagonal, $2Wn_j$ for an offset held over n_j steps. With only the hard limit on the first move, the plan is the regulator's move clipped at $\pm F_{max}$ with every later offset zero. Within its limits the plan is a finite-horizon, move-blocked form of the constrained linear quadratic regulator [20]. At the defaults $W = 0.04667$ against $R = 0.04$. Equation (7.19) holds to a relative error of 2.3×10^{-10} on 200 random states and offset sequences, the first move equals $-K\tilde{x}$ to 1.7×10^{-9} N over the 1599 plans of 16 s of case 1, and the clipped move is reproduced to 4.7×10^{-9} N.

Pause. Try Exercise 7.3 before reading on.

A plan is made at the regulator's own sample instants but at most every 10 ms, once in $n_d = \lceil 10\,\text{ms}/T_s \rceil$ executions, and between plans each execution applies $-K\tilde{x}$ plus the plan's offset for that instant. The offsets are held over six blocks of n_d times 1, 2, 4, 8 and 16 steps and the remainder of the horizon, which are 1, 2, 4, 8, 16 and 69 steps at the defaults. The horizon is 1.75 time constants of the slowest design pole, $-1.75 \pm 1.51i\,\text{s}^{-1}$, and P prices what lies beyond it. Horizons from 0.5 to 2 s stop the cart of case 10 within 1 mm of 0.768 m, and 0.3 s lets it run to 0.819 m.

Two hard rows hold the first move inside the drive limit, $|-K\tilde{x}_0 + c_0| \leq F_{max}$, with F_{max} read as the drive clamp reads it. Three limits are soft, each with one slack over the whole horizon, and they are judged every 50 ms of the plan, as shown in Equation (7.20).

$$|F_k| \leq F_{max} + \varepsilon_F, \qquad |x_k| \leq x_{max} + \varepsilon_x, \qquad |\theta_k| \leq \theta_{max} + \varepsilon_\theta \quad (7.20)$$

where $x_k = x_{SP} + \tilde{x}_{1,k}$ and θ_k are the predicted cart position from the center of the rail and the predicted angle. The rail limit $x_{max} = 0.8$ m is the challenge line, 0.2 m inside the bumpers, and the lean limit $\theta_{max} = 15°$ keeps the plan where the linearization holds, since up to 15° the sine of the angle stays within 1.1 % of the angle. The force is priced at $\rho_F = 10^4$ per N^2 and $\sigma_F = 10^4$ per newton, so its limit binds wherever it can. The rail and the lean are priced at $\rho = 10^3$ per m^2 or rad^2 and $\sigma = 10^2$ per meter or radian, which is 250 and 122 times the regulator's own weights on the cart and the angle. These prices and θ_{max} were chosen once, on case 10, from ρ in decades from 10 to 10^5 with $\sigma = \rho/10$ and θ_{max} of 15, 20, 30 and 45°, as the smallest cart excursion among the 14 of 20 settings that keep the pole up and the cart off the bumpers, while $\rho \leq 10^2$ behaves as the regulator. Each plan is a quadratic program in nine variables, the six offsets and three slacks, with 125 inequality rows at the defaults, solved by the dual active-set method of Section 1.6.3 with a ridge of $10^{-9}(1 + 2RN)$ on the diagonal of the offsets.

Offset-free predictive controllers usually add integrating disturbances to their model [21], as the Kalman filter of Section 1.6.3 does and as the boiler of Volume 9, Chapter 4 does for an unmeasured

fuel bias. This one has none, because no law of the application integrates, no constant force acts on the cart, upright rest on the target with zero force is an exact equilibrium of both the plant and the model, and an estimator would give the MPC dynamics that the regulator lacks and void Equation (7.19). The offset that results is shared. A constant force $d = 1\,\mathrm{N}$, added to the cart outside the application, leaves the cart 107.9 mm from the target under the regulator and 107.7 mm under the MPC, against $d/|K_1| = 108.0\,\mathrm{mm}$ from the cart entry of the gain.

A plan is discarded if the quadratic program ends in any status other than optimal, if the force it would apply at once is undefined or infinite, or if that force exceeds F_{max} in magnitude. The regulator then keeps the drive with the remaining offsets of the last accepted plan, alone once that plan has run out, a mark is written to the record, and control passes back to the MPC at the next accepted plan. Forced at each of the 101 decisions of case 10 between 2.00 and 3.00 s, a rejection leaves the command 0.21 N from the MPC's own on average and 2.97 N at worst, against 0.94 and 6.91 N for the bare regulator. No shipped case rejects a plan. The MPC also plans in manual operation. Handed the captured pole by the swing-up in the same capture region, it catches the pole of case 9 at 2.67 s, as the regulator does. Because a plan never asks past $\pm F_{max}$, the saturation flag counts the time the MPC's command sits at the limit.

7.5.6 Terminal Cost and the Stability of a Short Plan

The terminal weight is a setting of the MPC with two values. LQR, the default and the controller described above, prices the end of the plan with the regulator's P. Off gives the last step the stage weight instead, so that without slacks the plan minimizes Equation (7.21).

$$J_{\mathrm{off}} = \sum_{k=0}^{N-1} \left(\tilde{x}_k^T Q \tilde{x}_k + R F_k^2 \right) + \tilde{x}_N^T Q \tilde{x}_N \tag{7.21}$$

Nothing else changes. The model of Equation (7.16), the offsets, the blocks, the three soft limits, the hard rows on the first move and the rejection rule are the same objects under both values. The telescoping sum behind Equation (7.19) needs P at the last step, so under Off the zero offset is no longer optimal. With no limit binding the first planned force is still linear in the state, $F_0 = -K_H \tilde{x}_0$, with a gain K_H that depends on the horizon and the blocking, and the loop that the plan's own model predicts under that law is Equation (7.22).

$$\tilde{x}_{k+1} = (\Phi - \Gamma K_H)\, \tilde{x}_k \tag{7.22}$$

The receding-horizon law is stable on its model when the spectral radius $\rho(\Phi - \Gamma K_H)$, the largest eigenvalue magnitude, is below one.

With P at the last step $K_H = K$ at every horizon. The terminal weight is the value of the infinite-horizon problem from $\tilde{x}_N$, so the plan is the regulator's infinite-horizon problem cut at N with nothing lost, which is how a terminal cost equal to the cost-to-go makes a receding-horizon law stabilizing for stable and unstable linear plants alike [18, 22]. A condensed solution of the plan written independently of the application, with explicit predictions over the graduated blocks, gives $|K_H - K| \le 6.8 \times 10^{-9}$ for every N from 20 to 200 at $T_s = 10\,\mathrm{ms}$, and the application's own law agrees with that solution under Off to 2.3×10^{-6} in every gain entry. Without P the cart entry of the gain changes sign. At $T_p = 0.3\,\mathrm{s}$, $K_H = [+1.03, -1.45, -28.66, -4.36]$ against the regulator's $[-9.26, -9.44, -48.60, -8.28]$, so the plan holds the pole and pushes a displaced cart farther out, and $\rho = 1.00606$ per 10 ms, a real cart mode growing at $0.60\,\mathrm{s}^{-1}$.

The radius falls as the horizon grows, as shown in Figure 7.3. It is 1.00350 at 0.4 s, 1.000008 at 0.45 s and 0.99832 at 0.46 s, where the law first becomes stable, and every longer horizon to 2 s stays stable, with 0.99704 at 0.5 s and 0.98673 at 1 s against the regulator's 0.98261. With the drive lag and the rate filter added to the loop, as in the sampled-data model of Section 7.4.4, the radius moves by at most 2.0×10^{-3} and the crossing stays at 0.46 s, about 2.6 e-folding times of the unstable root and 0.8 time constants of the slowest design pole. A long enough horizon is the other route to

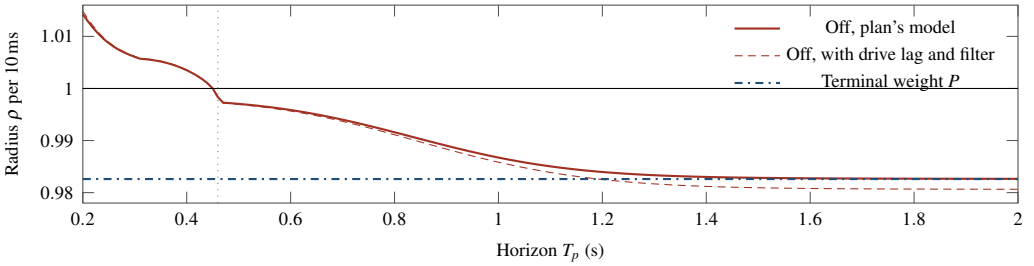


Figure 7.3 Spectral radius of the predictive law with no limit binding against the horizon, at $T_s = 10$ ms with the weights of case 10. Without the terminal weight (thick solid on the plan's model, dashed with the drive lag and the rate filter added) the law is unstable below the dotted line at 0.46 s. With it (dash-dotted) the law is the regulator's at every horizon. The thin solid line is a radius of one.

stability without a terminal cost [23], and 0.46 s is that length for this law. The pure finite-horizon law, with one-step blocks instead of the graduated ones, is a different gain, that of the backward Riccati recursion from $P_N = Q$. Its radius is 1.0105 at 0.3 s, and it crosses one at 0.48 s.

The linear crossing is necessary for a hold and not sufficient. Just past it the slowest mode is the cart, decaying at 0.296 s^{-1} at 0.5 s, a time constant of 3.4 s, so a push that sends the cart toward the end of the rail reaches the bumper before that mode brings the cart back. The drive limit and the 30 ms lag, which the plan's model lacks, take further margin, and the nonlinear rig needs a longer horizon than the linear crossing, as Section 7.6.9 measures.

7.5.7 Manual Operation

In manual operation the command is the drive-force slider, ±20 N in 0.5 N steps, plus pushes of $F_{push} = 2$ N that last 50 ms of simulated time from the buttons, or as long as an arrow key is held, so a push is the same at every speed. In automatic operation the slider follows the command, so a switch to manual hands over the force being applied. The law keeps executing in manual operation. With no integrator there is nothing to initialize, and on the switch to automatic the law's latest force replaces the operator's, softened only by the drive lag. The transfer is bumpless at the upright rest state that the verification suite checks, and not in general.

7.6 Results Across the Scenario Sequence

Eleven cases and a challenge make up the sequence in Table 7.5. Each of the first nine departs from case 1 in one setting, case 10 changes the structure and the push, and case 11 runs the controller of case 10 on the push of case 7 with a 0.3 s horizon and no terminal weight. The values in Table 7.6 were measured by running the published engine headlessly on each scenario, with the one change named made after the scenario is loaded and before the run starts, as a learner makes it, and reading the scorecard and the trajectories at the 0.5 ms physics step. The engine draws no random numbers, so every value repeats exactly on one JavaScript engine. The scenarios, their completion rules and the scorecard are those of Section 1.6.4.

7.6.1 Balance, the Wrong-Way Start and a Human

Case 1 runs the cascade at its defaults. The 60 °/s tap at 2 s leans the pole 4.68° toward $+x$ at 2.12 s, and the cart loop, braking the cart as it runs out, holds its lean request on the −6° clamp for 0.29 s. The pole overshoots that request to −6.98° at 2.66 s, the cart reaches 0.292 m on a 5.77 N peak, the pole is back inside ±2° 1.06 s after the tap and the cart is within 2 cm of center 2.46 s after it. At 8 s the target steps to 0.5 m. The lean request jumps to 3.5° and rises to 4.86° while the cart backs 21.3 mm the wrong way in the first 0.28 s, and the cart then reaches 90 % of the move 1.90 s after the step with a 1.0 mm overshoot (Table 7.6), and the cart trend of Figure 7.4 shows the move. The

Table 7.5 Scenario Sequence

#	Scenario	Change from case 1	Lesson
1	Tuned balance	—	catch, wrong-way start
2	Balance it yourself	manual, 1° start	the growth rate
3	Too much angle gain	$K_{p\theta} = 75\,\mathrm{N/rad}$	the ceiling
4	The angle loop alone	no cart loop, 2°	friction runaway
5	A longer pole	$L = 1.0\,\mathrm{m}$	damping lost
6	A heavier pole	$m_p = 1.0\,\mathrm{kg}$	stiffness lost
7	A hard push	tap 150 °/s	recoverable push
8	Delayed sensing	T_s 20, τ_d 60 ms	delay closes the window
9	Swing-up from hanging	hanging start	capture, handover
10	LQR against MPC on the rail	MPC, tap 205 °/s	worth of the limits
11	Terminal cost at a short horizon	MPC 0.3 s, Off, tap 150 °/s	worth of P
C	Challenge	5°, 125 °/s at 3 s	a limit off the CV

Table 7.6 Outcomes Read From the Running Engine

Case and condition	Result
1, tap 60 °/s at 2 s	4.68°, then −6.98°, 0.292 m, 5.77 N
1, back in ±2° / cart in 2 cm	1.06 s / 2.46 s after the tap
1, target 0 to 0.5 m at 8 s	21.3 mm back, 90 % at 1.90 s, 1.0 mm over
1, regulator	0.188 m, 0.83 s, 35.6 mm back, 1.42 s
2, no force, L 0.3 / 0.5 / 1.0 m	down at 0.741 / 0.944 / 1.326 s
2, automatic from 1°	peak 1.005°, 0.022 m
2, switch to automatic at 0.5 s	peak 10.8°, 0.464 m
3, 75 N/rad	±16.7° at 2.10 Hz, 43.3 % saturated
3, 55 / 40 N/rad	±9.3° / ±0.08°
3, $K_{d\theta}$ 8 / τ_m 10 ms	±0.08° / ±8.8°
4, angle loop alone	bumper at 4.11 s, down at 4.54 s
4, cascade / weak cart loop	0.046 m / bumper at 61.8 s
5, cascade	±8.3° at 0.70 Hz, first swing 0.970 m
5, 40 and 7 / soft cart loop	±1.00° / bumper at 3.28 s
5, regulator	1.01 s, 0.461 m
6, cascade	±9.8° at 0.60 Hz, cart ±0.178 m
6, 40 N/rad / regulator	3.28 s, 0.295 m / 0.87 s, 0.194 m
7, cascade	bumper at 2.88 s, down at 3.29 s
7, 15° clamp / regulator	0.745 m / 0.495 m, 19.0 N
7, largest tap without a bumper	130 / 183 / 229 °/s
8, cascade	bumper at 6.53 s, down at 6.84 s
8, $\tau_d = 20\,\mathrm{ms}$	peak 7.6°, 0.321 m
8, 40 / 20 / 15 N/rad	down 3.11 s / 14.9° / down 4.01 s
8, regulator / MPC	down 2.98 s / down 3.00 s
9, a_{sw} 4 / 6 / 2 m/s^2	2.67 s / 2.06 s / never caught
10, regulator / MPC	0.809 m, line at 2.46 s / 0.768 m
11, Off / terminal P	down 3.04 s / 0.509 m
C, shipped cascade	0.927 m, held 24.1 s, fails
C, regulator	0.392 m, held 25.4 s, passes
C, 40 and 6 N s/rad	0.721 m, passes

regulator on the same case catches the tap with 0.188 m of travel instead of 0.292 m, back inside ±2° in 0.83 s with an 8.48 N peak, and it makes the move faster, 90 % in 1.42 s, at the price of a deeper wrong-way start of 35.6 mm and a 10.0 mm overshoot. A faster cart loop pays for its speed in the depth of the wrong-way start, which is the right-half-plane zero at work.

🕮 **Pause.** Try Exercise 7.2 before reading on.

In case 2 the operator holds the drive with the pole 1° from upright. With no force applied the pole



Figure 7.4 The inverted pendulum application in scenario 1, Tuned balance, captured at $t = 12.00$ s with the cascade selected, a 10 ms sample interval and no measurement delay. The figure is a composite of the screen: the badge row and gauges, the scene and the three trends, with the header, the control panel and the scorecard cut. (1) The timing budget reads an e-fold of 179 ms against a sensing lag of 5 ms. (2) The cart stands at 0.498 m beside the target marker at 0.5 m with the pole upright. (3) The dashed ± 0.8 m challenge lines lie inside the bumpers at the ends of the rail. (4) The trends run from 2 to 12 s. At their left edge the 60 °/s tap leans the pole about 4.7° toward $+x$, the dashed lean request on the angle strip is pinned at −6° while the cart, pushed out to 0.29 m on a 5.8 N peak, is braked, and the pole swings to about −7°. (5) At 8 s the target steps to 0.5 m, the lean request jumps to 3.5°, and the cart dips below zero before it runs to the target. The force command carries spikes of a few tenths of a newton from single angle counts. The badge row reads 0.498 m of maximum travel and no rail hits.

lies on the deck at 0.944 s, at 1.326 s for a 1.0 m rod and at 0.741 s for a 0.3 m rod. At full speed every correction has to be seen, decided and applied inside one e-fold of 179 ms, while at the 0.1 speed setting the same fall lasts 9.44 s of wall time. A 2 N push lasting 50 ms, applied to the pole at rest upright, leaves it turning at 19.9 °/s toward $-x$ and the cart moving at 0.087 m/s 0.2 s after it starts. In automatic operation the cascade catches the same start with a 1.005° peak and 0.022 m of travel. A switch to automatic 0.5 s into the unaided fall, with the pole at 8.35° and 46 °/s, still recovers it with a 10.8° peak and 0.464 m of travel, and the latest switch that recovers it without touching a bumper is at 0.576 s, with the pole at 12.7°.

7.6.2 Floor, Ceiling and Delay

Case 3 raises the angle gain from 30 to 75 N/rad. The angle loop alone would stay stable at this gain, with its pendulum pair at $-2.42 \pm 17.42i\,\mathrm{s}^{-1}$ in the sampled-data model. In the cascade the cart loop's force gains grow with $K_{p\theta}$, so the s^3 coefficient of Equation (7.10) falls from 3.25 to 0.64 and

even the lag-free loop has a pair at $+0.14 \pm 13.21i\,\mathrm{s}^{-1}$. With the lags the pair is $+2.52 \pm 12.06i\,\mathrm{s}^{-1}$, and between 15 and 30 s the pole rocks $\pm 16.7°$ at 2.10 Hz, with the command beyond the 20 N limit 43.3 % of the time, the lean request on its clamp 73 % of the time and the cart shaking ± 0.129 m. Because the angle loop alone is stable, the clamp, which cuts off the outer loop, is what bounds the oscillation. At 55 N/rad the lag-free loop is stable at $-1.40 \pm 10.47i\,\mathrm{s}^{-1}$ and the sampled loop is not, at $+0.55 \pm 10.80i\,\mathrm{s}^{-1}$, so the pole rocks $\pm 9.3°$ at 1.75 Hz without saturating. Lowering the gain to 40 N/rad settles the loop to a $\pm 0.08°$ residual, and raising the damping to 8 N s/rad at 75 N/rad restores the s^3 coefficient to 3.64 and settles it to the same $\pm 0.08°$. A drive three times faster, 10 ms, does not, leaving $\pm 8.8°$. With the drive lag as the only lag, a 5 ms drive would bring the continuous loop to the edge at $+0.02 \pm 13.1i\,\mathrm{s}^{-1}$, but the rate filter and the hold keep the sampled loop unstable at $+0.52 \pm 13.0i\,\mathrm{s}^{-1}$, and the pole still rocks $\pm 7.9°$.

Case 8 replaces the encoder angle with a camera-like sensor that delivers a reading every 20 ms, 60 ms late, so the law sees the pole 70 ms in the past on average. The model puts an unstable pair at $+1.05 \pm 8.04i\,\mathrm{s}^{-1}$. After the tap the oscillation grows with a period of 0.68 s to a 32.7° swing, with the drive saturated 29 % of the time, the cart reaches the right bumper at 6.53 s, and the pole falls at 6.84 s. With a 20 ms delay the loop holds with a 7.6° peak and 0.321 m of travel. At 60 ms, 40 N/rad falls at 3.11 s, 20 N/rad holds with a 14.9° peak and 0.633 m of travel, and 15 N/rad falls at 4.01 s. The last failure is not a delay effect. The same gain falls at 5.03 s with the delay set to zero, because 15 N/rad is below the 18.6 N/rad floor that the cart loop imposes on the cascade although it is above the static floor of 11.3 N/rad.

7.6.3 The Cart Runaway

In case 4 only the angle loop runs, from a 2° start. The catch peaks at 2.01°, and from 1 s until the cart first meets the right bumper the pole leans between 0.38 and 1.53° forward, while the cart speed grows from 0.128 to 0.520 m/s between 1 and 4 s, a rate of $0.47\,\mathrm{s}^{-1}$ that matches the friction root of Section 7.4.2. The cart passes 0.5 m at 2.93 s and strikes the bumper at 4.11 s at 0.55 m/s. The 0.80 N s impulse through the pivot turns the pole by 132 °/s, the cart meets the bumper four more times in the next 0.14 s, and the pole falls at 4.54 s. The cascade from the same start keeps the cart within 0.046 m. A cart loop weakened to 1° per meter and 2° per m/s does not overcome the runaway. Its linearization has a pair at $+0.024 \pm 0.531i\,\mathrm{s}^{-1}$, and the cart swings with an 11.9 s period and a growing amplitude, 0.596 m by 40 s, crosses the 0.8 m line at 50.5 s, meets the bumper at 61.8 s and loses the pole at 62.5 s.

> **TRY IT IN THE APP** **4 · The Angle Loop Alone**
>
> Load the scenario and press Run. Watch the pole angle first: after the catch it never leans more than 1.53° forward. Then watch the cart. It creeps to the right, speeds up, strikes the right bumper at 4.11 s, and the jolt throws the pole down at 4.54 s.

7.6.4 A Longer and a Heavier Pole

Case 5 doubles the rod to 1.0 m at the same 0.15 kg. The unstable root falls from 5.58 to $3.98\,\mathrm{s}^{-1}$, so an unaided fall from 1° takes 1.326 s instead of 0.944 s, and yet the loop tuned for the short rod no longer settles. After the tap the cart swings out to 0.970 m, and from 20 to 60 s the pair rocks $\pm 8.3°$ at 0.70 Hz with the cart swinging ± 0.171 m and the lean request on its clamp 41 % of the time, against a predicted pair at $+0.50 \pm 4.61i\,\mathrm{s}^{-1}$, or 0.73 Hz. A stiffer angle loop of 40 N/rad and 7 N s/rad leaves $-0.061 \pm 5.76i\,\mathrm{s}^{-1}$, a mode that decays by a factor of e only every 16 s, so the pole stays inside $\pm 2°$ only from 15.12 s after the tap and still rocks $\pm 1.00°$ between 30 and 60 s. A softer cart loop of 4° per meter and 7° per m/s moves the pair to $-0.96 \pm 4.49i\,\mathrm{s}^{-1}$ and ends the rocking, but the tap then carries the cart into the bumper at 3.28 s. The regulator, redesigned for the new rod, is back inside $\pm 2°$ 1.01 s after the tap with 0.461 m of travel.

> **DID YOU KNOW? A Slower Pole the Loop Cannot Hold**
>
> From the simulator: a rod twice as long falls more slowly, reaching the deck at 1.326 s instead of 0.944 s with no force applied. Yet the cascade tuned for the short rod no longer settles, and the pole rocks ±8.3° at 0.70 Hz (Table 7.6).

Case 6 keeps the length and raises the rod mass to 1.0 kg. The unstable root moves only from 5.58 to 6.71 s^{-1}, but $(M+m)g$ rises from 11.3 to 19.6 N/rad and the ideal cascade now needs $K_{p\theta}$ above 29.5 N/rad, so the shipped 30 N/rad leaves a pair at $-0.042 \pm 3.99i\ s^{-1}$ that the lags push to $+0.022 \pm 3.92i\ s^{-1}$. The pole rocks ±9.8° at 0.60 Hz with the cart swinging ±0.178 m from 20 to 60 s. At 40 N/rad it is back inside ±2° 3.28 s after the tap with 0.295 m of travel, and the redesigned regulator takes 0.87 s and 0.194 m. The two cases show the same rocking with different causes, the damping coefficient in case 5 and the stiffness coefficient in case 6, and only the second answers to more angle gain.

7.6.5 The Recoverable Push

Case 7 raises the tap to 150 °/s. The cascade runs out of rail. For 0.77 s of the 0.88 s between the push and the first impact its lean request sits on the −6° clamp, with no drive saturation, and the cart strikes the right bumper at 2.88 s, meets it five more times by 3.04 s, and loses the pole at 3.29 s. With a 15° clamp the cart stops at 0.745 m after a 17.9° peak. The regulator, which has no clamp and weighs the cart position directly, stops it at 0.495 m with a 19.0 N peak and 0.07 s of saturation. The largest tap each configuration recovers without touching a bumper, found by bisection on the tap size, is about 130 °/s for the cascade, 183 °/s with the 15° clamp and 229 °/s for the regulator, and on the 5 °/s steps of the slider 130, 180 and 225 °/s are the largest that pass. The clamp decides the cascade's limit and the rail decides the regulator's, which from 230 °/s strikes a bumper without losing the pole and from 250 °/s loses it as well (beyond 245.2 °/s by bisection).

7.6.6 Swing-Up

Hanging at rest the pole is 0.736 J short of upright. The pump, limited to 4 m/s^2, needs three turning points, and the regulator captures the pole at 2.67 s with the cart never more than 0.496 m from center and a 6.07 N force peak, then returns the cart to center without a release. With a 2 m/s^2 limit the swings stall about 27° short of upright, just outside the capture wedge, and the pole is never caught in 30 s. With 6 m/s^2 it is caught at 2.06 s, but the cart travels 0.743 m and the force reaches 9.18 N. A 5° capture wedge, or an energy gain of 20 m/s^2 per J, never catches it, and the pump keeps swinging the pole.

> **TRY IT IN THE APP 9 · Swing-Up From Hanging**
>
> Load the scenario and press Run, and watch the pole energy gauge and the capture wedge above the pivot. After three turning points the regulator captures the pole at 2.67 s. Press Reset, set Pump acceleration limit to 2 m/s^2 and run again: the swings stall about 27° short of upright, just outside the wedge, and the pole is never caught in 30 s.

7.6.7 The Graded Challenge

The challenge starts the pole 5° from upright and pushes it at 125 °/s at 3 s. Over 30 s the cart must never reach ±0.8 m or touch a bumper, the pole must stay up, and it must be held inside ±2° for at least 20 s. Every change is made before the run, and a changed model parameter or disturbance voids the attempt. The shipped cascade holds the angle for 24.1 s but carries the cart to 0.927 m, 0.127 m past the line, which it crosses at 3.89 s, a pass on the angle and a failure on the hard limit. The suggested regulator, which uses the default weights ($x_s = 0.5$ m, $\theta_s = 20°$, $F_s = 5$ N), settles 1.63 s after the push, holds for 25.4 s, 5.4 s more than required, and keeps the cart within 0.392 m, 0.408 m inside the line, with 0.03 s of saturation. It spends more actuator travel to do so, 138.1 N of

Table 7.7 Regulator and MPC With One Model and One Set of Weights, Runs of 30 s

Case	Law	θ_{pk} (°)	RMS (°)	x_{pk} (m)	t_{lim} (s)	Travel (N)	t_{set} (s)	Plans
5, long pole	LQR	11.9	1.41	0.461	0	167.5	1.01	—
	MPC	11.9	1.41	0.461	0	160.2	1.01	3000
7, 150 °/s	LQR	20.0	2.07	0.495	0.07	129.5	1.85	—
	MPC	17.6	1.98	0.504	0.08	138.3	1.97	3000
8, delay	LQR	fell 2.98 s		1.000	0.66	67.6	—	—
	MPC	fell 3.00 s		1.000	0.68	68.3	—	154
10, 205 °/s	LQR	31.2	3.22	0.809	0.15	142.5	2.04	—
	MPC	28.0	2.97	0.768	0.24	152.7	2.20	3000
C[a]	LQR	16.2	1.72	0.392	0.03	138.1	1.63	—
	MPC	15.1	1.67	0.396	0.03	146.2	1.75	3000

θ_{pk} is the largest angle after the push and x_{pk} the largest cart distance from center, both from the true state at the physics step, t_{lim} the time with the command at ±20 N, Travel the summed change of the command, and RMS and t_{set} the scorecard's angle RMS and recovery inside ±2°. Plans are those of the MPC, none rejected. In case 8 both carts reach the bumper. [a]Challenge, for information.

summed force change over the 30 s run against 81.7 N for the shipped cascade. The cascade with 40 N/rad and 6 N s/rad also passes, at 0.721 m.

7.6.8 Regulator Against Predictive Control

Table 7.7 sets the regulator beside the MPC on case 10 and on the three earlier cases in which a limit or a model error could separate them, the long pole of case 5, the hard push of case 7 and the delay of case 8, with the challenge for information. Each case is loaded, the structure is selected before the run, and every run lasts 30 s. The regulator keeps its shipped weights, which are also the suggested setting of the challenge, and the MPC keeps the settings of case 10 in every case. Both read the same quantized instruments, and the application draws no random numbers, so both meet the same push at the same instant. The lean limit and the prices were chosen on case 10 itself, and its push was chosen as the first step of the 5 °/s push slider past the largest push that the regulator holds inside the line, the step at which the two laws differ in kind.

In case 10 the 205 °/s push at 2 s carries the regulator's cart across the line at 2.46 s to 0.809 m, while the MPC stops it at 0.768 m, as shown in Figure 7.6, and Figure 7.5 sets the two laws side by side on every case of the table. Both laws drive at +20 N from 2.022 s to get the cart under the pole. The first plan that the rail limit shapes comes at 2.060 s, 0.40 s before the regulator's cart reaches the line. The MPC then holds +20 N to 2.18 s against 2.17 s and brakes at −20 N from 2.29 to 2.37 s, where the braking force of the regulator peaks at 16.8 N, and the pole leans back 28.0° instead of 31.2°. Only 64 of the 3000 plans have a limit binding. The cost is a recovery 0.16 s slower and 10.2 N more force travel. The plans are not forecasts. A plan made at 2.17 s put the cart at 0.918 m, 118 mm past the line on its slack, and the cart stopped at 0.768 m, because each plan is replaced 10 ms later by one made from a new measurement, and the model leaves out the drive lag, the slew limit, the delay of the readings and the sine of the lean.

The pushes over which only the MPC keeps the cart inside the line form a narrow band. On the 5 °/s steps of the push slider the largest push held inside the line rises from 200 to 205 °/s, and by bisection from 203.6 to 209.6 °/s. By bisection the largest push kept off the bumpers rises from 229.4 to 230.2 °/s, and the largest push that keeps the pole from 245.2 to 248.6 °/s. At 240 °/s the rail slack of the plan reaches 414 mm, and both carts meet the bumper once and keep the pole. The travel after a push is set mostly by the push and the 20 N drive, because the cart must run under the pole and then lean it back to stop.

Where no limit binds the two laws are one law. In case 5 no limit shapes any plan, and the two commands agree to 3.8×10^{-9} N until the first encoder reading differs at 5.05 s. From there the quantizer dither of Section 7.7 gives the two runs different count crossings and a force travel of 160.2 against 167.5 N, while the peak angle, the RMS, the excursion and the recovery agree to the digits printed. On the rig of case 10, pushes of 60 to 100 °/s reproduce the regulator's excursion to 0.1 mm.

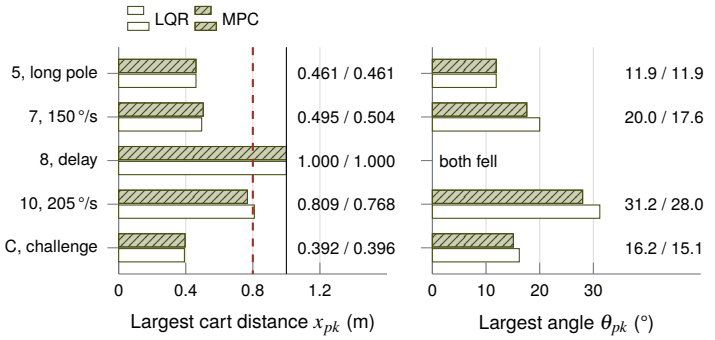


Figure 7.5 The regulator (white) and the predictive controller (hatched) of Table 7.7, with one model and one set of weights: the largest cart distance from center against the 0.8 m line (dashed) and the bumper at 1.0 m (thin solid), and the largest angle after the push. Values read LQR / MPC. Notice where the bars part: the predictive controller keeps the cart of case 10 inside the line, gives up 8.5 mm of travel to its lean limit in case 7, and falls with the regulator in case 8.

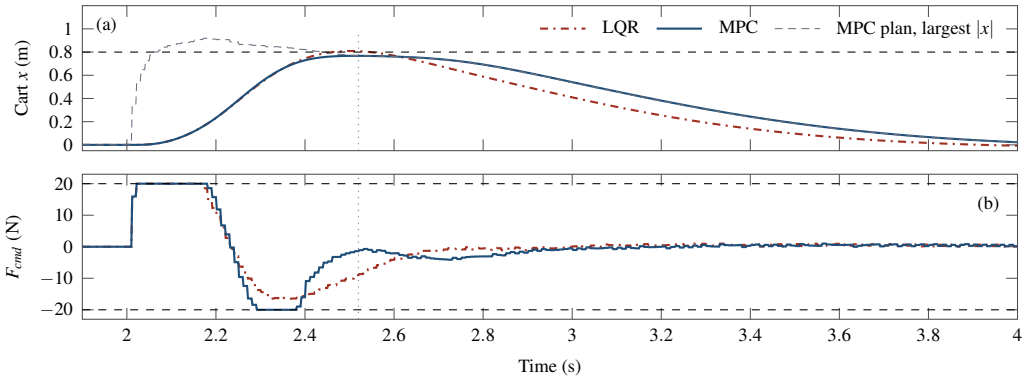


Figure 7.6 Case 10 under the regulator and under the MPC, which share one model and one set of weights, from 1.9 to 4.0 s with the 205 °/s push at 2 s, the regulator dash-dotted and the MPC thick solid. (a) Cart position against the 0.8 m line (thin dashed), with the largest cart position in each plan of the MPC (gray, densely dashed). (b) Force command against the 20 N drive limits (thin dashed). At the dotted line, $t = 2.52$ s, the regulator's cart is at 0.808 m past its peak of 0.809 m and the MPC's at 0.767 m.

The MPC loses case 7 on travel and recovery while it leans less. The cart never nears the line, and only the 15° lean limit binds, on the backward lean that brakes the cart. The MPC leans back 17.6° against 20.0°, with 9.4° forward in both runs, so it brakes the cart less, and the cart runs to 0.504 m against 0.495 m and recovers in 1.97 against 1.85 s. With $\theta_{max} = 45°$ the MPC reproduces the regulator's 0.495 m and 1.85 s. The same mechanism carries the MPC's cart 0.8 to 10.3 mm further than the regulator's for pushes of 120 to 180 °/s in steps of 20 °/s. Neither law survives case 8, whose 60 ms delay and 20 ms samples are in neither model. Both carts cross the line at 2.87 s and strike the bumper, the regulator loses the pole at 2.98 s and the MPC at 3.00 s. On the challenge both pass, with 0.392 and 0.396 m of travel and the pole held for 25.4 and 25.3 s.

The 30 ms drive lag, which neither model contains, is also why the limits cannot be priced higher. At ten times the shipped rail and lean prices the cart reaches 0.813 m with 0.53 s at the drive limit, and at a hundred times the pole falls at 2.68 s. With the drive lag cut to 5 ms the hundredfold prices keep the pole and stop the cart at 0.767 m, but the regulator on that faster drive stops it at 0.740 m, so

Table 7.8 The 150 °/s Push of Case 11 With and Without the Terminal Weight, Runs of 15 s

Law	T_p (s)	Pole		θ_{pk} (°)	x_{pk} (m)	t_{lim} (s)	Plans
Cascade	—	fell	3.29 s	—	1.000	0.30	—
Angle PD	—	fell	2.98 s	—	1.000	0.22	—
LQR	—	held		20.0	0.495	0.07	—
MPC, P	0.3	held		17.4	0.509	0.07	1499
MPC, P	0.5, 0.6, 1	held		17.6	0.504	0.08	1499
MPC, Off	0.3	fell	3.04 s	—	1.000	0.24	303
MPC, Off	0.5	fell	7.54 s	—	1.000	0.43	754
MPC, Off	0.6	held		11.9	0.793	0.00	1499
MPC, Off	1	held		15.9	0.547	0.05	1499

Columns as in Table 7.7, with θ_{pk} for the runs that hold. Every fall follows bumper contact, first at 2.88, 2.76, 2.79 and 6.61 s. No plan is rejected.

a quicker actuator buys more rail than the predictive controller does. Case 10 makes 3000 plans in 30 s, and 48 runs at the corners of the horizon, the two limits of the plan and the sample interval, on cases 1 and 10 and the challenge, reject no plan and need at most 28 iterations of the solver. A plan takes about 90 µs in Node.js and 105 µs in the Chrome browser, the difference between whole runs of case 10 under the MPC and under the regulator divided by the plans, or 1.05 % of one processor core at real time. At $T_s = 2$ ms with $T_p = 2$ s a plan has 1000 steps and takes 0.92 ms, 9.1 % of a core at real time and 18.3 % at twice real time.

7.6.9 Terminal Cost at a Short Horizon

Case 11 is the rig of case 1 with the 150 °/s push of case 7 at 2 s, under the MPC with the settings of case 10 except a horizon of 0.3 s, the shortest the application offers, and no terminal weight. Table 7.8 sets every shipped structure, the MPC with P and the MPC without it on that push, each run for 15 s from the same instant with the same instruments, the terminal weight and the horizon being the only settings varied. The regulator alone is the shipped structure given the same signal, because its cost-to-go is the terminal weight.

As loaded, the plan without P catches the push with at most 14.2 N before the cart reaches the bumper, against the 20 N that the plan with P applies, and it never leans the pole back to brake. The cart reaches 1.70 m/s, the rail limit first shapes a plan at 2.38 s, and from then on paying the rail slack inside 0.3 s costs the plan less than a backward lean, so at 2.50 s its predicted pole ends 11.4° forward. The cart strikes the right bumper at 2.79 s, the impacts throw the pole, and it falls at 3.04 s. Each of its 303 plans was optimal and accepted, because the plan solved the problem that it was given. Restoring P at the same horizon holds the pole with 0.509 m of travel, against 0.495 m for the bare regulator. The difference is the 15° lean limit, which shapes its plans from 2.13 to 2.43 s, and with a 45° limit the plan reproduces the regulator's 0.495 m. The cascade and the angle loop lose the same push at 3.29 and 2.98 s.

Longer horizons without P hold the pole only well past the linear crossing of Section 7.5.6. On case 11's push, 0.45 s falls at 3.65 s, 0.46 s at 5.20 s and 0.5 s at 7.54 s, after the cart creeps into the bumper at 6.61 s, while 0.55 s holds with 0.823 m of travel. Run for 30 s at every horizon of the setting from 0.3 to 2 s in steps of 0.1 s, the shortest horizon from which Off holds the pole at that and every longer horizon is 0.5 s for pushes of 60 and 100 °/s, 0.6 s for 150 and 205 °/s and 0.9 s for 240 °/s. With P every horizon holds every one of these pushes, with one to three bumper contacts at 240 °/s, where the bare regulator also strikes the bumper once. Halving the physics step to 0.25 ms leaves every verdict and every travel unchanged and moves the fall at 0.5 s from 7.54 to 7.57 s.

The hold at 0.3 s is therefore owed to the terminal weight. The constraints are not the cause, because Off keeps them and falls, and the horizon is not the cause, because with P it hardly matters, at 0.509, 0.504 and 0.504 m for 0.3, 0.6 and 1 s. The terminal weight is the regulator's cost-to-go, so the gain of the plan with P over the plan without it is the regulator's. A horizon long enough to do

without *P* buys stability and not the regulator's performance. Off carries the cart 0.793 m at 0.6 s and 0.547 m at 1 s, against 0.504 m with *P*.

The same holds across the whole sequence. Over the ten earlier cases and the challenge, with the challenge counted once at its shipped settings and once at the suggested ones, twelve runs of 30 s for each structure, the pole falls in 11 of 12 runs under the angle loop, 4 under the cascade, 2 under the regulator, 2 under the MPC with *P* and 2 under Off at the shipped 1 s horizon, the two being the manual case and the delay of case 8 for the last three, and in 11 under Off at 0.3 s, which loses every case that any law balances. The swing-up start counts as no balancing case for these laws, because the watchdog holds the drive at zero while the pole hangs. Off at 1 s holds every case that the MPC with *P* holds but carries the cart farther in each, 0.539 m against 0.510 m in case 1, 0.432 against 0.396 m on the challenge, and 0.823 against 0.768 m in case 10, across the 0.8 m line, with a smaller peak lean, and in case 8 it falls at 2.94 s against 3.00 s. Off is thus less safe than the shipped plan wherever the two differ, and the application states this beside the setting.

The terminal weight costs nothing to compute, because it changes the weights of the last step and not the size of the quadratic program. In Node.js, measured while other simulations loaded the computer, a plan of case 11 takes 32 µs as loaded and 59 µs at 0.6 s, and case 10 takes 95 and 96 µs with and without *P* at 1 s. The cost grows with the horizon. Timed in Chrome while the application animates itself under the same load, the whole controller of case 10 as shipped occupies 2.6 to 3.2 % of one processor core at 1× speed and 4.8 to 5.7 % at 2× over two passes. The 2 s horizon without *P* takes 4.8 to 5.2 % and 8.7 to 9.3 %, close to the limit of 10 % of one core at the top speed that the apps are held to, and with *P* at 2 s it takes 4.3 to 5.1 % and 7.5 to 9.8 %, so the cost belongs to the horizon and not to the terminal weight.

7.7 Repeatability of a Quantized Loop

Reset (Section 1.6.1) restarts a case from $t = 0$ with the settings in force and replays the run bit for bit on one JavaScript engine, and yet the encoder quantizers make the fine detail of a run sensitive to arithmetic. Perturbing the case 1 tap by one to twenty parts in 10^{13}, in either direction, moves the cart position at 12 s between 0.4962 and 0.4984 m and the summed force change between 37.2 and 40.0 N. A last-bit difference changes the sample at which the angle crosses a count boundary, and the rate estimate turns a single count into a kick of 0.37 N through $K_{d\theta}$. Linear feedback of quantized measurements on an unstable plant is known to produce pseudorandom rather than convergent trajectories [24], and here the result is a deterministic dither that keeps the angle within 0.14° of upright and the command within ±0.42 N at rest. The screenshot behind Figure 7.4 was captured by stepping the same engine in a browser, and its scorecard, cut from the figure, matches the headless run behind the tables to the last digit, with 0.498 m of travel and a summed force change of 39.85 N, because the browser and the headless run perform the same arithmetic.

7.8 Technology in Practice

What Measures It

Laboratory carts and poles read both outputs with quadrature encoders, one on each cart and one on the pendulum [6]. A rig that must also see a flexing pole adds a rate gyro at its tip, because an angle read at the base alone can match a pole that bends and a pole that does not. Comparing the tip's rate with the angle at the base then reveals the flex [25]. The chapter's encoders round the cart to 0.02 mm and the pole to 4096 counts a turn, carry no noise and miss no count, and the delay line of case 8 stands in for a camera (Section 7.3).

TECHNOLOGY Sensor: Encoders Give Position, Not Speed

On one laboratory rig at a space flight center, every cart and the pendulum carry a high-resolution quadrature encoder. The encoders measure position only, so a state estimator supplied the velocities that full state feedback needs [6]. A second rig read its pole with a rate gyro instead, on ranges from ± 250 to $\pm 2000\,°/s$, where a wider range costs precision. Its angle came from integrating the rate, and the gyro's bias showed at rest as a steady walk to $-10°$ after 25 s, cut to $\pm 1°$ once the bias was subtracted [25]. The chapter's encoders are exact to one count and drift-free, and its rates are differences filtered over $\tau_f = 15\,\text{ms}$.

What Moves It

The motor that moves the cart is a direct-current motor on one rig and servos turning the wheels of a mobile base on another. The first rig commands its direct-current motor with a voltage and converts volts to newtons with a hardware gain identified on the rig [6]. The second commands continuous-rotation servos whose signal drops out past the ends of its range, so a shaping function maps the controller's command onto the usable band on each side of the dead band [25]. The chapter's drive is the idealized end of that chain, a force source with a 30 ms lag, a slew limit of 1000 N/s and a limit of 20 N, and its Assumptions name the backlash and back electromotive force that it lumps into the lag.

TECHNOLOGY Final Element: A Voltage-Driven Motor With a Dead Band

The space-center rig drives its carts with a direct-current motor commanded by a voltage, and a hardware gain turns volts into newtons in the model. The gain found by experiment proved wrong once the feedback gains were tried, and the pendulum stayed up only about ten seconds while the gain was checked [6]. On a wheeled rig a band of commands around neutral gives no motion, the band shifts from time to time, and a routine steps the command until the motor encoders report motion [25]. On a trend, friction of this kind shows as a cart that waits until the pole has fallen far enough, then moves, in a repeating oscillation [26]. The chapter's drive is a lagging force source that lumps all of this.

Where the Controller Runs

On the space-center rig a data acquisition board first ran the control algorithm. It was replaced by a graphical programming environment on a computer with interface cards, and a design script discretized the two-cart model at a sampling period of 0.001 s [6]. The chapter's law runs every 10 ms on a 2 ms clock. The state feedback $F = -K\tilde{x}$, with the lean clamp of the cascade or the plan of the predictive controller around it, is the one block that holds the chapter's structure, and the swing-up adds a switch on the capture region. Figure 7.7 draws the loop on rig hardware.

What Is Modernizing It

The first development is sensing of the pole's flexure. A wheeled rig with a flexible aluminum pole, built as an analogue of a launch vehicle, added the tip gyro, and work continued on sensor blending and direct output adaptive control to cancel the flex dynamics. Its models were not yet accurate enough to move a simulated controller onto the hardware [25]. The second is the safety filter. A model predictive safety certification changes the input of a learned controller only when the cart and pole approach its constraints, and on a simulated cart and pole it kept a learned controller inside constraints that the controller broke alone [27]. The chapter's predictive controller acts in the same way, departing from the regulator only when a limit binds. The third is open benchmarks. Simulated cart and pole environments now randomize inertial properties and initial states and set state and input constraints, and a reinforcement learner trained there used over nine hours of simulated data [27].

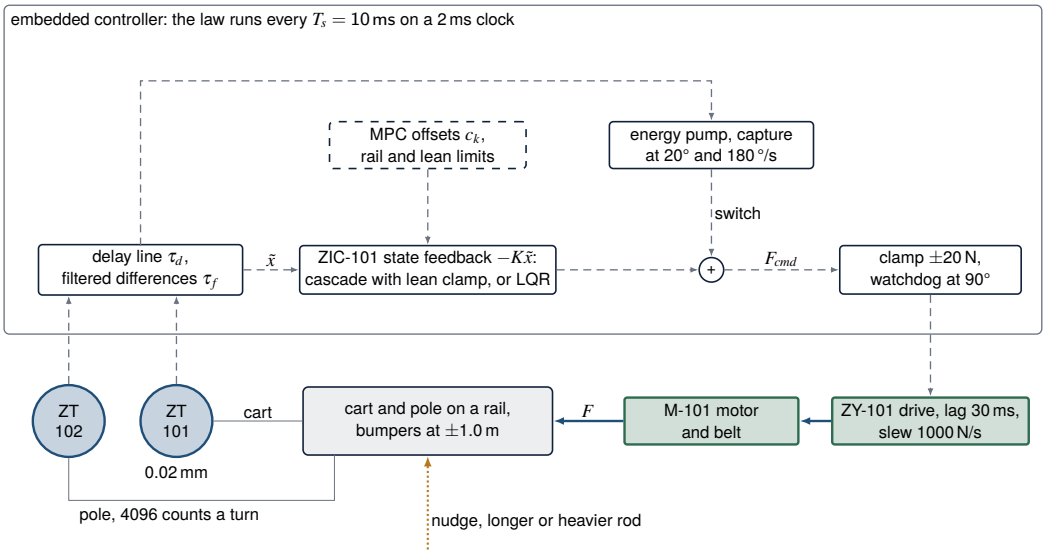


Figure 7.7 The balance loop on rig hardware. The encoders ZT-101 on the cart and ZT-102 on the pole feed a delay line and filtered differences, which form the state $\tilde{x}$. The state feedback ZIC-101 holds the cascade or the regulator, the predictive controller (dashed) adds offsets to the regulator's move, and in the swing-up the energy pump drives until the pole enters the capture region. The clamped command goes to the drive ZY-101 and the motor M-101. Notice that one force serves both outputs, so every block above the drive shares one limit of 20 N.

> **IN PRACTICE** **A Fast Pair for the Pole, a Slow Pair for the Cart**
>
> A textbook design for a cart and pendulum splits the closed loop into fast dynamics that hold the pendulum upright and slower dynamics that place the cart. The fast pair is set from the natural frequency of the pendulum hanging down, with a damping ratio of 0.5. That first design gave the desired step response but asked for a peak force almost three times the force of gravity, so the fast pair was slowed by about a factor of three and the cart was moved more slowly still [26]. The peak force, not the step response, decided the redesign. The chapter's cascade, a fast angle loop inside a slow cart loop, makes the same split.

7.9 Limitations

The rod is rigid, the friction is viscous and the encoders are noise-free, so the small limit cycles that stick-slip friction produces on real carts do not appear [16].

The drive is a lagging force source. A belt or rack drive with backlash, back electromotive force and a voltage limit would add lag and lower the ceiling below the 50.7 N/rad that the present lags allow.

None of the laws has integral action, and no constant force acts on the cart, so the steady offset that a tilted rail would cause is never exposed.

The regulator is designed with a perfect model of the pole and redesigned instantly, which makes cases 5 and 6 more favorable to it than a real redesign would be.

The linear analysis leaves out the quantizers, the clamps and the limits, so its window is necessary and not sufficient for recovery, as the exception in Figure 7.2 shows.

OPEN CHALLENGE **Certifying Where a Learned Controller Is Safe**

Learning-based control is often certified with a region of attraction, the states from which the controller is guaranteed to bring the system back to its equilibrium. A review of safe learning finds that the guarantees rest on assumptions, such as bounded disturbances, hard to verify before a robot runs. It calls for systematic ways to check them with little data [27]. Here the linear window agrees with 39 of 40 nonlinear runs, and the exception shows it necessary for recovery, not sufficient.

The predictive controller plans on the regulator's model, linear at upright and without the drive lag, the slew limit, the measurement delay or the sine of the lean, and it is given the regulator's instant redesign after every change of pole. Its plans are therefore not forecasts, and its limits priced a hundred times higher ask for force changes that the 30 ms drive delivers too late. It has no disturbance estimator, its rail and lean limits are soft and judged every 50 ms of the plan, and its prices were chosen on case 10, the case it wins.

The crossing at 0.46 s is a linear, unconstrained result on the plan's own model at the default sample interval, weights and pole, and it moves when any of them changes. It is a lower estimate of the horizon that a push needs without P and not that horizon, because the bumpers, the drive limit and the lag lie outside the analysis. The application offers no terminal constraint set, the other ingredient of the usual stability proof for constrained plans [18], so the argument for the plan with P covers only states from which no limit is active at the end of its horizon.

7.10 Conclusion

A cart and pole is an unstable plant with one actuator, two outputs and a right-half-plane zero, and the simulator of this chapter puts each of those properties within reach of an experiment. Its unstable root of $5.58\,\mathrm{s}^{-1}$, one e-fold every 179 ms, leaves a hand on the drive that much time for each correction at full speed and ten times as long at the 0.1 speed setting. An angle loop alone needs more than $(M+m)g = 11.3\,\mathrm{N/rad}$, and with rail friction it runs the cart into a bumper 4.11 s after a 2° start while the pole leans no more than 1.53°. The cascade that cures the runaway starts every move the wrong way, 21.3 mm for a 0.5 m step, and its characteristic polynomial explains why its floor is 17.4 N/rad and why one tuning rocks a 1.0 m rod through its damping and a 1.0 kg rod through its stiffness. An exact sampled-data linearization agrees with 39 of 40 nonlinear runs and closes the gain window at about 105 ms of average sensing lag. The regulator stops a 150 °/s push at 0.495 m where the clamped cascade loses the pole, and the energy pump brings the hanging pole to capture in 2.67 s. A predictive controller on the regulator's model and weights equals the regulator until a limit binds, stops the cart of case 10 at 0.768 m where the regulator's crosses the line to 0.809 m, gives up 8.5 mm of travel to its own lean limit in case 7, and falls with the regulator in case 8. Without its terminal weight the same plan is an unstable law below a 0.46 s horizon and loses at 0.3 s the 150 °/s push that it holds with the weight, while a horizon long enough to hold without it carries the cart farther.

The model is rigid, viscous and noise-free. A fuller model would add stick-slip friction, encoder noise, drive backlash and a regulator designed on a mismatched pole. Within those bounds the application is a laboratory for balancing an underactuated machine within hard limits.

KEY POINTS

- A pole on a cart falls on its own, and only moving the cart under it holds it up.
- Rail friction runs the cart of an angle loop alone into a bumper 4.11 s after a 2° start.
- The cascade that cures the runaway starts every move the wrong way, 21.3 mm for a 0.5 m step.
- The energy pump brings the hanging pole to capture in 2.67 s.
- A predictive controller on the regulator's model equals the regulator until a limit binds.

Where else this applies. Volume 3, Chapter 10 meets the same wrong-way start in a rocket.

References

[1] F. Grasser, A. D'Arrigo, S. Colombi, and A. C. Rufer, "JOE: a mobile, inverted pendulum," *IEEE Transactions on Industrial Electronics*, vol. 49, no. 1, pp. 107–114, 2002.

[2] S. Kajita, F. Kanehiro, K. Kaneko, K. Yokoi, and H. Hirukawa, "The 3D linear inverted pendulum mode: a simple modeling for a biped walking pattern generation," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, 2001, pp. 239–246.

[3] K. H. Lundberg and T. W. Barton, "History of inverted-pendulum systems," *IFAC Proceedings Volumes*, vol. 42, no. 24, pp. 131–135, 2010.

[4] J. L. Cabrera and J. G. Milton, "On-off intermittency in a human balancing task," *Physical Review Letters*, vol. 89, no. 15, p. 158702, 2002.

[5] O. Boubaker, "The inverted pendulum benchmark in nonlinear control theory: A survey," *International Journal of Advanced Robotic Systems*, vol. 10, no. 5, p. 233, 2013.

[6] C. L. Wilson and P. Capo-Lugo, "Robust control algorithm for a two cart system and an inverted pendulum," NASA Marshall Space Flight Center, Huntsville, AL, Internship final report NTRS 20110015819, 2011. [Online]. Available: https://ntrs.nasa.gov/citations/20110015819

[7] J. Freudenberg and D. Looze, "Right half plane poles and zeros and design tradeoffs in feedback systems," *IEEE Transactions on Automatic Control*, vol. 30, no. 6, pp. 555–565, 1985.

[8] K. J. Åström, "Limitations on control system performance," *European Journal of Control*, vol. 6, no. 1, pp. 2–20, 2000.

[9] G. Stépán, "Delay effects in the human sensory system during balancing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 367, no. 1891, pp. 1195–1212, 2009.

[10] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[11] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[12] S. Mori, H. Nishihara, and K. Furuta, "Control of unstable mechanical system Control of pendulum," *International Journal of Control*, vol. 23, no. 5, pp. 673–692, 1976.

[13] K. Furuta, M. Yamakita, and S. Kobayashi, "Swing-up control of inverted pendulum using pseudo-state feedback," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 206, no. 4, pp. 263–269, 1992.

[14] K. J. Åström and K. Furuta, "Swinging up a pendulum by energy control," *Automatica*, vol. 36, no. 2, pp. 287–295, 2000.

[15] A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuronlike adaptive elements that can solve difficult learning control problems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 5, pp. 834–846, 1983.

[16] S. A. Campbell, S. Crawford, and K. Morris, "Friction and the inverted pendulum stabilization problem," *Journal of Dynamic Systems, Measurement, and Control*, vol. 130, no. 5, p. 054502, 2008.

[17] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control: Optimization, Estimation, and Control*. Routledge, 2018.

[18] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[19] B. Kouvaritakis, J. A. Rossiter, and J. Schuurmans, "Efficient robust predictive control," *IEEE Transactions on Automatic Control*, vol. 45, no. 8, pp. 1545–1549, 2000.

[20] P. O. M. Scokaert and J. B. Rawlings, "Constrained linear quadratic regulation," *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1163–1169, 1998.

[21] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[22] J. B. Rawlings and K. R. Muske, "The stability of constrained receding horizon control," *IEEE Transactions on Automatic Control*, vol. 38, no. 10, pp. 1512–1516, 1993.

[23] A. Jadbabaie and J. Hauser, "On the stability of receding horizon control with a general terminal cost," *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 674–678, 2005.

[24] D. F. Delchamps, "Stabilizing a linear system with quantized state feedback," *IEEE Transactions on Automatic Control*, vol. 35, no. 8, pp. 916–924, 1990.

[25] N. A. Janmohamed and M. A. Dupuis, "Modeling and control for a flexible inverted pendulum robot," NASA Kennedy Space Center, Kennedy Space Center, FL, Internship final report NTRS 20205009993, 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205009993

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

[27] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

7.11 Exercises

7.1 ○ Concept check · 2 min

In case 4 the angle loop alone holds the pole within 1.53° of upright, yet the cart runs into the right bumper at 4.11 s. Which statement explains the runaway?

(a) Rail friction makes the constant term of the loop's polynomial negative, so one real root is positive at any gain

(b) The angle gain of 30 N/rad is below the floor $(M+m)g$, so the pole is held only loosely and drags the cart

(c) The quantized angle reading adds a small mean force through the derivative term that pushes the cart to one side

(d) The 30 ms drive lag delays every correction, and the delayed corrections add up to a growing cart speed

7.2 ○ Concept check · 1 min

In case 1 the cart target steps from 0 to 0.5 m at 8 s. What does the cart do first under the cascade?

(a) It moves toward the target at once, while the pole leans away from the target to balance the push

(b) It backs about 21 mm away from the target, which leans the pole toward the target before the cart follows

(c) It stays still until the lean request reaches the 6° clamp, and then it moves toward the target

(d) It moves slowly toward the target, because a cart gain of 7° per meter asks for very little lean

7.3 ○ Concept check · 2 min

The MPC of case 10 plans offsets c_k from the regulator's move $-K\tilde{x}$, with the regulator's Riccati solution P as its terminal weight. When no limit binds anywhere in the plan, what force does it apply?

(a) The regulator's move, but only once the horizon is longer than the crossing at 0.46 s

(b) A force smaller than the regulator's, because the weight W on the offsets is larger than R

(c) The regulator's move at every horizon, because the cost is $\tilde{x}_0^T P \tilde{x}_0$ plus $W\sum c_k^2$

(d) A force that already brakes for the 0.8 m line, because the plan looks one second ahead

7.4 ▭ Theory and calculation · 6 min

The default rig has a cart of $M = 1.0\,\mathrm{kg}$ and a uniform rod of $L = 0.5\,\mathrm{m}$ and $m_p = 0.15\,\mathrm{kg}$ with no tip mass, and $g = 9.81\,\mathrm{m/s^2}$. Neglect both frictions.

(a) Find the first moment $m\ell$, the inertia J about the pivot, the length $\ell_{eq} = J/m\ell$ and $\Delta_0 = (M+m)J - (m\ell)^2$. *Checkpoint:* 0.0375 kg m, 0.0125 kg m^2, 0.333 m and 0.012 97 kg^2 m^2.

(b) Find the unstable root λ_0 from Equation (7.7), its e-folding time, and the force per radian $(M+m)g$ that holds a tilted pole still. *Checkpoint:* 5.71 s^{-1}, 175 ms and 11.3 N/rad.

(c) Find the wrong-way zero $z = \sqrt{g/\ell_{eq}}$, how far it lies below λ_0, and its value for the 1.0 m rod of case 5. *Checkpoint:* 5.42 rad/s, 5 % below, and 3.84 rad/s.

7.5 Theory and calculation · 6 min

The controller sees the pole through an encoder of 4096 counts per turn, and it forms the rate as the filtered difference of Equation (7.6) with $\tau_f = 15\,\text{ms}$.

(a) At the default $T_s = 10\,\text{ms}$, find the filter coefficient a and the size of one count in degrees. *Checkpoint:* 0.487 and 0.0879°.

(b) One count changes between two samples. Find the step this causes in the rate estimate $\hat{\omega}$, and in the force through $K_{d\theta} = 5\,\text{N s/rad}$ and through $K_{p\theta} = 30\,\text{N/rad}$. *Checkpoint:* 0.075 rad/s, 0.37 N and 0.046 N.

(c) In case 8 ($T_s = 20\,\text{ms}$, $\tau_d = 60\,\text{ms}$), find the average sensing lag as a fraction of the 179 ms e-folding time, and compare it with the lag at which the window of Figure 7.2 closes. *Checkpoint:* 70 ms, 0.39, against about 105 ms, 0.59.

7.6 Theory and calculation · 7 min

Inside its lean clamp the cascade is four-state feedback with $k_x = K_{p\theta}K_{px}$ and $k_v = K_{p\theta}K_{dx}$, and with the frictionless model it closes into the quartic of Equation (7.10). Use the defaults $K_{p\theta} = 30\,\text{N/rad}$, $K_{d\theta} = 5\,\text{N s/rad}$, $K_{px} = 7^\circ$ per meter and $K_{dx} = 10^\circ$ per m/s, with the values of Δ_0, $m\ell$ and ℓ_{eq} for the default rod.

(a) Find k_x, k_v and the five coefficients. *Checkpoint:* 3.67 N/m, 5.24 N s/m, and 0.346, 3.25, 17.5, 51.4, 36.0.

(b) Find the s^3 coefficient for the 1.0 m rod of case 5 ($\ell_{eq} = 0.667\,\text{m}$) and the s^2 coefficient for the 1.0 kg rod of case 6 ($M + m = 2.0\,\text{kg}$). *Checkpoint:* 1.51 and 9.16.

(c) Find the s^3 coefficient with $K_{p\theta} = 75\,\text{N/rad}$ (case 3), first with $K_{d\theta} = 5$ and then with 8 N s/rad. Explain why more angle gain cannot restore this coefficient. *Checkpoint:* 0.64 and 3.64.

7.7 App challenge · 8 min

Load scenario 4, The angle loop alone, and change no setting. Use the default seed (the app draws no random numbers).

(a) Press Run and let the run go until the app pauses itself, about 0.6 s after the pole falls. Read the Pole and Rail hits tiles. *Checkpoint:* Fell at 4.54 s, with 5 rail hits.

(b) Export the CSV file and read the cart speed, column `v_mps`, at 1 and 4 s. Find the average growth rate of the speed between them and compare it with the estimate $b_c g/(K_{p\theta} - (M+m)g)$. *Checkpoint:* 0.128 and 0.520 m/s, a rate of 0.47 s^{-1} within 0.01 s^{-1}, against 0.52 s^{-1}.

(c) Press Reset, set Controller structure to Cascade before pressing Run, and run to 10 s. Read the Cart travel tile. *Checkpoint:* 0.046 m, within 0.002 m.

7.8 App challenge · 8 min

Load scenario 7, A hard push, which taps the pole at 150 °/s at 2 s. Use the default seed and run each part until the app pauses itself after a fall or until 8 s.

(a) Run the scenario as shipped. Read the Pole and Rail hits tiles. *Checkpoint:* Fell at 3.29 s, with 6 rail hits.

(b) Press Reset, switch to Advanced, raise Lean limit on the cart card from 6 to 15° and run. Read the Cart travel tile. *Checkpoint:* 0.745 m, within 0.005 m.

(c) Press Reset, set Controller structure to LQR and run. Read the Cart travel and Force saturated tiles. *Checkpoint:* 0.495 m within 0.005 m, and 0.07 s.

State which limit decides the push each configuration survives.

7.9 App challenge · 8 min

Load scenario 11, Terminal cost: Off against LQR at a short horizon. The MPC plans 0.3 s ahead with no terminal weight, and the pole takes a 150 °/s push at 2 s. Use the default seed.

(a) Run as loaded until the app pauses itself. Read the Pole tile. *Checkpoint:* Fell at 3.04 s.

(b) Press Reset, switch to Advanced, set Terminal cost to LQR on the MPC card and run to 8 s. Read the Cart travel tile. *Checkpoint:* 0.509 m, within 0.005 m.

(c) Press Reset, set Terminal cost back to Off, set Prediction horizon to 0.5 s and run to 8 s. Then press Reset, set the horizon to 0.6 s and run to 8 s. Read the Pole and Cart travel tiles each time. *Checkpoint:* 0.5 s: fell at 7.54 s. 0.6 s: up, with 0.793 m, within 0.005 m.

Compare the shortest horizon that holds the pole with the linear crossing of Section 7.5.6.

7.10 Application area · 4 min

A team plans to balance the default pole with a camera in place of the angle encoder. The camera delivers 30 frames per second, each image arrives 50 ms late, and the law runs once per frame.

(a) Estimate the average sensing lag $\tau_d + T_s/2$ and its fraction of the 179 ms e-folding time. *Checkpoint:* About 67 ms, or 0.37.

(b) Compare the result with case 8 and say whether the shipped gain of 30 N/rad can be kept. Name two effects of a real camera that the chapter's sensing model leaves out, and say which way each moves the conclusion. *Checkpoint:* Close to case 8's 70 ms, so the gain must come down.

7.11 Application area · 3 min

The bumpers return the cart with restitution $e = 0.4$, and through the pivot the impact changes the pole rate by $\Delta\omega = -(m\ell\cos\theta/J)\,\Delta v$ with $\Delta v = -(1+e)v$, as in Equation (7.3). For the default rod $m\ell/J = 3\,\mathrm{m}^{-1}$.

(a) Find the change in pole rate for a cart that reaches the bumper at 0.55 m/s with the pole near upright. *Checkpoint:* 132 °/s.

(b) Find the largest arrival speed that keeps the change below the 60 °/s tap of case 1, and the change at 0.55 m/s for a bumper with no rebound ($e = 0$). *Checkpoint:* 0.25 m/s and 95 °/s.

7.12 Review the agent · 5 min

Case 8 replaces the encoder angle with a camera-like reading every 20 ms that arrives 60 ms late, and the shipped cascade loses the pole at 6.84 s. An AI assistant proposes a fix: "The delay does not change the plant, whose unstable root stays at $5.58\,\mathrm{s}^{-1}$. The loop fails because a stale measurement makes it react too weakly. Raise the angle gain from 30 to 40 N/rad so that it pushes harder for the same reading, and keep the angle damping at 5 N s/rad." Find the flaw and say how you would check it.