INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 8

Double Inverted Pendulum Simulator for Model-Based Balance and Computed Swing-Up

Pages 188–216

Edited by John D. Hedengren

How to cite this chapter. "Double Inverted Pendulum Simulator for Model-Based Balance and Computed Swing-Up," 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. 8, pp. 188–216.

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

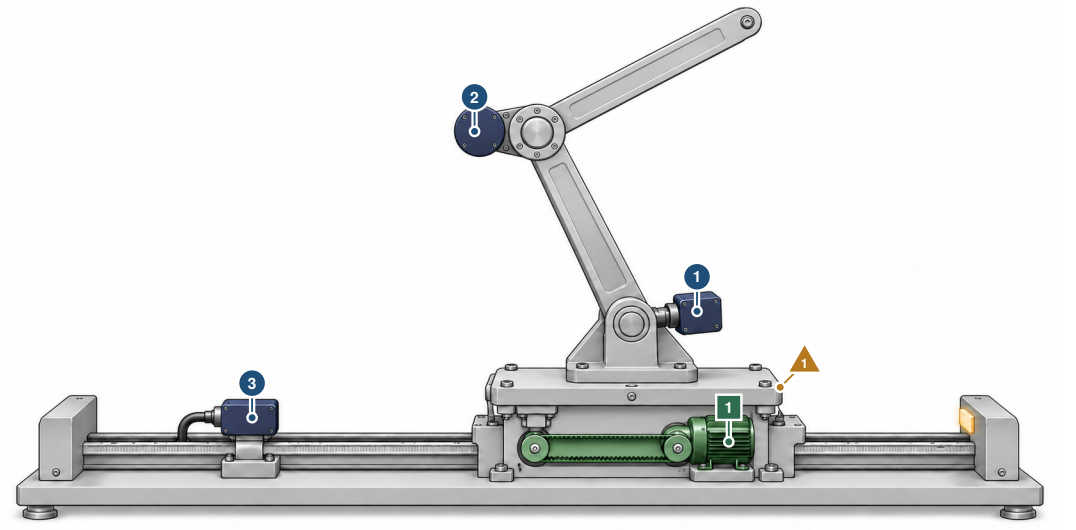
The app for this chapter
https://apmonitor.com/apps/double-inverted-pendulum/

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

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

CHAPTER 8

Double Inverted Pendulum Simulator for Model-Based Balance and Computed Swing-Up



The controller moves 1 motor and belt that push the cart

The controller reads 1 lower link's angle at its pivot 2 angle between the links at the elbow 3 where the cart is on the rail

What disturbs it 1 a push on the cart

Notice that one motor moving one cart must hold both hinged links upright while the cart stays between the bumpers at the ends of the rail. Illustration, not to scale.

IN PLAIN TERMS **Two Rods Stacked on One Moving Cart**

The main idea. A motor pushes a cart along a short rail. Two rods stand on the cart, the second balanced on the tip of the first, each on a free hinge. The controller keeps both rods upright by moving only the cart.

The challenge. The top rod folds back faster than a person can react. Pushing the cart toward the side each rod leans never works. The controller needs a model of the whole machine and a good guess of how fast each rod turns.

Why it matters. Standing robots, cranes with swinging loads and people keeping their balance face the same problem.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 8.2, the fact on the optimizer in Section 8.5, the try-its and the fact on sampling in Section 8.6, 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 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the open challenge

The control question Two rods stand on a cart that one motor pushes, and the faster of their two unstable modes is too quick for a hand to hold. What must a stabilizing law know and do that single loops cannot, and what fails when each of those requirements is taken away?

Learning objectives

- Show from the low-order coefficients of the closed-loop polynomial why no pair of angle loops that push toward both leans can hold the pendulum.
- Explain how differenced encoder counts make the force chatter, and why a Kalman filter on the same readings does not.
- Relate the longest stable sample interval to the drive lag the design leaves out, and compare it with the Nyquist interval.
- Describe how time-varying feedback along a stored swing-up absorbs a model error that the stored force alone cannot.
- Compare a constrained predictive controller with the regulator it is built from, and say where knowing the limits helps and where a linear plan hurts.

Prerequisites Sections 1.4.4 and 1.4.5 (state space, LQR), 1.4.6 (observer, Kalman filter), 1.4.7 (RHP poles and zeros, lag limits) and 1.6.3 (predictive control). Chapter 7 for the single pendulum on the same cart.

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

See also Volume 9, Chapter 4 for an offset-free predictive controller whose Kalman estimator carries a disturbance, and Volume 9, Chapter 2 for predictive control of a slow plant behind a dead time. Volume 5, Chapter 13 for an extended Kalman filter that estimates where an underactuated vehicle is when nothing measures its position.

8.1 Introduction

A double inverted pendulum on a cart was under computer control by 1978 [1], and in 1993, when its stabilization was already a familiar demonstration of state-space design, it was also swung up from hanging [2]. What makes the machine a hard test is that one motor must hold two unstable modes. The simulator of this chapter keeps the cart, rail, drive and encoders of the single-pendulum application of Chapter 7 and adds a 0.3 m, 0.10 kg rod on a free pivot at the tip of the 0.5 m pole. Linearized at upright, the rig has unstable roots at $+11.07\,\mathrm{s}^{-1}$ and $+4.49\,\mathrm{s}^{-1}$. The fast mode folds the upper link back against the lower one, $\theta_2 = -3.86\,\theta_1$, and multiplies by e every 90 ms. Released 0.5° from upright with the drive idle, the upper link passes 2° after 0.205 s, and the angle watchdog trips the drive at 0.615 s.

Three habits learned on single loops fail on this plant, and the simulator makes each failure measurable. The first is to give each link a loop of its own. Proportional-derivative loops that push the cart toward each link's lean, beside a loop that pulls the cart to its target, are never stable here, and two coefficients of the closed-loop polynomial say why: the two angle gains may sum to no more than $(M+m_1+m_2)g = 12.26\,\mathrm{N/rad}$, while the lower link's gain alone must exceed 58.6 N/rad. The second habit is to size the sample interval by the Nyquist frequency of the fastest motion. A regulator redesigned for 20 ms samples loses a 3.9 Hz oscillation, about a sixth of that sampler's 25 Hz Nyquist frequency, because its model leaves out the drive's 30 ms lag, and with that lag no interval longer than 12.3 ms is stable. The third is to trust a trajectory computed offline. Played without feedback, the stored swing-up reaches the balance regulator only on the exact rig it was computed for, while feedback along it absorbs an error in the upper link's mass from −27 % to +27 %.

These difficulties belong to a class of machines in which several modes share one actuator, so that no mode can be corrected without disturbing the others. People standing on a support shorter than their feet change from an ankle strategy, which moves the body about the ankle joints, to a hip strategy, whose motion is focused about the hips [3]. A model of that balance built from linearized mechanics, full-state feedback and state estimation reasonably matched experimental data once its objective took advantage of the fast modes that the body's inertia and geometry dictate [4]. A crane

whose payload swings as a double pendulum has two oscillation frequencies to suppress through the commands to its motors [5], where the single-pendulum crane of Chapter 12 has one. Maneuvers in which a humanoid robot gets up from the ground or recovers from a large disturbance acrobatically have been synthesized by online trajectory optimization [6]. Each of these needs commands that coordinate several modes at once, works from a state that is partly unmeasured, or plans its large motions instead of regulating them.

However, the experiments cited here document designs that succeeded on each rig, such as a digital servo that held one on an inclined rail [7], swing-ups by energy and by a limit cycle [8], and trajectories followed with feedforward and feedback [9, 10], and a failed alternative, such as a loop pair that cannot hold or a stored force played without feedback, is a fall that a physical rig is built to prevent. This chapter identifies, on a simulator described equation by equation, which requirements of full-state design this plant enforces and what goes wrong when each one is removed. The regulator, the Kalman filter, iterative LQR, time-varying tracking and constrained model predictive control are all standard methods from the literature.

Section 8.2 traces the history of the benchmark. Section 8.3 gives the nonlinear model, whose impacts pass through the mass matrix, and the elbow encoder that reads the angle between the links. Section 8.4 derives a closed form for the low-order coefficients of the closed loop under any single-input law, which excludes every loop pair that pushes toward both leans. Section 8.5 describes the controllers as coded: three estimators, among them a Kalman filter that predicts with the nonlinear model, the loop pair, the regulator, a stored swing-up that a repository script regenerates byte for byte, and a predictive controller built on the regulator's own model and told only what the regulator is not, the force the drive can deliver and where the rail ends (Section 8.5.6). Section 8.6 runs the ten cases and the challenge, each of which takes one requirement away or, in the tenth, adds one. It maps the longest stable sample interval against the drive lag that the design leaves out, checked against 11 nonlinear runs, and tests a linear rule that predicts the tilts the force limit can catch. Its comparison of the regulator with the predictive controller (Section 8.6.10) separates what planning within the force limit buys on a noisy estimate from what a linear plan costs near that limit. Section 8.7 describes the encoders, drives and real-time controllers of laboratory rigs, and Sections 8.8 and 8.9 state the limitations and the conclusions, and the exercises follow in Section 8.10.

8.2 A Benchmark Held, Then Swung Up

Computer control of a double inverted pendulum was reported in 1978 [1]. Two years later a digital controller designed from linear servo theory kept a double inverted pendulum on a cart at a given reference point while the cart stood on an inclined rail, removing the effect of the incline on the cart position, and its authors gave theoretical analyses of the pendulum and of the loop [7]. By 1993 the stabilization had become a standard demonstration of the state-space approach, and swinging the pendulum from its stable hanging state to the upright one, a transfer between equilibria, was posed as the next problem and shown to work in experiments [2].

Energy came next. A swing-up of the double pendulum built on an energy principle was followed by a more robust transfer through a limit cycle, and both were demonstrated on hardware [8]. Passivity and energy shaping were later used to drive a double pendulum on a cart from any initial position into a region of attraction around upright, where a balance controller designed on the linearized model takes over, a scheme shown in simulation [11]. The application keeps that division into a swing-up and a balance law joined by a handover, but it computes its swing-up instead of shaping energy.

Trajectories offered the other route. A reference trajectory obtained by optimizing a spline-interpolated initial guess was tracked with a gain-scheduled linear quadratic controller designed for that reference, and it swung a self-erecting double pendulum up in simulation and on a rig [9]. Feedforward and feedback together swung up a double pendulum on a cart with experimental validation [10], and a swing-up with experimental validation was later reported for a triple pendulum on a cart [12].

The optimizer used here improves a whole control sequence at once. A second-order method that successively improves such a sequence for a nonlinear discrete-time system, converging in one step when the system is linear and the cost quadratic, was derived in 1966 with an argument motivated by dynamic programming [13]. Iterative LQR, which repeats a linear quadratic approximation about the current trajectory, was introduced for locally optimal feedback control and computed energy-optimal reaching movements of a musculoskeletal arm with 10 states and 6 controls, converging substantially faster than three earlier methods [14]. Online trajectory optimization later synthesized getting-up and recovery maneuvers for humanoid robots at seven times slower than real time and solved the acrobot problem in real time [6]. A stored trajectory is only as useful as the feedback around it. Trajectories stabilized by LQR, with stability regions computed by sums-of-squares programming, have been assembled into a tree whose region of attraction probabilistically covers the controllable part of the state space [15]. The handover test of Section 8.5 is a far simpler and unverified relative of such a region.

> ⧖ HISTORY **The Funnel Picture of Feedback, 1985 to 2010**
>
> A robotics paper of 1985 popularized the funnel as a picture of a feedback policy: a law that collapses a large set of starting states into a smaller set of final ones. In 1999 feedback motion planning was drawn as a sequence of such funnels, each a locally valid policy, that carries a broad set of starting states to a goal. The weak point then was the mouth of each funnel, hard to compute and often estimated by trial and error. A Lyapunov function joins the two pictures, because a function that falls along every motion inside a ball proves that the ball drains to the goal. The cost-to-go of a regulator supplies such a candidate automatically, which is why a method of 2010 built its funnels around trajectories stabilized by LQR [15]. The handover test of this chapter is a level set of the same kind of cost-to-go.

8.3 The Cart and Its Two Links

The rig is the single-pendulum application's cart on a 2 m rail, with the same bumpers, drive and encoders, carrying a second uniform rod on a free pivot at the tip of the first, as shown in Figure 8.1. It is the upright cart of Section 1.4.1 with a second link on top and a force in place of a commanded acceleration. The cart position x counts from the middle of the rail, and both link angles are absolute, taken from upright and positive toward $+x$, so θ_2 is not the angle between the links. The eight states are x, $\dot{x}$, θ_1, $\dot{\theta}_1$, θ_2, $\dot{\theta}_2$, the force F that the drive has developed, and a latch for a lower link lying on its deck stop. Table 8.1 sorts the signals by role.

8.3.1 The Equations for Three Coordinates

The kinetic energy of the cart and both rods is $\frac{1}{2}\dot{q}^{\mathsf{T}}M(q)\dot{q}$ with $q=(x,\theta_1,\theta_2)$, and Lagrange's equations with viscous friction take the form of Equation (8.1).

$$M(q)\,\ddot{q}=\begin{bmatrix} F-b_c\dot{x}+h_2 s_1\dot{\theta}_1^2+h_3 s_2\dot{\theta}_2^2 \\ h_7 s_1-h_5 s_{12}\dot{\theta}_2^2-b_p\dot{\theta}_1+b_{p2}\dot{\theta}_{21} \\ h_8 s_2+h_5 s_{12}\dot{\theta}_1^2-b_{p2}\dot{\theta}_{21} \end{bmatrix} \tag{8.1}$$

where $s_i=\sin\theta_i$, $c_i=\cos\theta_i$, $s_{12}=\sin(\theta_1-\theta_2)$, $c_{12}=\cos(\theta_1-\theta_2)$ and $\dot{\theta}_{21}=\dot{\theta}_2-\dot{\theta}_1$, the elbow damping acting on the relative rate. The mass matrix and its constants are given in Equation (8.2).

$$\begin{gathered} M(q)=\begin{bmatrix} h_1 & h_2 c_1 & h_3 c_2 \\ h_2 c_1 & h_4 & h_5 c_{12} \\ h_3 c_2 & h_5 c_{12} & h_6 \end{bmatrix} \\ h_1=M+m_1+m_2,\quad h_2=l_1\left(\tfrac{m_1}{2}+m_2\right),\quad h_3=\tfrac{1}{2}m_2 l_2 \\ h_4=l_1^2\left(\tfrac{m_1}{3}+m_2\right),\quad h_5=\tfrac{1}{2}m_2 l_1 l_2,\quad h_6=\tfrac{1}{3}m_2 l_2^2 \\ h_7=g\,h_2,\qquad h_8=g\,h_3 \end{gathered} \tag{8.2}$$

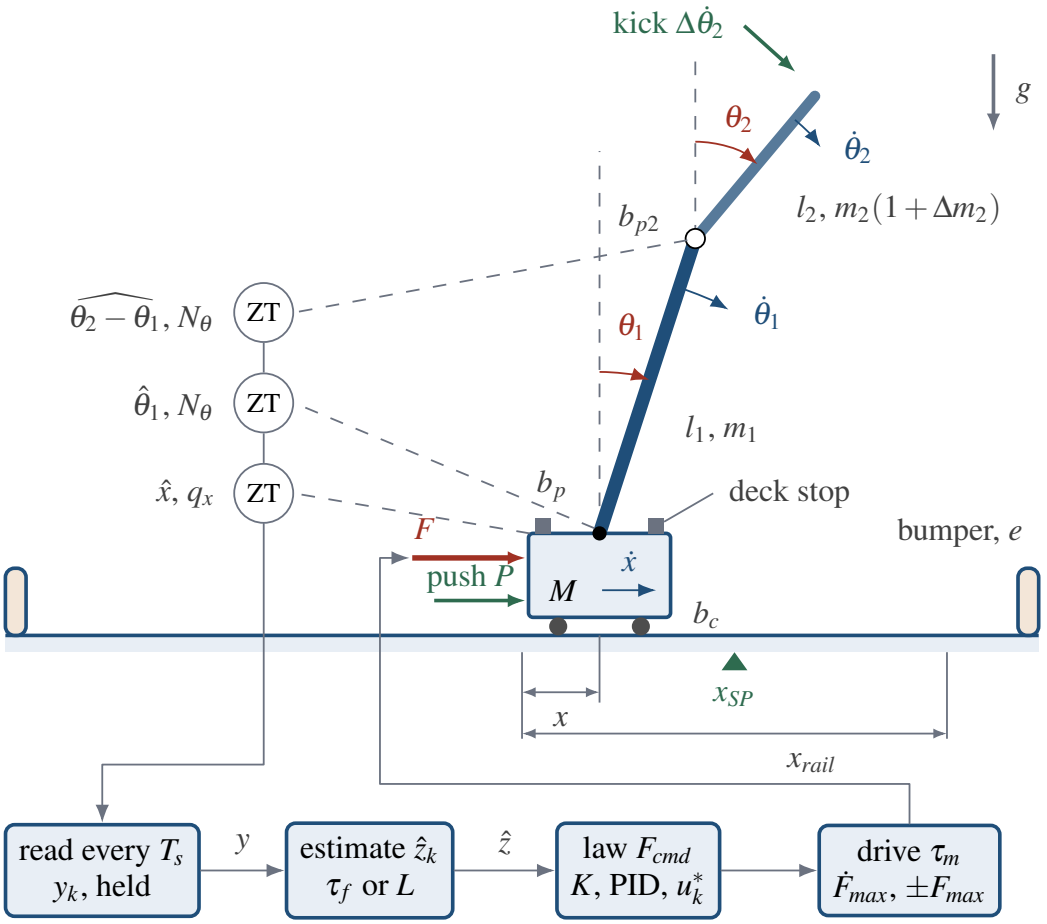


Figure 8.1 The cart and two links on a rail with bumpers, and the chain from the three encoders to the drive, not to scale. The cart of mass M on viscous friction b_c is pushed by the developed force F. The lower rod (l_1, m_1) turns on a pivot with damping b_p and the upper rod (l_2, m_2) on a free elbow with damping b_{p2}, and both angles are measured from upright, positive toward $+x$. Deck stops catch a lower link that passes 90°, and the cart center meets the bumpers at $\pm x_{rail}$ with restitution e. The encoders ZT read the cart to q_x, and the lower link and the angle between the links to N_θ counts per turn. Every T_s the readings y become an estimate $\hat z$, a law computes F_{cmd}, and the drive follows it with lag τ_m, slew limit $\dot F_{max}$ and limit $\pm F_{max}$. A push P strikes the cart, a kick $\Delta\dot\theta_2$ the upper link, and Δm_2 changes the upper link's mass.

Its top row balances the forces on the cart in newtons, and the lower two rows balance moments in newton meters about the cart pivot and the elbow. For the default rods $h_1 = 1.25\,\mathrm{kg}$, $h_2 = 0.0875\,\mathrm{kg\,m}$, $h_3 = 0.015\,\mathrm{kg\,m}$, $h_4 = 0.0375\,\mathrm{kg\,m^2}$, $h_5 = 0.0075\,\mathrm{kg\,m^2}$ and $h_6 = 0.003\,\mathrm{kg\,m^2}$. Because $M(q)$ is the matrix of a kinetic energy that is positive for every nonzero velocity, it is positive definite in every configuration, and the application solves the 3×3 system from its cofactors at each evaluation of the right-hand side. Its verification suite checks the definiteness over a 5° grid of both angles. Without force or friction the total energy $\frac{1}{2}\dot q^{\mathsf T} M \dot q + h_7 c_1 + h_8 c_2$ holds to one part in 10^6 of the swing-up energy $2(h_7 + h_8) = 2.011\,\mathrm{J}$ over 10 s, from three starts.

Table 8.1 Signals of the Rig by Role

Symbol	Role	Nominal	Range
θ_2, θ_1	CVs, link angles	0°	band ±2°
x	CV, cart position	0 m	SP ±0.6, band ±0.02
F_{cmd}	MV, drive command	0 N	±20
x, $\dot{x}$	states, cart	0	$\lvert x\rvert \le 1.0$ m
θ_i, $\dot{\theta}_i$	states, links	0	full circle
F	state, developed force	0 N	±20
—	state, deck latch	0	0 or ±1
$\hat{x}$, $\hat{\theta}_1$, $\hat{\phi}$	measurements	—	q_x, $2\pi/N_\theta$ steps
$\hat{z}$	estimate in charge	—	3 estimators
P	DV, push on the cart	0.5 N s	0 to 3
$\Delta\dot{\theta}_2$	DV, kick on link 2	30 °/s	0 to 300
Δm_2	DV, mass of link 2	0 %	−30 to +60
Δb_c	DV, rail friction	0 %	−100 to +300
$\theta_i(0)$	DV, release tilts	0°	±20°
—	DV, lost sample	—	event
λ_f, λ_s	outputs, roots	11.07, 4.49 s^{-1}	—
E	output, link energy	0 J	−2.011 J hanging

8.3.2 Drive, Rail, Deck Stop and Disturbances

The drive is a force source, and Equation (8.3) gives the lag and the rate bound through which its developed force follows the command.

$$\dot{F} = \mathrm{sat}\Big(\frac{F_{cmd} - F}{\tau_m},\ \dot{F}_{max}\Big), \qquad |F_{cmd}| \le F_{max} \tag{8.3}$$

where $\mathrm{sat}(v,c) = \max(-c, \min(c,v))$, $\tau_m = 30$ ms, $\dot{F}_{max} = 1000$ N/s and $F_{max} = 20$ N. Every impact is an impulse applied between integration steps through the inverse mass matrix, as shown in Equation (8.4).

$$\Delta\dot{q} = \frac{\Delta\dot{q}_j}{e_j^{\mathsf{T}} M(q)^{-1} e_j} M(q)^{-1} e_j \tag{8.4}$$

where e_j selects the coordinate struck and $\Delta\dot{q}_j$ is the change imposed on it. At a bumper, j is the cart and $\Delta\dot{x} = -(1+e)\dot{x}$ when the cart moves into the stop at more than $v_{rest} = 0.02$ m/s, with $e = 0.4$, and below that speed the restitution is zero. The same impulse reaches both links through the off-diagonal terms of $M(q)^{-1}$, which the verification suite checks against the closed form, and with $e = 1$ it conserves kinetic energy exactly. A lower link that passes 90° lands on a deck stop, $j = \theta_1$ and $\Delta\dot{\theta}_1 = -\dot{\theta}_1$, and from then on it moves with the cart while the upper link swings on from its tip, so the dynamics reduce to the two coordinates x and θ_2. The swing-up and the balance with a hanging lower link have no deck stop.

Three disturbances act on the state directly and one on the measurements. A push is an impulse P on the cart, $\Delta\dot{q} = P M(q)^{-1} e_1$. A kick adds $\Delta\dot{\theta}_2$ to the upper link's rate and leaves the cart and the lower link as they were. A release sets both links to chosen tilts at rest, and a dropout withholds one encoder sample. The mass change Δm_2 and the friction change Δb_c alter the rig that the physics runs, and only a present-links design hears of Δm_2. The energy of the links relative to upright rest, $E = \frac{1}{2}(h_4\dot{\theta}_1^2 + h_6\dot{\theta}_2^2) + h_5 c_{12}\dot{\theta}_1\dot{\theta}_2 + h_7(c_1 - 1) + h_8(c_2 - 1)$, is −2.011 J hanging at rest and zero balanced.

8.3.3 What the Controller Reads

The controller never reads the state unless the estimate is set to perfect, which the application labels a teaching comparison. When the law executes, the three encoders are read at that instant, as shown in

Table 8.2 Plant Constants. 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 to 3	a
b_c	cart friction	1.0 N s/m	0 to 8	a
l_1, m_1	lower rod	0.5 m, 0.15 kg	0.3 to 0.8, 0.05 to 0.5	s
l_2, m_2	upper rod	0.3 m, 0.10 kg	0.15 to 0.8, 0.02 to 0.4	s
b_p	pivot damping	0.002 N m s/rad	0 to 0.05	a
b_{p2}	elbow damping	0.0005 N m s/rad	0 to 0.01	a
g	gravity	9.81 m/s^2	1.6 to 24.8	a
F_{max}	drive force limit	20 N	5 to 40	a
τ_m	drive lag	30 ms	5 to 100	a
$\dot{F}_{max}$	drive slew limit	1000 N/s	100 to 5000	a
x_{rail}	bumper position	1.0 m	fixed	h
e	restitution	0.4	0 to 1	a
v_{rest}	resting-contact speed	0.02 m/s	fixed	h
q_x	cart encoder step	0.02 mm	0.001 to 2	a
N_θ	joint encoder counts	4096/turn	256 to 65536	a
λ_f, λ_s	unstable roots	11.07, 4.49 s^{-1}	derived	
z_1, z_2	zeros of x/F	10.95, 4.20 s^{-1}	derived	

Equation (8.5).

$$\hat{x} = q_x \,\mathrm{rnd}\Big(\frac{x}{q_x}\Big), \qquad \hat{\theta}_1 = q_\theta \,\mathrm{rnd}\Big(\frac{\theta_1}{q_\theta}\Big)$$
$$\hat{\phi} = q_\theta \,\mathrm{rnd}\Big(\frac{\theta_2 - \theta_1}{q_\theta}\Big), \qquad \hat{\theta}_2 = \hat{\theta}_1 + \hat{\phi} \tag{8.5}$$

where rnd rounds to the nearest integer, $q_x = 0.02\,\mathrm{mm}$ and $q_\theta = 2\pi/N_\theta$ with $N_\theta = 4096$, one count being 0.0879°. The elbow encoder reads the angle $\hat{\phi}$ between the links, so a count of either joint moves the controller's $\hat{\theta}_2$. The readings carry no noise and no transport delay, and nothing in the application is random, so a reset, which restarts from $t = 0$ with the settings in force, replays a run exactly, and the seed of Section 1.6.1 changes nothing in this app.

8.3.4 Constants and Integration

Table 8.2 lists the plant constants with their classes. The engine of Section 1.6.1 advances the eight states by the fourth-order Runge-Kutta method on a 1 ms step. The controller clock ticks on that same step and runs the law once per T_s, the log keeps 200 samples per simulated second, and a run can be played at speeds from 0.1 to 2 seconds of simulated time per second of real time. The bumper and deck-stop impulses follow each integration step, before the outputs are computed and logged. Halving the step moves the peak upper-link angle and the peak cart travel of the case 2 catch by less than 0.5 %, and the application passes all 52 of its verification checks.

8.3.5 Assumptions

Rigid, uniform rods. The mass matrix of Equation (8.2) holds for two rigid rods of uniform section. A long and light upper link would also bend, and the model has no way to show it. The elbow damping is a quarter of the cart pivot's because the upper rod's inertia about the elbow is only $h_6 = 0.003\,\mathrm{kg\,m^2}$, and the pivot's value placed at the elbow would by itself move the fast root from 11.37 to 10.42 s^{-1}.

Viscous friction. The cart and both pivots resist motion in proportion to speed. Stick-slip on a real rail makes a balanced cart hunt about its target in a small limit cycle, which nothing here can produce.

A lumped drive. One lag and one slew limit stand for the amplifier, the motor and the belt together. Backlash and a voltage limit would lengthen that lag, which Section 8.6 charges against the sampling margin.

Ideal impacts. A bumper contact takes zero time and has a single restitution coefficient, a lower link landing on the deck loses all of its rate, and the cart never lifts from the rail.

Encoders that only round. Each encoder errs by rounding to its count and in no other way, without noise, lost counts or delay. A camera, or noise on the counts, would spend part of the same lag margin.

Instantaneous disturbances on a level rail. Releases, kicks and pushes act in zero time. The rail is horizontal, which leaves the cart no steady force to resist, and gravity, the masses and the frictions change only when a disturbance changes them.

One stored trajectory. The swing-up was computed for the default rig, the default drive and 5 ms samples, and on any other rig it is a trajectory for a different machine.

8.4 What the Linearization Allows

Sections 1.4.4 and 1.4.7 set out the tools used here on a single pendulum. Three models are kept apart. The nonlinear Equations (8.1) to (8.5) are the simulation model. This section's linearizations are used for analysis and nowhere else, while the controllers of Section 8.5 carry design models of their own, built from the default links unless the reader chooses the present ones.

8.4.1 Two Modes and Two Wrong-Way Zeros

About upright rest, where $q = 0$, small motions obey $M_0\ddot{q} = Gq - D\dot{q} + e_1F$ with $M_0 = M(0)$, $G = \mathrm{diag}(0, h_7, h_8)$ and D the friction matrix. Its unstable roots are +11.37 s^{-1} and +4.57 s^{-1} without friction and +11.07 s^{-1} and +4.49 s^{-1} with it, one e-fold every 90 and 223 ms. The fast mode folds the links, $\theta_2 = -3.86\,\theta_1$, while the slow one leans them together, $\theta_2 = 1.50\,\theta_1$, and as $m_2 \to 0$ without friction one root becomes the single pendulum's 5.71 s^{-1}. The fast mode grows by a factor of 9.15 in 200 ms, 1.057 in one 5 ms sample and 1.248 in one of 20 ms.

Writing the linearized dynamics as $P(s)\,q(s) = e_1F(s)$ with $P(s) = M_0s^2 + Ds - G$, the transfers from force to the three coordinates share the denominator $\det P(s)$, and their numerators are the cofactors of the first row of $P(s)$, as shown in Equation (8.6).

$$\frac{x}{F} = \frac{C_{11}(s)}{\det P(s)}, \qquad \frac{\theta_1}{F} = \frac{C_{12}(s)}{\det P(s)}, \qquad \frac{\theta_2}{F} = \frac{C_{13}(s)}{\det P(s)} \tag{8.6}$$

The cofactor C_{11} is the characteristic polynomial of the two links with the cart held still, so the links' own unstable roots, +10.95 s^{-1} and +4.20 s^{-1}, are right-half-plane zeros of the force-to-cart transfer. Friction moves their stable partners to –4.24 s^{-1} and –11.52 s^{-1}, off the mirror image. One such zero makes a cart move start the wrong way, as Section 1.4.7 shows for the single pendulum. Two make a quick move of the cart start the right way, reverse, and reverse again before heading out, as cases 2 and 9 measure.

8.4.2 Why No Loop Pair Holds It

With the force applied without lag or sampling, any law $F = -Kz$ on the six mechanical states closes the loop on the characteristic polynomial of Equation (8.7), because the force enters only the first row of $P(s)$, so the closed-loop determinant expands along that row.

$$\Delta(s) = \det P(s) + (K_x + K_{\dot{x}}s)\,C_{11}(s) + (K_{\theta_1} + K_{\dot{\theta}_1}s)\,C_{12}(s) + (K_{\theta_2} + K_{\dot{\theta}_2}s)\,C_{13}(s) \tag{8.7}$$

A pair of angle loops that push the cart toward each link's lean, beside a cart loop that pulls it toward its target, is the law $F = K_{p1}\theta_1 + K_{d1}\dot{\theta}_1 + K_{p2}\theta_2 + K_{d2}\dot{\theta}_2 - K_{px}x - K_{dx}\dot{x}$, with every gain positive in the intuitive tuning. Without friction, C_{12} and C_{13} share their s^2 term gh_2h_3, because a slow acceleration

of the cart tilts both links alike, the pendulum leaning as one body in the accelerating frame. The two angle loops enter the lowest coefficients only through their sums, as shown in Equation (8.8).

$$\begin{aligned} a_2 &= g h_2 h_3 \left[(M + m_1 + m_2) g - K_{p1} - K_{p2} \right] - g \gamma K_{px} \\ a_3 &= -g h_2 h_3 \left(K_{d1} + K_{d2} \right) - g \gamma K_{dx} \\ a_4 &= \alpha K_{p1} - \beta K_{p2} + \delta K_{px} - a_4^0 \end{aligned} \tag{8.8}$$

where a_k multiplies s^k in $\Delta(s)$, $\gamma = h_3 h_4 + h_2 h_6$, $\alpha = h_2 h_6 - h_3 h_5$, $\beta = h_2 h_5 - h_3 h_4$, $\delta = h_4 h_6 - h_5^2$, all positive for uniform rods, and $a_4^0 = g[h_1 \gamma - h_2 h_3 (h_2 + h_3)]$ is the deficit of the open loop. A stable polynomial needs every coefficient positive. At the default rig $a_2 > 0$ caps the summed angle gain, $K_{p1} + K_{p2} < 12.26\,\mathrm{N/rad} - (0.629\,\mathrm{m}) K_{px}$, and $a_4 > 0$ demands $K_{p1} - 0.625 K_{p2} > 58.6\,\mathrm{N/rad} - (0.375\,\mathrm{m}) K_{px}$. Both hold only if $K_{p2} < -28.5\,\mathrm{N/rad} - (0.156\,\mathrm{m}) K_{px}$, a loop that pushes against the upper link's lean. The conflict is not peculiar to the default rig. With $K_{px} = 0$ the threshold exceeds the cap by $a_4^0/\alpha - (M + m_1 + m_2) g = g h_3 (h_1 h_4 - h_2^2 + h_1 h_5 - h_2 h_3)/\alpha$, and both differences in the bracket are positive for any uniform rods, while a cart gain only widens the gap because $\delta/\alpha < \gamma/(h_2 h_3)$. Over 20 000 random rigs the threshold was never less than 1.08 times the cap. The third coefficient asks independently for $K_{d1} + K_{d2} < -(0.629\,\mathrm{m}) K_{dx}$, so on balance the derivative action must also push against the links' rates. With the default friction the cap becomes 12.20 N/rad and the threshold 58.4 N/rad, and the largest a_3 that nonnegative gains reach while a_2 stays positive is -0.0081, still below zero.

The contrast with a single pole is the sign pattern of the open loop. With one unstable root λ, $\det P(s)$ is proportional to $s^2 (s^2 - \lambda^2)$, its s^2 coefficient is negative, and a loop that pushes toward the lean repairs it, which is the familiar floor of $(M + m_1) g$ on the angle gain that Chapter 7 derives for the single pole. With two, the polynomial is $s^2 (s^2 - \lambda_f^2)(s^2 - \lambda_s^2)$, the s^2 coefficient is already positive and the deficit sits at s^4, so pushing toward both leans only spends the s^2 coefficient. The shipped pair, 40 N/rad and 4 N s/rad on each link, sums to 80 N/rad. With its cart loop removed, on the frictionless model, it holds the lean-together motion as a single-pole loop would and leaves one real root at $+15.25\,\mathrm{s}^{-1}$ whose shape, $\theta_2 = -1.58\,\theta_1$, folds the links faster than the open loop's fast mode, and the default friction adds a second real root at $+0.14\,\mathrm{s}^{-1}$. The regulator of Section 8.5 has $K_{p1} = 153.6\,\mathrm{N/rad}$, $K_{p2} = -315.7\,\mathrm{N/rad}$ and $K_{d1} + K_{d2} = -52.7\,\mathrm{N\,s/rad}$ in these terms, and on the frictionless model its coefficients from a_0 to a_6 are 3.79, 4.00, 2.00, 0.422, 0.0455, 0.00233 and 5.86×10^{-5}, all positive. A search of 257 600 pairs that push toward both leans on the lag-free linear model finds none stable, the best leaving a root at $+9.12\,\mathrm{s}^{-1}$, and the application's own verification grid of 3125 pairs agrees.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.5 The Digital Controller

8.5.1 Sampling, Command and Watchdog

The law executes when a whole sample interval T_s has passed on the simulation clock, 5 ms by default and adjustable from 1 to 40 ms, and its force is held until the next execution. At each execution the encoders are read, all three estimators are updated, the law computes a force toward $+x$, and the command is clipped at $\pm F_{max}$, saturation being flagged whenever the unclipped command exceeds that limit. A watchdog trips the drive once either angle reading lies within one count of 90° from its target, and in the swing-up only after the handover. The settings are collected in Table 8.3.

8.5.2 Three Estimates of the State

The law needs $z = (x, \dot{x}, \theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2, F)$ and has three position readings, so the three estimators process the same samples in parallel and the trends show them together. The perfect estimate is the true state, sampled. The differenced estimate takes the positions as read, forms each rate as a filtered

Table 8.3 Controller Settings

Symbol	Meaning	Default	Range
T_s	sample interval	5 ms	1 to 40
τ_f	rate filter, differenced	5 ms	0 to 30
σ_F	Kalman force noise	2 N	0.1 to 20
σ_τ	Kalman joint torque noise	2 mN m	0 to 50
r	measurement variance, $\times q^2/12$	1	0.1 to 100
x_s	LQR cart excursion	0.3 m	0.05 to 1
θ_{1s}, θ_{2s}	LQR angle excursions	10°	1 to 45
F_s	LQR force scale	10 N	1 to 40
$\dot{x}_{ref}$	cart reference rate	0.25 m/s	0.05 to 2
K_{p1}, K_{p2}	link loop gains	40	−200 to 400
K_{d1}, K_{d2}	link loop damping	4	−20 to 40
K_{px}, K_{dx}	cart loop	5, 5	0 to 50
θ_h	handover ellipsoid	5°	1 to 20
t_0	earliest swing-up start	0.5 s	0.2 to 5
F_{push}	manual push	2 N	0.5 to 20
T_d	MPC decision period	10 ms	fixed
T_p	MPC horizon	1 s	fixed
δ	MPC rail margin	0.05 m	fixed

Loop gains in N/rad, N s/rad, N/m and N s/m.

backward difference, Equation (8.9), and replays the drive's lag on the commands sent to estimate F.

$$\hat{v}_k = \hat{v}_{k-1} + \eta\left(\frac{\hat{p}_k - \hat{p}_{k-1}}{T_s} - \hat{v}_{k-1}\right), \qquad \eta = 1 - \mathrm{e}^{-T_s/\tau_f} \tag{8.9}$$

where $\hat{p}$ is a reading and $\hat{v}$ its rate, and $\eta = 0.632$ at the defaults. One count differenced over 5 ms is a rate step of 17.6 °/s. Backward differences of encoder data respond well to velocity transients, while least-squares estimators filter imperfect measurements better [16], and the trade between chatter and lag appears in case 3.

The Kalman filter, whose filtering problem is the dual of the regulator problem [17], has the form of the observer of Section 1.4.6. It predicts the readings from the model and corrects the prediction by how far each reading misses. Its prediction runs the design model's full nonlinear equations over one sample, in Runge-Kutta steps of about 1 ms, with the command that was held, and its correction uses a steady-state gain, as shown in Equation (8.10).

$$\bar{z}_k = f(\hat{z}_{k-1}, u_{k-1}), \qquad \hat{z}_k = \bar{z}_k + L\,(y_k - H\bar{z}_k) \tag{8.10}$$

where $y = (\hat{x}, \hat{\theta}_1, \hat{\phi})$ and H maps z to $(x, \theta_1, \theta_2 - \theta_1)$. The gain L comes from the dual Riccati iteration of the shared module `shared/src/framework/core/riccati.js` on the seven-state model that includes the drive lag, sampled with a hold at T_s. Its process noise is a white force of $\sigma_F = 2$ N at the cart and white torques of $\sigma_\tau = 2$ mN m at each joint, entering through the columns of M_0^{-1}, and its measurement noise is each encoder's quantization variance $q^2/12$. The iteration converges in 232 sweeps, and the slowest error mode of the filter decays at $-7.6 \pm 6.3i\ \mathrm{s}^{-1}$. A lost sample leaves the readings as they were, so the differenced rates see a zero step and then a double one, while the Kalman filter only predicts.

8.5.3 The Loop Pair

The loop pair sums three copies of the shared velocity-form PID of Section 1.6.2, one on each estimated angle with the action reversed so that the loop pushes toward its link's lean, and one on the cart that pulls it toward the cart reference. Each derivative term differences the measurement after

smoothing it with a lag of τ_f, and each emitted output is clipped at $\pm F_{max}$. The integral is disabled in all three loops, so nothing is back-calculated. In manual mode the angle loops track their own proportional and derivative terms while the cart loop absorbs the rest of the operator's force in its integral state, so switching back to automatic causes no bump, and with no integral gain that state then stays as a fixed bias.

8.5.4 The Regulator

The regulator applies $F = -K(\hat{z} - z^*)$ with $z^* = (x_{ref}, 0, \theta_1^*, 0, \theta_2^*, 0, 0)$ holding the cart reference and the target angles. Its design model is the linearization at the target for the default links, without the drive lag (six states, F taken as F_{cmd}) unless the lag is modeled as a seventh state, discretized with a zero-order hold at T_s through one matrix exponential evaluated by scaling and squaring. Each weight is one over the square of the excursion the designer will accept in that quantity, the rule of Section 1.4.5, as shown in Equation (8.11).

$$Q = \mathrm{diag}\Big(\frac{1}{x_s^2}, 0, \frac{1}{\theta_{1s}^2}, 0, \frac{1}{\theta_{2s}^2}, 0\Big), \qquad R = \frac{1}{F_s^2} \tag{8.11}$$

The fixed-point Riccati iteration of the same shared module runs on the cost-to-go matrix S from $S = Q$ until the relative change falls below 10^{-11}, 1384 sweeps at the defaults, and gives $K = [30.0, 31.7, -153.6, 18.0, 315.7, 34.7]$ on $(x, \dot{x}, \theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2)$, with a zero appended on F. The force goes toward the lower link's lean and against the upper link's, at about twice the gain. The gain is recomputed whenever the sample interval, a weight, a plant constant or the design model changes. On the rig with its drive lag, the closed loop sampled at 5 ms has poles at $-1.56 \pm 1.23i$, $-1.94 \pm 27.48i$, $-5.62 \pm 4.73i$ and $-13.31\ \mathrm{s}^{-1}$, a largest magnitude in the z-plane of 0.992, which the Modes view draws with the eigenvalue routine of `shared/src/framework/core/eig.js`. The cart reference walks to the target at $\dot{x}_{ref}$, and the law carries no velocity reference, so a ramp is chased rather than tracked.

The advanced controls offer two other balance targets, and each has a single unstable root. Chosen in case 3 on the Kalman filter, either target starts the upper link 2° from it, and a 30 °/s kick on the upper link was fired at each whole second from 3 to 8 s, once that start had been caught. With the upper link hanging below an upright lower one the root is $+5.68\ \mathrm{s}^{-1}$, and the regulator designed there holds every kick within 2.65° on the lower link and 38.8 mm of cart travel. With the lower link hanging and the upper one standing on its tip the root is $+8.84\ \mathrm{s}^{-1}$, with $\theta_2 = 6.55\,\theta_1$, and the angle gains grow to 980.0 and 1129.1 N/rad in magnitude. The 2° start alone swings the links 21.7° and 16.4°. Five of the six kicks swing them up to 67.0° and 54.1° from the target and the cart up to 0.635 m, with both links back within 0.75° 10 to 12 s later, and the kick at 7 s tips them past 90°.

8.5.5 The Stored Swing-Up

The swing-up has four phases. While the links swing far from hanging, an energy sink asks for the cart acceleration $k_E(h_2 c_1 \dot{\theta}_1 + h_3 c_2 \dot{\theta}_2)$ with $k_E = 12$, which in the cart's frame drains the links' energy at the rate $k_E(h_2 c_1 \dot{\theta}_1 + h_3 c_2 \dot{\theta}_2)^2$, plus a centering spring at 1.2 rad/s with damping ratio 0.9. Inside 15° and 1 rad/s of hanging, and until the links leave 30° or 3 rad/s, an LQR about the hanging equilibrium settles them and parks the cart, both laws capped at 8 N. At the first sample after t_0 at which both links are within 2° and 0.1 rad/s of hanging at rest, and the cart within 1 cm and 0.02 m/s of its start, the law begins to follow the stored trajectory with time-varying feedback, as shown in Equation (8.12).

$$u_k = u_k^* - K_k(\hat{z}_k - z_k^*), \qquad (\hat{z} - z_{up})^{\mathsf{T}} S (\hat{z} - z_{up}) \le S_{55}\,\theta_h^2 \tag{8.12}$$

where z_k^* and u_k^* are the stored states and forces, K_k the stored gains, z_{up} is upright at rest, and the second relation is the handover test, a level set of the balance regulator's cost-to-go matrix S that reaches θ_h on the upper link's angle alone. From the first sample inside it the balance regulator takes

over, with its cart reference starting where the handover found the cart. A trajectory that ends outside the set trips the drive rather than fling the cart after a lost pendulum.

The trajectory comes from iterative LQR run offline by the repository script `double-inverted-pendulum/tools/swingup.mjs` on the application's own simulation model. One step of its dynamics, $z_{k+1} = f(z_k, u_k)$, is one 5 ms sample of Equations (8.1) to (8.3), integrated by the same Runge-Kutta steps of 1 ms, so replaying the stored force on the default rig reproduces the stored states to better than 10^{-6}. Over $N = 500$ samples, 2.5 s, from hanging at rest to upright at rest at the rail center, it minimizes Equation (8.13).

$$J = \sum_{k=0}^{N-1} \Big[\tfrac{1}{2} w_u u_k^2 + \tfrac{1}{2} q_w (\dot\theta_{1,k}^2 + \dot\theta_{2,k}^2) + \tfrac{1}{2} w_x \big(|x_k| - x_{max}\big)_+^2 + \tfrac{1}{2} w_s \big(|u_k - F_k| - \tau_m \dot F_{max}\big)_+^2 \Big] + \tfrac{1}{2} \|z_N - z_{up}\|_{Q_f}^2 \quad (8.13)$$

subject to $|u_k| \le 17\,\mathrm{N}$, which leaves the feedback a margin of the drive's 20 N. Here $(\cdot)_+$ is the positive part, $w_u = 10^{-3}$ per N^2, $q_w = 2 \times 10^{-3}$ per $(\mathrm{rad/s})^2$, $w_x = 2 \times 10^4$ per m^2 beyond $x_{max} = 0.6\,\mathrm{m}$, $w_s = 10$ per N^2 beyond the 30 N that the drive can slew within its lag, and $Q_f = \mathrm{diag}(2000, 200, 50000, 2000, 50000, 2000, 1)$. Each iteration takes finite-difference Jacobians of f along the nominal, runs a backward pass with a regularization that grows when a step fails and shrinks when it succeeds, clamps a step that would leave the force bound and gives it no feedback, and accepts the forward pass by a backtracking line search, whose form the script's comments credit to [6]. From a fixed guess, a 6 N sine at 0.8 Hz ramped in over 0.5 s, it stops after 446 iterations at a cost of 18.497, when an accepted step improves the cost by less than one part in 10^9, and the check mode regenerates the stored file byte for byte in about 6 s. The nominal force never exceeds 15.21 N, the cart reaches 0.600 m from center at 1.20 s, and both links first swing further from upright, to 216.7° and 230.5°, before they come up and are within 90° of upright from 1.51 s.

> **+ DID YOU KNOW? The First Race of the Swing-Up Optimizer**
>
> When iterative LQR was introduced in 2004, one of its three test problems was swinging up a simple pendulum. It converged in 0.26 s of computing and 5 iterations, against 1.61 s and 20 iterations for differential dynamic programming in the same environment [14]. The two links on a cart of this chapter took 446 iterations.

The gains K_k come from a final backward Riccati pass along the nominal with the balance weights, starting from the cost-to-go of the seven-state balance design, so they end on that design's gain, $[30.0, 31.3, -248.0, 14.0, 404.1, 43.4, 1.18]$. At the start they read $[30.1, 24.7, 113.2, 5.4, -21.1, 2.5, 0.7]$, and both angle gains change sign 1.43 s into the trajectory, as the nominal links pass 115.9 and 98.9° from upright, near horizontal. Because the shipped balance design leaves the drive lag out, the gain steps at the handover.

8.5.6 Model Predictive Control on the Regulator's Model

The fourth balance structure, the last value of the controller choice and the one that case 10 selects, is a constrained model predictive controller (MPC). It works in the receding-horizon manner of Section 1.6.3, whose stability and optimality under constraints are reviewed in [18]. The regulator asks for whatever force its gain gives and the drive clips the excess, and none of the laws above knows where the rail ends. The MPC is told both limits and nothing else that the regulator lacks. It acts where the regulator acts, at a balance target from a start near it, and the swing-up and its handover stay as described.

The prediction model is the regulator's design model. With $\tilde z_k = \hat z_k - z^*$ the estimate measured from the target and u_k the force command, that model is $\tilde z_{k+1} = \Phi \tilde z_k + \Gamma u_k$, the six-state linearization, or seven with the drive lag modeled, held over T_s. The MPC rebuilds Φ, Γ, K and S with the regulator's own design routines from the same design constants and weights Q and R, so its matrices equal

the regulator's bit for bit, and a change of sample interval, weight, plant constant or design model redesigns both laws on the same sample, on which the MPC also plans afresh. With the cart mass raised to 1.3 kg in the middle of case 10, the two gains change on the same law sample and agree bit for bit on all 599 law samples of the run. Both controllers are therefore given the present plant, as stated for the regulator above, while a disturbance that the design does not hear, such as Δb_c, is model error to both. Nothing is identified. The MPC reads the estimate selected for the regulator, the target, the force held since the last sample, the limit F_{max} as the drive clamp reads it and the fixed bumper position x_{rail}. The plant of Equations (8.1) to (8.3), the encoders and estimators of Equations (8.5) to (8.10), and this linear model remain three separate things.

The decision variables are departures from the regulator's force, a pre-stabilized parametrization in the manner of [19], and each plan predicts the regulator's own closed loop, as shown in Equation (8.14).

$$u_k = -K\tilde{z}_k + v_k, \qquad \tilde{z}_{k+1} = (\Phi - \Gamma K)\tilde{z}_k + \Gamma v_k \tag{8.14}$$

where v_k is the departure at step k, counted in samples of T_s from the present estimate $\tilde{z}_0$. The prediction thus starts from a stable loop instead of the open-loop roots at $+11.07\,\mathrm{s}^{-1}$ and $+4.49\,\mathrm{s}^{-1}$. Because K and S satisfy the Riccati equation of the regulator, completing the square in each stage cost and summing makes the terms in S telescope, and the regulator's infinite-horizon cost of any plan whose departures stop after N steps reduces to Equation (8.15).

$$\sum_{k=0}^{\infty} \left(\tilde{z}_k^{\mathsf{T}} Q \tilde{z}_k + R u_k^2\right) = \tilde{z}_0^{\mathsf{T}} S \tilde{z}_0 + W \sum_{k=0}^{N-1} v_k^2 \tag{8.15}$$

where $W = R + \Gamma^{\mathsf{T}} S \Gamma$, $0.012\,31\,\mathrm{N}^{-2}$ at the design of case 10 against $R = 0.01\,\mathrm{N}^{-2}$. Every step after the horizon is the regulator's own, so the cost-to-go S acts as the terminal weight and the sum of squared departures is the only part a plan can change. With no limit in reach the best departures vanish whatever the horizon or the blocking, and the MPC is the regulator. Within its limits the plan is a move-blocked, finite-horizon form of the constrained linear quadratic regulator [20].

The horizon holds $N = 200$ samples, 1 s at 5 ms and shorter at faster sampling, and the departures are held over 13 blocks, four of 2 samples, two of 4, and then 6, 8, 12, 18, 30, 50 and 60 samples. Each plan solves the quadratic program of Equation (8.16).

$$\begin{aligned} &\min_{v,\ \varepsilon \geq 0} \ W \sum_{k=0}^{N-1} v_k^2 + \rho\, \varepsilon^2 \\ &|u_k| \leq F_{max},\ k \in \mathscr{K}_F, \qquad |x_k| \leq x_{rail} - \delta + \varepsilon,\ k \in \mathscr{K}_x \end{aligned} \tag{8.16}$$

where u_k and the predicted cart position $x_k = x_{ref} + \tilde{z}_{1,k}$ follow Equation (8.14), $\mathscr{K}_F$ holds every sample of the first 0.1 s and every fourth sample after it, 65 of the 200, $\mathscr{K}_x$ every fourth sample, $\delta = 0.05\,\mathrm{m}$ is the margin inside each bumper and ε is one slack for the whole horizon. The force rows are hard. The rail band is soft, priced at $\rho = 10 N Q_{11} = 10 \times 200 \times 11.1\,\mathrm{m}^{-2} = 22\,222\,\mathrm{m}^{-2}$, ten times the regulator's own cart weight summed over the horizon. That rule replaced, once and openly, a first choice of a 0.1 m margin priced at 10^5 per m^2 and 10^3 per meter, which lost a kick that the regulator holds at a 40 N limit. The program has 14 variables and 230 inequality rows, 130 for the force and 100 for the rail. The dual active-set solver shared across the series handles it, with a ridge of 4.4×10^{-5} added to its diagonal. Between the judged samples a planned force is free, and in 98 of 3473 accepted plans from eight closed-loop runs one passes F_{max} by up to 17.5 N. None is applied, because the samples a plan applies before the next decision are always judged.

A plan is made every $T_d = 10\,\mathrm{ms}$, once in round(T_d/T_s) executions of the law and at least once per execution, which gives 9 ms at 3 ms samples and 7 ms at 7 ms, and the priming call at $t = 0$ makes the first plan. Between decisions the regulator keeps correcting from the fresh estimate with the planned departure for that sample added, $u_j = -K\tilde{z}_j + v_j$, so the departure is held and the force is

not. Holding the first planned force alone for 10 ms drops the pendulum at 1.79 s after a 30 °/s kick and at 1.82 s after a 0.4 N s push, with the cart target at +0.6 m, the Kalman filter and the drive lag left out of the design, while the regulator and the held departure both catch these disturbances, and with the lag modeled all three hold. The drive receives $F_{cmd} = \mathrm{sat}(u_j, F_{max})$, the same clamp as under the regulator, so the correction between decisions may still ask past the limit, as it did on 154 of 6989 balancing samples in the eight runs, never on a decision sample with the plan fully in force.

The MPC uses the estimate in charge and adds no disturbance estimator. Offset-free predictive control usually augments the model with an integrating disturbance, as the boiler of Volume 9, Chapter 4 does with a fuel bias in its Kalman estimator, but no law here integrates, the rail is level, and upright rest at the target with zero force is an equilibrium of the plant and of the model for any masses, while an estimator would break Equation (8.15) and the identity with the regulator. A plan is rejected if the solver does not report an optimum, if a departure, the slack or the cost is not finite, if the first force or any judged force lies outside $\pm F_{max}$, or if the plan has the wrong length. The regulator then takes the drive with the departure ramped out over 50 ms, counted in whole samples, the record is marked, and after five accepted plans in a row the departure is ramped back in. A state from which no plan meets every force row is rejected in the same way, and the verification suite forces each path of the rejection.

An independent condensation of the explicit objective, with the stage costs and S as the terminal weight, reproduces the application's plans on 3473 accepted decisions of the eight runs to 2.0×10^{-8} N with the same ridge. In the 3158 of them with no limit in reach every departure is exactly zero. A decision takes 0.255 ms on average and 0.334 ms at the 99th percentile in Node.js, most of it spent building the prediction. On the application's frame loop in the Chrome browser the whole controller uses 3.2 to 5.4 % of one core at real time and 6.0 to 7.4 % at twice real time over two sets of runs on a shared, loaded host, against a budget of a tenth of a core and 0.22 to 0.40 % for the regulator.

8.5.7 Manual Operation and the Models Kept Apart

Manual operation hands the operator a drive-force slider of ± 20 N in 0.5 N steps and pushes of F_{push} that last 0.05 s from a button, or while an arrow key is held down. While a law is in charge the slider tracks its command, so taking manual control starts from the force already applied. The estimators run on the commands actually sent in either mode, which leaves the regulator ready the moment the law resumes. Equations (8.1) to (8.5), run with the true links, form the simulation model. The design models are the regulator's linearization, the Kalman filter's nonlinear prediction with the drive lag, and the swing-up's simulation model of the default rig. The MPC predicts with the regulator's linearization and has no model of its own. None hears about Δb_c, and only a present-links design hears about Δm_2. The analysis models are the linearizations of Section 8.4 and the sampled closed loop of the true rig under the design's gain that the Modes view plots.

A balance target chosen before Run, or carried through a reset, starts the links at the case's tilt from that target, and the swing-up starts hanging at rest. Either way the case and event markers in the trends are kept.

8.6 The Cases, One Requirement at a Time

Table 8.4 lists the ten cases and the challenge, and Table 8.5 collects the outcomes. Every value was measured by running the application's engine headlessly on the named scenario, with the one change stated made before the run, which gives the same run as that change carried through a reset, and reading the trajectories and the scorecard the application computes. Each case starts from the reference settings of Table 8.3, with the regulator on the Kalman filter at 5 ms unless its row says otherwise.

8.6.1 No Hand and No Loop Pair

In case 1 the pendulum stands exactly at upright for 1 s, and then the upper link is let go 0.5° from upright with the drive at zero. It passes 2° after 0.205 s and 10° after 0.359 s, the watchdog trips the

Table 8.4 Scenario Sequence

#	Change from the reference	Requirement exposed
1	manual, loop pair in reserve, θ_2 released 0.5° at 1 s	a model, not a hand
2	perfect states, $\theta_2(0) = 2°$, target 0.5 m at 3 s, 2 m/s	signs from the model
3	differenced rates, $\theta_2(0) = 2°$	estimated rates
4	$\theta_2(0) = 1°$, T_s 5 to 20 ms at 3 s	margin for unmodeled lag
5	perfect states, released at 6° and 2.5°	a force budget
6	$l_2 = 0.6$ m, design for 0.3 m	a current model
7	swing-up from hanging	a computed trajectory
8	swing-up, $\Delta m_2 = 10\,\%$, no feedback	feedback along it
9	target 0 to 0.5 m at 1 s	reference shaping
10	MPC, lag modeled, target 0.6 m, 60 °/s kick at 1 s	the worth of known limits
C	swing-up, differenced, target 0.5 m	all at once

Table 8.5 Outcomes Read From the Running Engine

Case and condition	Result
1, no force after the release	θ_2 past 2° at 0.205 s, tripped at 0.615 s
1, loop pair / reversed pair / LQR	falls 0.530 s after / within 0.50° / within 0.638°
2, catch of the 2° start	cart left 27.0 mm, θ_1 to 4.17°, 11.82 N, in 0.1° from 2.46 s
2, target step at 3 s	right 5.9 mm, left to −51.3 mm, settled in 2.00 s
3, differenced / Kalman, 10 s	chatter 7.17 / 0.43 N, rate error 5.76 / 0.53 °/s
3, τ_f 0 / 10 / 15 ms	chatter 12.89 N / falls at 1.038 s / at 0.855 s
4, switch to 20 ms	falls at 6.47 s, pair $+1.64 \pm 24.28i$ s^{-1}
4, set after the switch, 12 / 13 / 15 ms	holds / falls at 14.96 s / at 7.95 s
4, lag modeled, 20 / 30 / 40 ms	holds at all three
5, released at (6°, 2.5°) / (6, −2.5°)	caught on 10.1 N, 117 mm / falls 0.576 s after
5, band at 10 / 20 / 40 N	8.4 / 26.1 / 52.4 % of the grid
6, design for 0.3 / 0.6 m	falls at 2.83 s / within 0.33°
6, l_2 0.4 / 0.5 m, old design	holds 10 s / falls at 9.01 s
7, handover at θ_h 2 / 5 / 10°	2.165 / 2.030 / 1.900 s into the trajectory
7, tracking error, Kalman / differenced	0.142° / 2.18°
8, feedforward alone / with feedback	misses, 225.6 times the level / hands over
8, rescued mass changes	−27 to +27 %, present-links design +20 to +60 %
9, $\dot{x}_{ref}$ 0.25 / 1 / 2 m/s	dip 7.1 / 27.6 / 50.8 mm, settled 3.23 / 2.14 / 2.01 s
10, LQR / MPC	both 0.114 m from the bumper, at the limit 25 / 15 ms
C, shipped / suggested	5.93 / 4.59 s, chatter 6.55 / 0.59 N

drive 0.615 s after the release, and the lower link lands on its deck stop at 2.097 s. At the 0.1 speed setting the same 0.615 s lasts 6.15 s of wall time, which gives a hand on the slider time to see the fold but no force to stop it. Handed to the loop pair in automatic operation, the release falls 0.530 s later on a peak command of 6.7 N, as Section 8.4.2 predicts. A pair with its upper loop reversed, 80 N/rad and 3 N s/rad on the lower link, −80 N/rad and −10 N s/rad on the upper and 1 N/m and 1 N s/m on the cart, satisfies all three conditions and holds the release within 0.50°, and the regulator on the Kalman filter holds it within 0.638° with 14.8 mm of cart travel.

> **TRY IT IN THE APP** **1 · Balance It by Hand**
>
> Load the scenario, leave it in Manual and press Run. Watch the upper link after it is let go at 1 s: it folds against the lower link, passes 2° after 0.205 s, and the watchdog trips the drive 0.615 s after the release. At the 0.1 speed setting the same fall lasts 6.15 s.

8.6.2 Signs From the Model

Case 2 runs the regulator on perfect states from a 2° tilt of the upper link. The cart's first move is to the left, 27.0 mm in 0.136 s, which swings the lower link 4.17° toward the upper link's lean and brings the elbow back under it. The peak force is 11.82 N, and both links are inside 0.1° from 2.46 s. At 3 s the target moves to 0.5 m with a reference walking at 2 m/s. The cart goes right 5.9 mm in 0.161 s, back to −51.3 mm at 0.460 s, and only then out to the target, inside 2 cm and 2° 2.00 s after the step, with the lower link leaning up to 6.51°. Two reversals where a single pole makes one are the two right-half-plane zeros of Equation (8.6) at work.

8.6.3 Rates From Three Encoders

Case 3 repeats the catch on the encoder readings. With differenced rates through the 5 ms filter, one count of the elbow encoder becomes a rate step of 0.307 rad/s, which the gain of 34.7 N s/rad on $\dot{\theta}_2$ turns into 10.6 N, 6.7 N after the first filter step. Over a 10 s run the command then changes by 7.17 N RMS from one sample to the next, the upper link wobbles ±0.64° at 4.7 Hz between 3 and 10 s with the force swinging ±9.10 N, and the estimate of $\dot{\theta}_2$ is wrong by 5.76 °/s RMS. On the same readings the Kalman filter gives 0.43 N of chatter, ±0.11° and 0.53 °/s. Filtering the differences harder is no escape. With a 10 ms filter the pendulum falls at 1.038 s and with 15 ms at 0.855 s, while with no filter the chatter reaches 12.89 N. A longer filter delays the rates the law uses, the same kind of lag that limits the sample interval in case 4.

Retuning either estimator moves it along a trade between chatter and lag, and the two trades do not overlap. Differenced through filters of 0, 2 and 5 ms, the chatter is 12.89, 11.23 and 7.17 N and the rate error 9.88, 8.66 and 5.76 °/s, and an 8 ms filter drops the pendulum at 1.190 s. The Kalman filter with joint torque noise of 0.5, 2, 5 and 10 mN m gives 0.35, 0.43, 0.64 and 1.09 N of chatter and 0.32, 0.53, 0.94 and 1.44 °/s of rate error, so the worst of those settings still improves on the best differenced one by a factor of 6.6 in chatter and 4.0 in rate error. A single lost sample at 4 s hardly shows in either. Between 3.9 and 4.2 s the largest change of the command from one sample to the next is 0.73 N on the Kalman filter, against 0.93 N in the same window without the loss, and 15.80 N on differenced rates, whose own chatter there reaches 15.92 N without the loss.

8.6.4 The Sampling Margin

Case 4 balances on the Kalman filter after a 1° start and at 3 s lengthens the sample interval from 5 to 20 ms, redesigning the regulator and the filter for the new interval at once. The pendulum falls at 6.47 s, 3.47 s after the switch. On the rig, the lightly damped pair at $-1.94 \pm 27.48i\,\mathrm{s}^{-1}$ crosses to $+1.64 \pm 24.28i\,\mathrm{s}^{-1}$, an oscillation of 3.9 Hz against the 25 Hz Nyquist frequency of the sampler, so the sampler still resolves the motion that escapes. With the interval set again just after the switch, 10 and 12 ms hold through 15 s, while 13 ms falls at 14.96 s and 15 ms at 7.95 s, and with the drive lag modeled the same loop holds after a switch to 20, 30 or 40 ms.

> **DID YOU KNOW?** **Fast Enough to See the Motion, Too Slow to Hold It**
>
> From the simulator: switched from 5 to 20 ms samples, the regulator loses the pendulum at 6.47 s on a growing pair at $+1.64 \pm 24.28i\,\mathrm{s}^{-1}$, a motion the sampler still resolves. With the drive lag in the design, the same loop holds at 20, 30 and 40 ms (Table 8.5).

The design is not blind to the hold, since it is redone for each interval on the zero-order-hold model, and the redesign helps, because the 5 ms gains kept at 20 ms reach a largest magnitude of

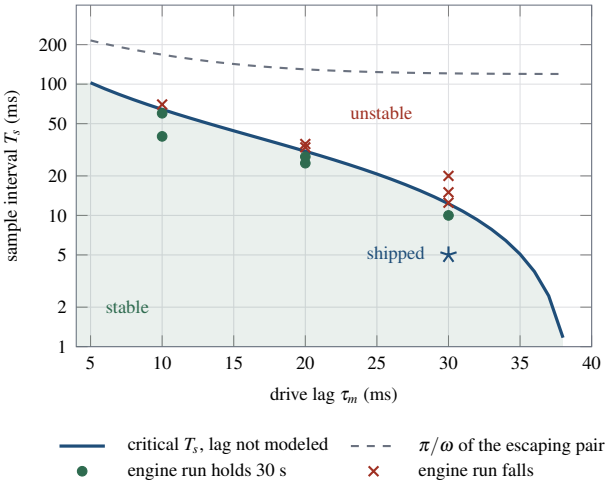


Figure 8.2 Longest stable sample interval of the regulator designed without the drive lag, against the drive lag, with the regulator redesigned at every interval. Below the solid curve the sampled loop on the rig is stable. The dashed curve is the Nyquist interval π/ω of the pair that escapes at the critical interval. The markers are case 4 runs on the nonlinear engine with the lag set before the run, the marked interval taking effect at the scripted switch at 3 s, and each run continued to 30 s, a filled circle for a run that holds and a cross for one that falls. The star marks the shipped setting. With the lag modeled, every interval up to 200 ms is stable.

1.053 and new ones 1.033. What it leaves out is the drive's lag, and the lag it can absorb shrinks as the interval grows. Section 1.4.7 bounds the lag that a loop around an unstable pole can carry on the single pendulum, and the map here adds the sample interval to that bound. Bisecting the largest pole magnitude of the sampled loop on the rig gives the critical interval as a function of the drive lag, as shown in Figure 8.2. At the default 30 ms lag it is 12.28 ms, the pair that escapes there oscillates at 26.0 rad/s, and its Nyquist interval would be 120.8 ms, ten times longer. The tolerable lag falls from 38.1 ms at 1 ms samples to 35.1 ms at 5 ms. Over drive lags from 15 to 35 ms the lag plus half the critical interval stays between 35.2 and 37.5 ms, so the hold spends the margin roughly as half a sample of extra lag would, although the design accounts for it, and the relation fails at short lags, where the sum reaches 56.3 ms at 5 ms of lag. The nonlinear engine agrees at all 11 points tested, run to 30 s because the growth near the boundary is slow. With a 10 ms lag the loop holds at 40 and 60 ms and falls at 70 ms, with 20 ms it holds at 25 and 28 ms and falls at 33 and 35 ms, and with 30 ms it holds at 10 ms and falls at 12.5, 15 and 20 ms. The two falls closest to the boundary, at 33 ms with a 20 ms lag and 12.5 ms with a 30 ms lag, come only at 10.35 and 10.72 s. With the lag in the design model, every interval from 1 to 200 ms is stable at every drive lag from 5 to 40 ms, on a grid of 1 ms in both.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.6.5 The Catchable Band

Case 5 releases the links at rest at 6° and 2.5° with perfect states. The regulator catches them on a 10.1 N peak with 117 mm of cart travel, and both links settle 1.62 s later, a catch that Figure 8.3 shows on the trends, while a release at 6° and −2.5°, fired at a sample instant once that catch has settled, falls 0.576 s after letting go, and 0.575 to 0.596 s after it over the five 1 ms phases of the sample. The Modes view bakes the full nonlinear catch on perfect states over a 41×41 grid of releases spanning $\pm 16°$ on the lower link and $\pm 8°$ on the upper, and a cell counts as caught when, over 2.5 s, neither link passes 90°, the cart stays off the bumpers and both links end within 2° of upright. At 20 N the

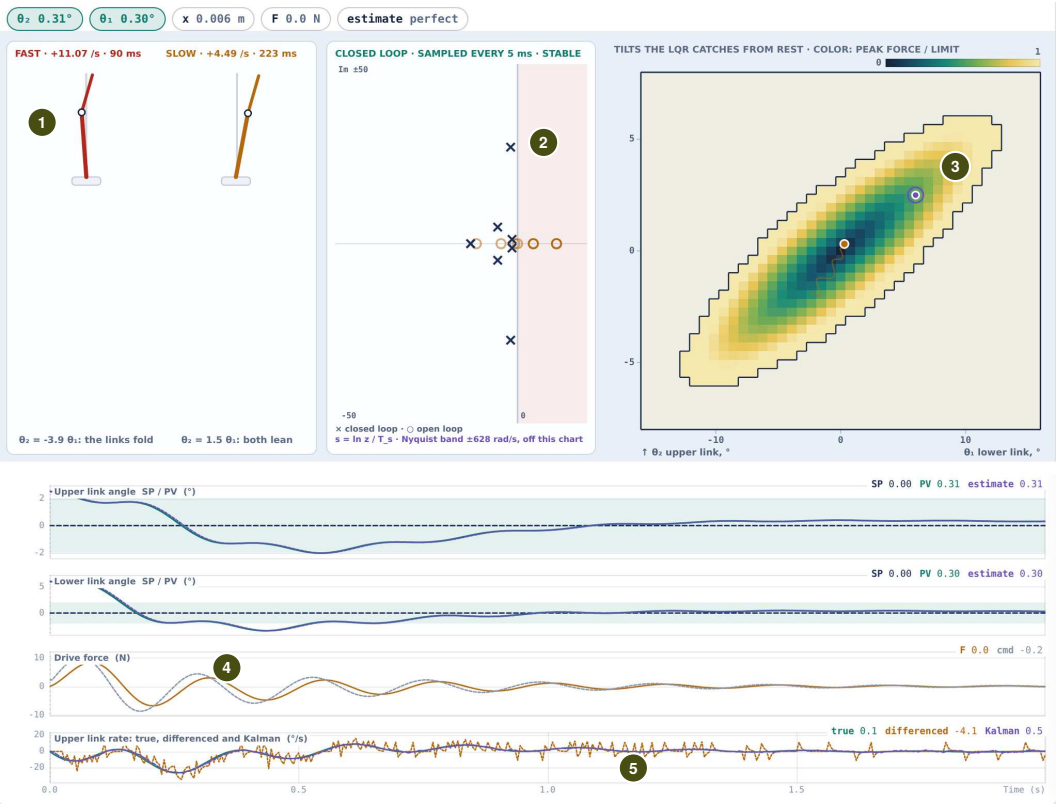


Figure 8.3 The application in case 5, The tilts it can catch, captured at $t = 2.00\,\text{s}$ on the Modes view, after a release at rest at 6° on the lower link and 2.5° on the upper with the regulator on perfect states. The figure is a composite of the badge row and the Modes view above the four trends, and the header, the control panel and the Learn panel with its scorecard are cut. (1) The left panel names the two unstable modes, the fast one at $+11.07\,\text{s}^{-1}$ that folds the links ($\theta_2 = -3.9\,\theta_1$) and the slow one at $+4.49\,\text{s}^{-1}$ that leans them together ($\theta_2 = 1.5\,\theta_1$), and animates each shape, here caught with the fold of the fast mode visible. (2) The middle panel maps the loop sampled every 5 ms into the s-plane, crosses for the closed loop and circles for the open-loop roots, with the lightly damped pair that the text places near ± 27.5 rad/s and the 628 rad/s Nyquist band far off the chart. (3) The right panel is the band of tilts caught from rest at the 20 N limit, colored by the peak force over the limit, with the release as a violet ring and the path of the live state back to the center. (4) Over the 2 s of trends the drive force rings down from its peak of 10.1 N (Table 8.5), while the upper link swings through −2.0° and the lower through −3.4° before settling. (5) The differenced estimate of the upper link's rate, plotted with the true and Kalman rates although perfect states are in charge, jumps each time a count arrives.

caught cells form a thin band through upright along $\theta_2 \approx 0.41\,\theta_1$ that covers 26.1 % of the grid and reaches ±6.4° on the lower link alone but only ±3.2° on the upper. At 40 N it covers 52.4 % and catches the second release on a 39.2 N peak, and at 10 N it shrinks to 8.4 %. That catch is marginal. Fired 1 ms after a sample, so that the drive holds the force computed before the release for 4 ms, the same release at 40 N falls 1.00 to 1.13 s later at six settled instants, while fired 2 to 4 ms after a sample it is caught with the command at the limit. About one release in five is lost, and the case's briefing tells the learner to reset and fire again.

The shape of the band follows from the linear loop. Released by one degree along a direction

in the (θ_1, θ_2) plane, the sampled loop with the drive lag asks for a peak force of 3.28 N along θ_1 alone, 7.34 N along θ_2 alone, and as little as 1.50 N along $\theta_2 = 0.43\,\theta_1$, so at 20 N its unsaturated reach is 6.09°, 2.72° and 13.4°. The farthest cell the nonlinear runs catch, at (−12.8°, −5.6°), lies 14.0° from upright along that direction. The band does not follow the direction that leaves the fast mode unexcited, which is $\theta_2 = 1.51\,\theta_1$. Applied to the grid, the rule that the linear peak force stay within the limit marks 331 cells at 20 N, 19.7 %, and the nonlinear runs catch every one of them plus 108 more that they recover through saturation. At 10 N the rule marks 83 cells, all caught. At 40 N it marks 961 and the nonlinear runs lose 80 of them, 72 of which are still lost with the slew limit removed, so the nonlinear terms of the largest tilts decide those cells. A linear rule sized by the force limit is a safe estimate of what a saturating drive can catch at the default limit, and an optimistic one at twice that limit.

8.6.6 A Model That Is Out of Date

Case 6 doubles the upper link to 0.6 m at the same 0.10 kg while the regulator and the filter keep the design for 0.3 m. The roots fall from 11.07 and 4.49 to 9.06 and 3.92 s^{-1}, an easier plant by its growth rates, and yet the old gains leave a closed-loop pair at $+4.46 \pm 20.68i\ \mathrm{s}^{-1}$, a 3.3 Hz oscillation that grows from the 0.5° start until the pendulum falls at 2.83 s. Redesigned for the present links, with 347.0 N/rad instead of 315.7 N/rad on the upper angle, it holds within 0.33° after the first second. At 0.4 m the old design is already unstable by a factor of 1.009 per sample but grows slowly enough to hold through a 10 s run, and at 0.5 m it falls at 9.01 s.

8.6.7 Swinging Up From Hanging

Case 7 starts hanging. The links are still from the start, so the trajectory begins at 0.50 s, and the handover comes at 2.53 s, 2.03 s into it, with the links at 13.9° and 18.0° but moving home at −85 °/s and −82 °/s, a state the cost-to-go counts as close. Both links are inside 2° and the cart within 2 cm of center from 4.15 s, the cart never passes 0.600 m, and the command peaks at 15.22 N. Played without its feedback on this rig, the stored force alone hands over at the same 2.03 s into the trajectory. With the Kalman filter in charge the links stay within 0.142° of the stored angles at every sample, against 2.18° on differenced rates. A 2° ellipsoid moves the handover to 2.165 s into the trajectory, at 4.1° and 8.7°, and a 10° one to 1.900 s, at 26.0° and 30.4°. On differenced rates the 2° ellipsoid is never entered and the swing-up misses.

> **TRY IT IN THE APP** **7 · Swing-Up Along the Stored Trajectory**
>
> Load the scenario and press Run on the Swing-up view. The links hang still until 0.50 s, then follow the stored trajectory, and the regulator takes over at 2.53 s with the links still 13.9° and 18.0° from upright but moving home. Press Reset, set Handover ellipsoid to 2° and run again: the handover moves to 2.165 s into the trajectory.

8.6.8 Feedback Along the Trajectory

Case 8 makes the upper link 10 % heavier, a payload the trajectory and the designs never hear of, and plays the stored force alone. The upper link goes over the top, and when the trajectory ends at 3.00 s it is at 286.5° with a cost-to-go 225.6 times the handover level, so the drive is tripped. Played alone the stored force hands over only on the exact rig, since of the mass changes from −5 % to +5 % in steps of one percent only zero succeeds. With the time-varying feedback the same +10 % run hands over 2.03 s into the trajectory on a 15.58 N peak, and the feedback rescues every change from −27 % to +27 %. Lighter than that the swing-up misses the ellipsoid. Heavier, it hands over but the balance design for the default links cannot hold the heavier link, and with the design for the present links every change tried from +20 % to +60 % is held. Balancing alone, in case 6 with the upper link set back to 0.3 m and its 0.5° start, the default design holds an upper link 25 % heavier for 10 s but drops one 30 % heavier at 7.46 s, although its sampled loop on that rig still has a largest magnitude of 0.999.

Table 8.6 Regulator, MPC, Governor and MPC Without the Rail Band, Runs of 10 s

	LQR				MPC					Governor	No rail band
Case and change	Fate	Hits	x_m (m)	t_{lim} (ms)	Fate	Hits	x_m (m)	t_{lim} (ms)	Plans	Fate	Fate
10 as shipped	held	0	0.114	25	held	0	0.114	15	1000/0	held	held
10, perfect states	held	0	0.168	90	held	0	0.164	20	1000/0	held	held
10, differenced rates	1.95	1	0	790	held	1	0	135	1000/1	1.95	held
10, 65 °/s	held	0	0.088	135	held	0	0.084	50	1000/0	held	held
10, 65 °/s, perfect states	held	0	0.114	255	2.27	12	0	600	228/1	held	2.27
10, 70 °/s	1.97	2	0	725	held	1	0	185	1000/1	1.97	held
10, 40 N, 80 °/s	held	0	0.028	0	held	0	0.039	0	1000/0	held	held
10, 40 N, 85 °/s	held	1	0	35	1.81	2	0	290	182/1	held	held
10, 40 N, 90 °/s	held	2	0	100	1.83	5	0	420	183/3	held	held
9, target 0.6 m, differenced, 0.4 N s push	held	0	0.363	510	1.94	1	0	405	194/1	held	1.94
3, down-up target	held	0	0.839	1029	1.82	0	0.658	1140	182/4	held	1.74

Fate is held, or the time in seconds at which the watchdog trips. Hits counts bumper contacts, x_m is the least distance of the true cart from a bumper, t_{lim} the time with the command at the drive limit, and Plans the MPC's accepted plans and rejections, the count stopping at a trip. The kicks act on the upper link at 1 s and the push on the cart at 1 s. The 40 N rows keep the drive lag modeled. Governor is the regulator with the reference governor, and No rail band the MPC without its rail rows.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.6.9 Moving the Cart

Case 9 moves the target from 0 to 0.5 m at 1 s while balancing on the Kalman filter, with the reference walking at 0.25 m/s. The cart edges right 0.83 mm, backs away 7.1 mm and settles inside 2 cm and 2° 3.23 s after the step, the links within 1.43°. At 1 m/s the dip is 27.6 mm and the settling 2.14 s with the lower link at 5.32°, and at 2 m/s the dip is 50.8 mm and the settling 2.01 s at 6.51°. Each doubling of the rate buys less time and costs a deeper dip.

8.6.10 The Regulator Against the Predictive Controller

Case 10 balances on the Kalman filter with the cart target at +0.6 m, 0.4 m from the right bumper, and the drive lag modeled in both designs, and at 1 s it kicks the upper link at 60 °/s toward $+x$, enough to hold the command at the drive limit. Table 8.6 compares the regulator and the MPC on that case, on its variants with the other estimators, larger kicks and a 40 N drive, and on the two other cases where the two laws part, and Figure 8.5 draws each run's time at the drive limit and its fate. Every run lasts 10 s from the named case, the stated change being applied before the start. The regulator keeps its shipped weights, and the MPC runs with the settings of case 10 throughout, with nothing retuned for a row. Both read the same encoders through the same estimator, and nothing in the application is random. Two further laws isolate the rail. The regulator with a reference governor [21] moves its cart reference at each sample to the value nearest the target for which its own unconstrained prediction keeps $|x| \leq 0.95$ m over 1 s, which gives the regulator the same rail knowledge by the classical route. The MPC with its rail rows removed is the counterfactual for the rail band. The loop pair falls in every row in which it was run, between 1.36 and 1.74 s.

The case as shipped is a tie. Both laws bring the cart to 0.114 m from the bumper, the MPC at 1.605 s and the regulator at 1.610 s, with the links at 13.8° and 6.3°, and they settle 2.18 s and 2.21 s after the start. The two carts never differ by more than 1.87 mm between 0.8 and 2.8 s, as shown in Figure 8.4. The difference lies in the force. At 1.042 s the regulator asks for −22.36 N and the drive delivers −20 N, while the MPC plans −20.00 N with a departure of +2.36 N, its largest. The first plan that the force limit shapes comes at 1.031 s and predicts the cart out to 0.72 m. The command sits at the limit for 15 ms under the MPC against 25 ms under the regulator, and no plan of the 1000 is rejected.

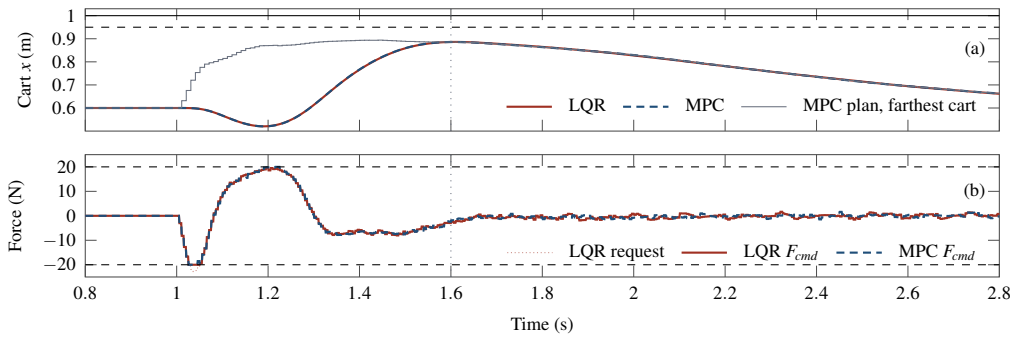


Figure 8.4 Case 10 run by the regulator and by the MPC on a common model with common weights, over a 2 s window from 0.8 to 2.8 s with the 60 °/s kick at 1 s, on every 1 ms step, the regulator as a thick solid line and the MPC as a thick dashed one. (a) Cart position against the bumper (thin horizontal line) and the MPC's soft limit at 0.95 m (thin dashed), with the farthest cart position predicted by the latest accepted plan as a thin stepped line. (b) The commanded force against the 20 N drive limit (thin dashed), with the regulator's unclipped request as a thin dotted line, which peaks at −22.78 N at 1.036 s. At the dotted line, $t = 1.60$ s, both carts are at 0.886 m, within 10 ms of their closest approach to the bumper.

Knowing the limit wins where the force is noisy or large. On differenced rates the regulator's chattering command sits at the limit for 790 ms and the pendulum falls at 1.95 s, while the MPC holds with one touch of the bumper, 135 ms at the limit and one rejected plan. With a 70 °/s kick on the Kalman filter the regulator falls at 1.97 s after two bumper contacts and the MPC holds after one. The same knowledge also costs, because the plan is linear. With perfect states and a 65 °/s kick the regulator holds with the lower link at 16.9°, and the MPC keeps the force at the limit longer than links leaning past 20° can take, reaches 21.2° on the lower link and falls at 2.27 s. The two runs part by 1° on the upper link at 1.523 s. In case 9 with the target at 0.6 m, differenced rates and a 0.4 N s push the MPC falls at 1.94 s where the regulator holds with 0.363 m to spare, the runs parting at 1.347 s with both links still under 6°. On the down-up target of case 3 the regulator holds within 21.8° and the MPC falls at 1.82 s after four rejected plans. At a 40 N limit the regulator survives the 85 and 90 °/s kicks with one and two slow contacts with the bumper, and the MPC falls at 1.81 and 1.83 s.

The rail band causes the two 40 N losses. Without its rail rows the MPC holds both kicks with one contact each, and it holds every case of the table that it holds with them. Avoiding a slow contact that the regulator survives costs a lean that the linear plan misjudges, and the band's only gain in the table is at 80 °/s, 0.039 m of rail left against 0.028 m. The governor changes the regulator's runs only at 80 and 85 °/s on the 40 N drive, leaving 0.035 m at 80 °/s, and it never ends a run worse than the regulator. The losses where the band plays no part, at 65 °/s on perfect states, after the differenced push and on the down-up target, also occur without it. A comparison built on a push that drives the regulator into the rail end is not available at the shipped limit. With the target at +0.6 m and the drive lag left out of the design, the regulator catches pushes up to 0.45 N s and kicks up to 30 °/s with at least 0.27 m of rail left, and every larger push or kick first holds the drive at 20 N and loses the links, the bumper being met only with a link 26.9° or more from upright or after the watchdog trip. The rail comes within reach of a kick that the drive can survive only at the 40 N limit.

Where no limit is in reach the laws do not part. The case 9 target steps at 0.25 and 2 m/s and from −0.6 m to +0.6 m at 2 m/s end alike on every estimator, the last on differenced rates falling under both at 3.27 and 3.21 s. On the ten existing scenarios the MPC is nowhere worse than the regulator. Among 49 changes a learner can make it does worse on seven, the five losses above, the down-up target with a kick, and a 24.8 m/s^2 gravity under which both fall, the MPC after four bumper contacts and the regulator after one, and better on three, the two wins and angle scales of 1°. The watchdog,

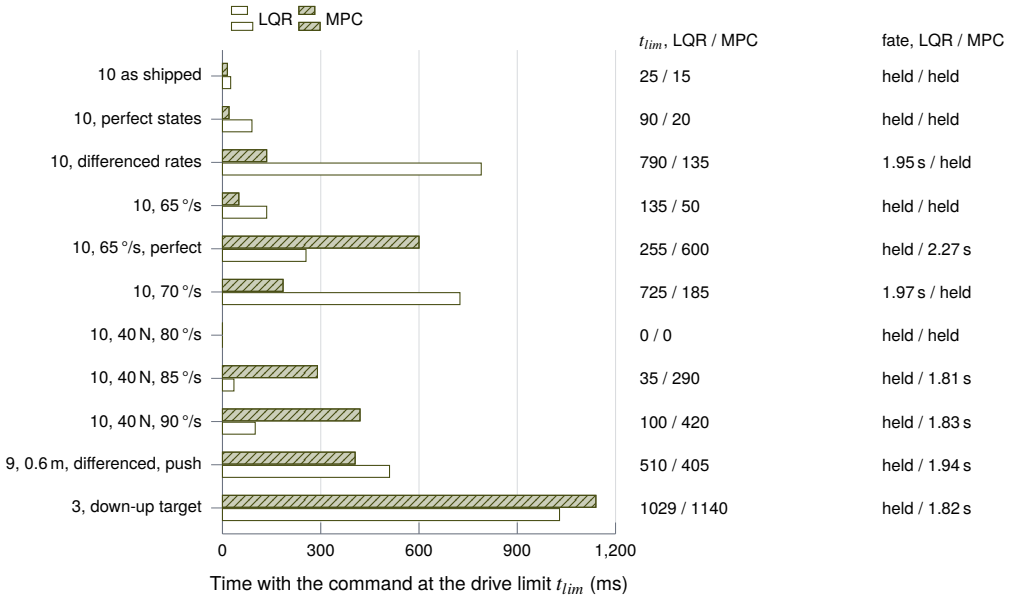


Figure 8.5 The regulator (white) and the predictive controller (hatched) of Table 8.6 on one model with one set of weights: the time each run spends with the command at the drive limit, and its fate, held or the time of the watchdog trip. Notice that the predictive controller holds both runs that the regulator loses, with less time at the limit, while in four of the five runs that it loses itself it stays at the limit longer than the regulator does.

the bumpers, the deck stop and the drive clamp act under every structure.

8.6.11 The Challenge

The challenge starts hanging with the target at +0.5 m and asks for both links inside $\pm 2°$ and the cart inside ± 2 cm of the target for good by 5.0 s of a 10 s run, with no bumper contact and a force chatter of at most 2.0 N RMS. As shipped, on differenced rates with the nominal weights, it settles at 5.93 s with 6.55 N of chatter. Each part of the target fails a different change. The Kalman filter alone quiets the force to 0.40 N but settles at 6.27 s, a 1 m/s reference alone settles at 4.72 s with 6.69 N of chatter, and a stiffer cart weight of $x_s = 0.15$ m on differenced rates drives the cart into a bumper. The suggested configuration, the Kalman filter with $x_s = 0.15$ m and a 1 m/s reference, settles at 4.59 s with 0.59 N of chatter and the cart never past 0.600 m. The shipped run is therefore 0.93 s late with 3.3 times the allowed chatter, and the suggested one meets the target with margins of 0.41 s and 1.41 N, whether its settings are made before the run or carried through a reset. As in other apps (Section 1.6.4), a change to a model parameter or a disturbance voids the attempt.

8.7 Technology in Practice

What Measures It

Laboratory pendulums on carts read their arm angles with encoders or with a motion-capture camera, and send the data over wires or a wireless link [22]. An encoder on a joint above the cart rides on a moving arm, so its signals must cross every joint below it. One open multi-link rig passes them through a slip ring, which keeps the latency low and needs no computing to recover the angle, and so suits high sampling rates [22]. The chapter's three encoders, on the cart, the lower pivot and the elbow, only round to their counts, with no noise, no lost count and no delay (Section 8.3).

> **TECHNOLOGY** **Sensor: Encoders That Ride on a Swinging Link**
>
> On an open multi-link rig each arm carries an optical encoder of 10 000 counts per revolution, and the cart's linear stage has a magnetic encoder with a resolution of 1 µm. An encoder above the first joint turns with the arm below it, so its signals leave through a slip ring, with the wires run through a hollow shaft. The slip ring adds no latency, unlike camera or wireless links, and needs no battery on the arm. Its brush contact adds friction at the joint, kept small with gold contacts that also cut electrical noise, and its five or eight channels leave little room to add a gyro [22]. The chapter's elbow encoder reads the angle between the links to 4096 counts a turn, with no friction beyond the elbow damping b_{p2}.

What Moves It

The same rig drives its cart with a linear motor whose stage carries its own encoder, and the servo drive closes a speed loop around it, so the controller's output is a commanded velocity rather than a force [22]. The chapter's drive takes a force command instead and develops it through one lag of 30 ms, a slew limit of 1000 N/s and a limit of 20 N, the lumped stand-in for the amplifier, the motor and the belt that its Assumptions name.

> **TECHNOLOGY** **Final Element: A Linear Motor in Place of a Belt**
>
> Carts on laboratory rigs are driven by a rotary motor, a servo motor with a belt, or a linear motor, and the choice affects how robust a control law proves. A linear motor avoids the backlash of a belt-driven cart, which makes a multi-link pendulum harder to control. On one open rig a servo drive in velocity mode reads the controller's analog voltage as a speed, compares it with the speed from the stage encoder and computes the motor current itself, over a stroke of 2 m [22]. The chapter's drive is a force source behind a 30 ms lag, and its Assumptions name the backlash that a real belt would add.

Where the Controller Runs

The open rig runs its law on a real-time target computer, with the design done on a separate host. Reading and sending its signals takes about 31.5 µs and a large block-diagram model about 56 µs, and in practice the rig samples at about 5 kHz while a time-varying LQR and a Kalman filter stabilize the double or triple pendulum. Its builders ask for at least 1 kHz for that stabilization [22]. The chapter's law runs every 5 ms, 200 Hz, on a 1 ms integration step, and Figure 8.2 shows why its unmodeled drive lag, not the sampler, sets the slowest safe rate. The state feedback $F = -K(\hat{z} - z^*)$ is the one block that holds the chapter's structure: the stored swing-up schedules its gain along the trajectory, and the predictive controller adds departures to its force. Figure 8.6 draws the loop on rig hardware.

What Is Modernizing It

The first development is the direct drive. A linear motor removes the belt and its backlash, and its stage encoder resolves 1 µm over the stroke [22]. The second is open hardware. The rig's drawings, tutorial code and data sets of the arms' angles and speeds, with and without control, are published so that system identification and machine learning methods can be tested on the same machine [22]. The third is the remote experiment. A user's controller would be checked first on a digital twin of the pendulum, sent back with a report if it broke a limit on position, speed or acceleration, and run on the hardware only after passing, with fail-safe limits that act outside the user's code [22].

> **IN PRACTICE** **Before the Cart Moves: Switches, Absorbers and the Power**
>
> On a documented open rig, limit switches at the ends of the travel tell the real-time controller to stop the motor. A shock absorber at the side of the rail takes up the cart's kinetic energy if the controller fails, and a protective panel guards people from parts thrown off the arm. Before a run the operator checks every wiring connection, and a run that does not go as planned is stopped at once with all power cut [22]. The chapter's watchdog and bumpers play the parts of the switch and the absorber.

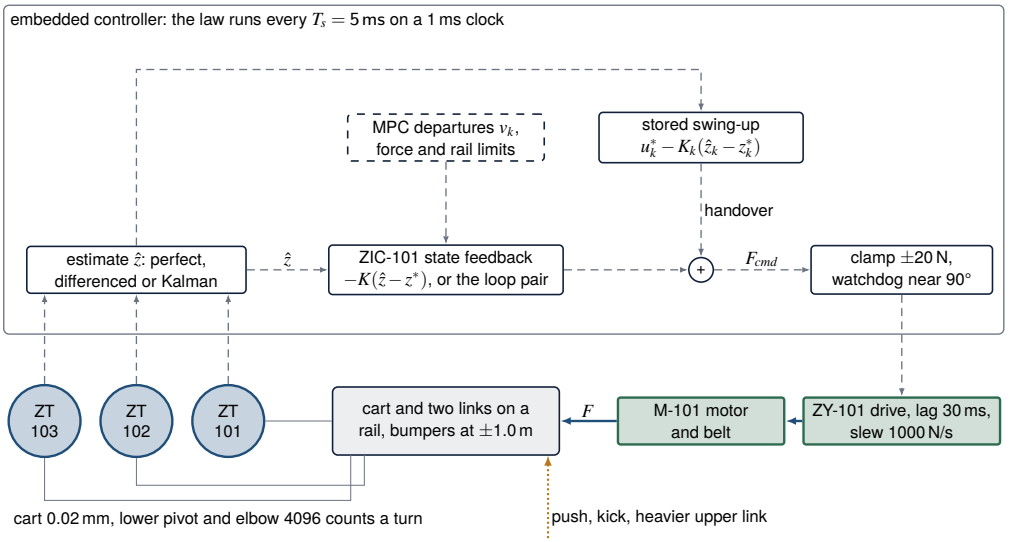


Figure 8.6 The balance loop on rig hardware. The encoders ZT-101 on the cart, ZT-102 at the lower pivot and ZT-103 at the elbow, which reads the angle between the links, feed three estimators that run in parallel. The state feedback ZIC-101 holds the regulator or the loop pair, the predictive controller (dashed) adds departures to the regulator's force, and in the swing-up the stored force with its time-varying feedback drives until the handover. The clamped command goes to the drive ZY-101 and the motor M-101. Notice that one force serves the cart and both links, so every block above the drive shares one limit of 20 N.

8.8 Limitations

Stick-slip hunting, the bending of a long upper link and sensor noise are all absent, because both rods are rigid, every friction is viscous and each encoder only rounds. A single lag and slew limit stand for the whole drive, and a real belt or rack would add backlash and voltage limits that shrink the margin of Figure 8.2.

The design models are perfect apart from the changes a scenario makes, and the present-links option redesigns instantly with the true links, which favors the regulator in cases 6 and 8 over any real redesign.

The analysis of Section 8.4.2 is a set of necessary conditions on the continuous, lag-free model, and the margin of Figure 8.2 is a linear stability boundary that says nothing about the size of disturbance a stable loop recovers. The handover set is a level set of a linear cost-to-go, not a verified region of attraction.

> ⚠ **OPEN CHALLENGE** **Proving a Handover Region for More Links**
>
> Verified feedback motion planning checks each funnel with sums-of-squares programs. Trajectory optimization and time-varying LQR scale well with the number of states, but this verification does not. It has succeeded on systems of up to five dimensions, with examples of six in the literature, and at high dimension the programs grow into very large semidefinite programs. Sampling can replace the proof at much lower cost, at the price of the guarantee [15]. The double pendulum on a cart already has six mechanical states, and the chapter's handover set is not a verified region of attraction.

The stored swing-up exists for one rig, one drive and one sample interval, and the application offers no re-solve in the browser.

The MPC plans on the regulator's linearization at the target and knows the drive limit and the rail, but not the nonlinear terms, the slew limit, the bumpers, the deck stop or the encoder quantization.

Near the edge of what the drive can do its plan can be confidently wrong, and it falls in five measured cases where the clipped regulator holds. Its force limit is judged on 65 of the 200 planned samples and its rail limit is soft, on one slack for the whole horizon. It has no disturbance estimator, which the level rail and the absence of a steady force on the cart make unnecessary here but a real rig with cogging or a tilted rail would need, and it is redesigned instantly from the present constants, as the regulator is.

8.9 Conclusion

A double inverted pendulum on a cart has two unstable modes and one motor, and the simulator of this chapter turns each requirement that follows into an experiment. The fast mode multiplies by e every 90 ms, so with nothing pushing, a pendulum released 0.5° from upright trips the drive 0.615 s later.

Loops that push toward both leans cannot hold it, because the angle gains may sum to no more than 12.26 N/rad while the lower link alone needs more than 58.6 N/rad, and the regulator that does hold it pushes against the upper link's lean at 315.7 N/rad. Differenced counts make its force chatter by 7.17 N where a Kalman filter on the same readings leaves 0.43 N. Its sample interval is limited by the lag it does not model, 12.3 ms at the drive's 30 ms, ten times shorter than the Nyquist interval of the motion that escapes.

A linear peak-force rule predicts the catchable band at 20 N without a false positive, and feedback along the stored swing-up absorbs upper-link mass errors from −27 % to +27 % where the stored force alone needs the exact rig. A constrained MPC on the regulator's own model ties the regulator in case 10, with 0.114 m of rail left under both, holds the differenced and 70 °/s kicks that the clipped regulator loses, and falls in five cases that the regulator holds, two of them because of its rail band.

The rods of the model are rigid and its encoders noiseless. A fuller model would add encoder noise and delay, backlash in the drive, a redesign from identified rather than true links, and a swing-up re-solved in the browser for the rig at hand. The simulator gives one plant on which every requirement of full-state design can be taken away and its failure measured.

KEY POINTS

- Two rods on one cart have two ways to fall, and the faster grows by e every 90 ms.
- Angle loops that push toward both leans cannot hold it: their gains may sum to no more than 12.26 N/rad, while the lower link alone needs more than 58.6 N/rad.
- Differenced counts make the force chatter by 7.17 N, a Kalman filter on the same readings by 0.43 N.
- The sample interval is limited by the lag the design leaves out, 12.3 ms at the drive's 30 ms.
- Feedback along the stored swing-up absorbs upper-link mass errors from −27 % to +27 %.

Where else this applies. The same cart and drive hold a single rod in Chapter 7.

References

[1] K. Furuta, T. Okutani, and H. Sone, "Computer control of a double inverted pendulum," *Computers & Electrical Engineering*, vol. 5, no. 1, pp. 67–84, 1978.

[2] M. Yamakita, K. Nonaka, and K. Furuta, "Swing up control of a double pendulum," in *1993 American Control Conference*. IEEE, 1993, pp. 2229–2233.

[3] F. B. Horak and L. M. Nashner, "Central programming of postural movements: adaptation to altered support-surface configurations," *Journal of Neurophysiology*, vol. 55, no. 6, pp. 1369–1381, 1986.

[4] A. D. Kuo, "An optimal control model for analyzing human postural balance," *IEEE Transactions on Biomedical Engineering*, vol. 42, no. 1, pp. 87–101, 1995.

[5] W. Singhose, D. Kim, and M. Kenison, "Input shaping control of double-pendulum bridge crane oscillations," *Journal of Dynamic Systems, Measurement, and Control*, vol. 130, no. 3, p. 034504, 2008.

[6] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 4906–4913.

[7] K. Furuta, H. Kajiwara, and K. Kosuge, "Digital control of a double inverted pendulum on an inclined rail," *International Journal of Control*, vol. 32, no. 5, pp. 907–924, 1980.

[8] M. Yamakita, M. Iwashiro, Y. Sugahara, and K. Furuta, "Robust swing up control of double pendulum," in *Proceedings of 1995 American Control Conference – ACC'95*, vol. 1, 1995, pp. 290–295.

[9] J. Rubí, Á. Rubio, and A. Avello, "Swing-up control problem for a self-erecting double inverted pendulum," *IEE Proceedings – Control Theory and Applications*, vol. 149, no. 2, pp. 169–175, 2002.

[10] K. Graichen, M. Treuer, and M. Zeitz, "Swing-up of the double pendulum on a cart by feedforward and feedback control with experimental validation," *Automatica*, vol. 43, no. 1, pp. 63–71, 2007.

[11] W. Zhong and H. Röck, "Energy and passivity based control of the double inverted pendulum on a cart," in *Proceedings of the 2001 IEEE International Conference on Control Applications (CCA'01)*. IEEE, 2001, pp. 896–901.

[12] T. Glück, A. Eder, and A. Kugi, "Swing-up control of a triple pendulum on a cart with experimental validation," *Automatica*, vol. 49, no. 3, pp. 801–808, 2013.

[13] D. Mayne, "A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems," *International Journal of Control*, vol. 3, no. 1, pp. 85–95, 1966.

[14] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *Proceedings of the First International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, vol. 1. SciTePress, 2004, pp. 222–229.

[15] R. Tedrake, I. R. Manchester, M. Tobenkin, and J. W. Roberts, "LQR-trees: Feedback motion planning via sums-of-squares verification," *The International Journal of Robotics Research*, vol. 29, no. 8, pp. 1038–1052, 2010.

[16] R. H. Brown, S. C. Schneider, and M. G. Mulligan, "Analysis of algorithms for velocity estimation from discrete position versus time data," *IEEE Transactions on Industrial Electronics*, vol. 39, no. 1, pp. 11–19, 1992.

[17] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[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] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.

[22] K. Kaheman, U. Fasel, J. J. Bramburger, B. Strom, J. N. Kutz, and S. L. Brunton, "The experimental multi-arm pendulum on a cart: A benchmark system for chaos, learning, and control," *HardwareX*, vol. 15, p. e00465, 2023.

8.10 Exercises

8.1 ○ Concept check · 2 min

The loop pair of case 1 puts 40 N/rad and 4 N s/rad on each link angle, each loop pushing the cart toward its own link's lean, beside a cart loop that pulls the cart toward its target. Why does no retuning with all gains positive hold the pendulum?

(a) Derivative action on both links cannot be made large enough to damp a mode that grows by e every 90 ms

(b) The 5 ms sample interval is too long for the fast mode, so every sampled pair of loops goes unstable

(c) The 20 N drive limit saturates before any tuning with positive gains can catch the fast mode

(d) The angle gains enter the s^2 coefficient only through their sum, which must stay below the gain the s^4 coefficient needs

8.2 Concept check · 2 min

In case 4 the sample interval goes from 5 to 20 ms at 3 s, and the regulator and the Kalman filter are redesigned for 20 ms on the zero-order-hold model at once. The pendulum still falls. Which explanation does the chapter support?

(a) The design leaves out the 30 ms drive lag, and the lag a design can absorb shrinks as the interval grows

(b) The 3.9 Hz oscillation that escapes lies above the Nyquist frequency of a 20 ms sampler

(c) The redesign ignores the zero-order hold, whose half sample of lag the new gains do not expect

(d) The redesigned Kalman filter converges too slowly at 20 ms to supply the rates the law needs

8.3 Concept check · 2 min

Case 8 plays the stored swing-up force alone, with no feedback along the trajectory, on a rig whose upper link is 10 % heavier, a payload of 10 g. Which statement explains why it misses the handover?

(a) The heavier link needs more force than the 17 N bound that the trajectory leaves below the drive limit

(b) The trajectory runs through unstable states, so a small model error grows unless feedback corrects it

(c) The swing-up reaches the top, but the balance design for the default links cannot hold the heavier link

(d) The trajectory starts before the links hang still, and that initial error rather than the mass decides

8.4 Theory and calculation · 7 min

The default rig has a cart of $M = 1.0\,\mathrm{kg}$, a lower rod with $l_1 = 0.5\,\mathrm{m}$ and $m_1 = 0.15\,\mathrm{kg}$, an upper rod with $l_2 = 0.3\,\mathrm{m}$ and $m_2 = 0.10\,\mathrm{kg}$, and $g = 9.81\,\mathrm{m/s^2}$.

(a) Evaluate h_1 to h_6 of Equation (8.2) and the energy $2(h_7 + h_8)$ that a swing-up from hanging at rest to upright at rest must supply. *Checkpoint:* $h_1 = 1.25\,\mathrm{kg}$, $h_2 = 0.0875\,\mathrm{kg\,m}$, $h_3 = 0.015\,\mathrm{kg\,m}$, $h_4 = 0.0375\,\mathrm{kg\,m^2}$, $h_5 = 0.0075\,\mathrm{kg\,m^2}$, $h_6 = 0.003\,\mathrm{kg\,m^2}$, and 2.011 J.

(b) With $K_{px} = 0$, find the cap on $K_{p1} + K_{p2}$ that $a_2 > 0$ sets and the threshold a_4^0/α that $a_4 > 0$ sets on the lower link's gain, with $\gamma = h_3 h_4 + h_2 h_6$, $\alpha = h_2 h_6 - h_3 h_5$ and $a_4^0 = g[h_1\gamma - h_2 h_3(h_2 + h_3)]$. *Checkpoint:* 12.26 N/rad and 58.6 N/rad.

(c) Check both conditions for the shipped pair, $K_{p1} = K_{p2} = 40\,\mathrm{N/rad}$, using $\beta/\alpha = 0.625$. *Checkpoint:* Both fail: the sum of 80 exceeds the cap by 67.7 N/rad, and $40 - 0.625 \times 40 = 15\,\mathrm{N/rad}$ is below 58.6 N/rad.

8.5 Theory and calculation · 6 min

The fast unstable root of the default rig is $\lambda_f = 11.07\,\mathrm{s^{-1}}$.

(a) Find the time in which the fast mode grows by a factor e, and the factor by which it grows over one sample of 5 ms, one of 20 ms, and 200 ms. *Checkpoint:* 90.3 ms, then 1.057, 1.248 and 9.15.

(b) After the switch to 20 ms in case 4 the lightly damped pair crosses to $+1.64 \pm 24.28i\,\mathrm{s^{-1}}$. Find its frequency in hertz and compare it with the Nyquist frequency of the 20 ms sampler. *Checkpoint:* 3.86 Hz against 25 Hz, 6.5 times lower.

(c) At the default 30 ms drive lag the critical interval is 12.28 ms, and the pair that escapes there oscillates at 26.0 rad/s. Find the Nyquist interval π/ω of that pair and the drive lag plus half the critical interval. *Checkpoint:* 120.8 ms, 9.8 times the critical interval, and 36.1 ms.

8.6 Theory and calculation · 5 min

The joint encoders count $N_\theta = 4096$ per turn, the law runs every $T_s = 5\,\text{ms}$, the differenced rate passes the filter of Equation (8.9) with $\tau_f = 5\,\text{ms}$, and the regulator's gain on $\dot{\theta}_2$ is 34.7 N s/rad.

(a) Find one count in degrees and the rate step that one count makes when it is differenced over one sample. *Checkpoint:* 0.0879°, and 17.6 °/s or 0.307 rad/s.

(b) Find η, and the step in the commanded force that one count of the elbow encoder causes after the first filter step and in full. *Checkpoint:* $\eta = 0.632$, with 6.7 N and 10.6 N.

(c) Find η for $\tau_f = 15\,\text{ms}$, and say why case 3 shows that a longer filter is no cure for the chatter. *Checkpoint:* $\eta = 0.283$. The pendulum falls 0.855 s into the run.

8.7 App challenge · 7 min

Load scenario 1, Balance it by hand, at speed 1 with the default seed. The upper link is let go 0.5° from upright at 1 s with the drive at zero.

(a) Press Run in Manual and touch nothing until the app pauses itself, about 1 s after the fall. Read the Pendulum tile, and find in the exported file the first row at which the upper link angle, column `th2_deg`, is beyond 2°. *Checkpoint:* Fell at 1.61 s, and the row at 1.205 s.

(b) Press Reset, switch to Auto so that the PID pair takes the drive, and run until the app pauses itself. Read the Pendulum tile. *Checkpoint:* Fell at 1.53 s.

(c) Press Reset, set Controller structure to LQR, keep Auto, and run to 5 s. Read the Cart travel tile, and find the largest $|$`th2_deg`$|$ after 1 s in the exported file. *Checkpoint:* 0.015 m and 0.64°, within 0.01°.

8.8 App challenge · 8 min

Load scenario 3, Encoders, differences and a Kalman filter. The upper link starts 2° from upright, and the regulator reads rates differenced from the encoder counts through the 5 ms rate filter. Use the default seed.

(a) Run as shipped to 10 s. Read the Force chatter and Rate error tiles. *Checkpoint:* 7.17 N and 5.8 °/s.

(b) Press Reset, set State estimate to Kalman, and run to 10 s again. *Checkpoint:* 0.43 N and 0.5 °/s.

(c) Press the reload button, switch to Advanced, set Rate filter to 15 ms with State estimate on Differenced, and run until the app pauses itself. Read the Pendulum tile. *Checkpoint:* Fell at 0.86 s.

Accept a chatter within 0.02 N. Explain the result of (c).

8.9 App challenge · 9 min

Load the Challenge, swing up, balance and move. Use the default seed and run each case for the full 10 s. Press the reload button before each part and make the part's settings before Run. After each run read the challenge result in the Learn panel, the settling instant and the force chatter.

(a) Run as shipped, on differenced rates with the nominal weights. *Checkpoint:* Settled at 5.93 s, chatter 6.55 N, not met.

(b) Set State estimate to Kalman and change nothing else. *Checkpoint:* 6.27 s, 0.40 N, not met.

(c) Set State estimate to Kalman and Cart reference rate to 1 m/s. *Checkpoint:* 4.81 s, 0.41 N, met.

(d) Set State estimate to Kalman and Cart excursion scale to 0.15 m, with the reference rate left at 0.25 m/s. *Checkpoint:* 6.00 s, 0.54 N, not met.

Say which part of the target each change addresses.

8.10 Application area · 4 min

A laboratory builds this rig with a belt drive. A step test shows that the belt's backlash and the amplifier's voltage limit stretch the effective drive lag from 30 to 35 ms. The regulator is designed without the lag and samples every 5 ms.

(a) Using Figure 8.2 and the results of case 4, say whether 5 ms samples are still safe. *Checkpoint:* No margin is left: the tolerable lag at 5 ms is 35.1 ms.

(b) Name two changes that restore a margin, and what to check after either one. *Checkpoint:* Model the lag in the design, or sample faster.

8.11 Application area · 3 min

A balancing robot built on this design lifts payloads that make its upper link up to 40 % heavier. It swings up along the one stored trajectory with the feedback on and balances with the design for the default links. Using the results of case 8, say which payloads it handles and what must be added for the rest.

8.12 Review the agent · 5 min

An AI assistant reviews case 10 and writes: "The MPC uses the regulator's own model, weights and Riccati solution, so with no limit in reach its force equals the regulator's, and near a limit it plans within the limit instead of being clipped. A controller that does everything the LQR does and also respects the limits cannot do worse. Make the MPC the default balance law on this rig." Find the flaw in the argument and say how to check it in the app.