INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 6

Magnetic Suspension Virtual Laboratory to Teach Control of Unstable Systems

Pages 131–157

Edited by John D. Hedengren

How to cite this chapter. "Magnetic Suspension Virtual Laboratory to Teach Control of Unstable Systems," 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. 6, pp. 131–157.

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

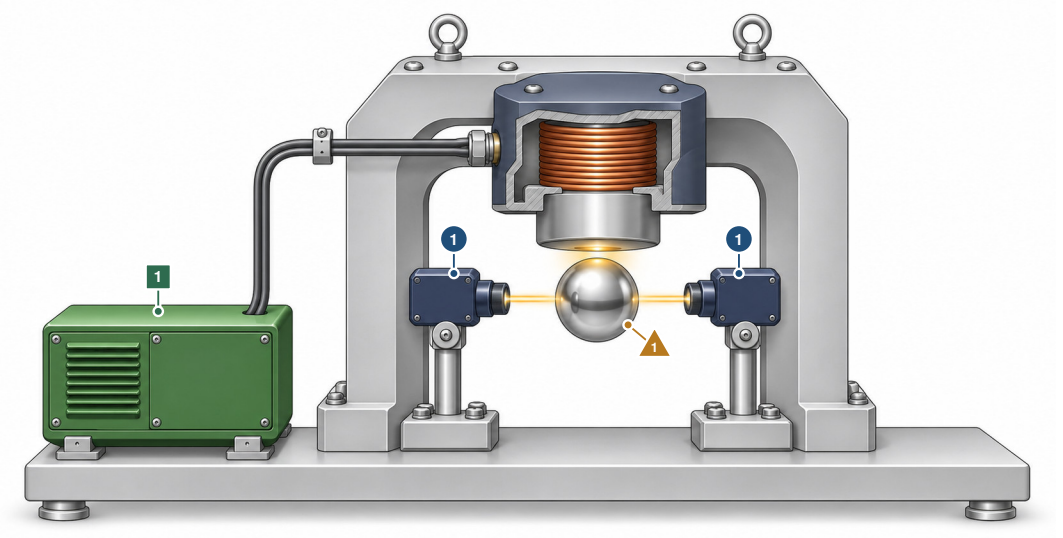
The app for this chapter
https://apmonitor.com/apps/magnetic-levitation/

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

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

CHAPTER 6

Magnetic Suspension Virtual Laboratory to Teach Control of Unstable Systems



The controller moves 1 current amplifier wired to the coil

The controller reads 1 infrared source and detector, a beam across the gap

What disturbs it 1 a heavier ball or a push away from the magnet

The controller resets the coil current a thousand times a second to hold the ball where the magnet's pull just balances its weight, a balance that does not last without feedback. Illustration, not to scale.

IN PLAIN TERMS A Steel Ball That Will Not Stay Put

The main idea. An electromagnet pulls a small steel ball upward against its weight. At one distance the pull and the weight balance, but if the ball drops a little the pull weakens and the ball falls faster.

The challenge. So a controller watches the ball with a beam of light and changes the current in the magnet a thousand times a second. It must react to how fast the ball moves, and it cannot ask for more current than the power supply can give.

Why it matters. Fast spinning machines and some trains float on the same kind of loop, with nothing touching and no oil.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 6.2, the fact on gaps in service in Section 6.3, the try-its in Sections 6.4 and 6.7, the fact on the derivative filter in Section 6.8, and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the gain limits and results (Sections 6.7 and 6.8), Technology in Practice (Section 6.10), Limitations (Section 6.11) and the open challenge

The control question A steel ball below an electromagnet has an equilibrium at every gap and none is stable, so even the exact balancing current loses it. Which terms of a PID hold the ball, how large must each be, and what bounds the gain from above?

Learning objectives

- Linearize the inverse-square force law about a gap and compute the unstable pole and the control effectiveness.
- Derive the three Routh conditions on K_c, τ_D and τ_I, and explain why proportional or integral action alone cannot hold the ball.
- Predict from the balancing current which gaps an amplifier ceiling can hold, and how much headroom is left for a disturbance.
- Explain how the derivative filter trades current chatter against the phase lead that holds the ball.
- Compare a PID and an explicit MPC on a step into the current ceiling and say what decides the outcome.

Prerequisites Sections 1.4.3 (linearization), 1.4.7 (lag limit, actuator headroom), 1.4.5 and 1.4.6 (LQR, Kalman filter), and Sections 1.6.1, 1.6.2 and 1.6.3 (engine, velocity-form PID, MPC option). Chapter 5 for the same lag limit on a servo.

The app https://apmonitor.com/apps/magnetic-levitation/

See also Volume 2, Chapter 2 for the velocity-form PID and its anti-windup, Volume 2, Chapter 3 for a setpoint outside the actuator range, Volume 2, Chapter 6 for one tuning across a nonlinear gain, and Volume 2, Chapter 12 for an explicit MPC on a machine axis. Volume 9, Chapter 4 for offset-free predictive control with an estimated disturbance.

6.1 Introduction

Almost every simulation used to teach feedback describes a plant that would come to rest somewhere definite if the controller were switched off. Tanks drain to a level, reactors settle at a temperature, exchangers reach a duty, and the loop then improves a condition the process already reaches on its own, which makes a first course in control a course about performance. A magnetic suspension has no such resting state. The upward force on a ferromagnetic ball falls as the square of the gap, so at every distance there is a current that balances the weight exactly, and at that current a small displacement away from the magnet weakens the force that was holding the ball while a small displacement toward the magnet strengthens it. The equilibrium is a maximum of the potential rather than a minimum. Setting the current to eleven correct digits and releasing the ball puts it either on the pole face or on the floor.

Hardware suspension rigs make the argument well and have been documented as undergraduate projects for four decades [1, 2]. A rig is also a power amplifier, a photodetector and a machined pole face, and the time available in a single laboratory period is usually spent assembling it rather than measuring it. However, the published rigs and the published control designs both begin after the instructive failure has been eliminated, so there is no inexpensive place to observe an unstable plant diverge. This chapter presents a simulated magnetic suspension in which manual control is attempted and fails first, and derives from the implemented loop a closed-form condition for when feedback succeeds.

The model rests on standard mechanics. Newton's second law on a single translational degree of freedom supplies one state equation, the reluctance force of a gapped magnetic circuit supplies the driving force as $\alpha i^2/h^2$, and the coil inductance divided by its resistance supplies a first-order actuator lag, so the whole plant is three states and one input [3]. Linearizing an attractive force that depends on the reciprocal square of a displacement produces a positive position stiffness rather than a negative one, which is the origin of the right-half-plane pole and the reason a controller has to supply the stiffness the plant does not have [4, 5]. The scenarios built on that model teach the distinction between an equilibrium and a stable equilibrium, linearization about an operating point and the meaning of an open-loop pole, the separate jobs of the proportional, integral and derivative terms, Routh stability

testing of a fourth-order characteristic polynomial, operating-point dependence of process gain, actuator saturation and anti-windup, and the amplification of sensor noise by derivative action [6, 7]. Industry runs the same loop wherever contact is removed on purpose, in the active magnetic bearings that hold turbomolecular pumps and gas compressors and in the suspension of a guideway vehicle, and in every one of those machines the feedback loop is a load-bearing member [8, 9].

The analysis gives a stability certificate for the loop as implemented. Three Routh inequalities fix the minimum proportional gain, the minimum derivative time and the maximum integral action in terms of the physical constants, and they agree with the nonlinear simulator to within one setting increment. The model is a three-state suspension covering forces, amplifier, sensor, hard stops and numerics, specified closely enough that a second implementation reproduces the same trajectories. Its scenarios make the two mechanisms that bound the gain from above switchable one at a time, so each can be met on its own. An explicit model predictive controller beside the PID adds a measured comparison on a step into the current ceiling, in which runs that remove the ceiling, the proportional kick or the look-ahead one at a time show what each of them decides. Suspension modeling, PID tuning and explicit predictive control are all mature, and the chapter uses each of them as it stands.

Section 6.2 traces the problem from the impossibility of a static suspension to magnetic bearings and guideway vehicles. Section 6.3 states every relation and constant required to rebuild the suspension from this text alone. Section 6.4 linearizes it about a gap and shows why manual operation at the exact balancing current fails. Section 6.5 describes the PID and the explicit MPC, Section 6.6 derives the three stability conditions and tests them against the simulator, and Section 6.7 meets the amplifier ceiling and the sensor noise that bound the gain from above. Section 6.8 sets out the scenarios, the graded challenge and the comparison of the MPC with the PID, and Section 6.9 gives the numerics, and Section 6.10 describes the sensors, amplifiers and controllers of real suspensions. Sections 6.11 to 6.13 state the limitations, the wider lesson and the conclusions, and the exercises of Section 6.14 repeat the measurements on the app.

6.2 Background

The impossibility result predates the engineering. A static arrangement of fixed magnets and ferromagnetic bodies cannot produce a stable suspension, and Braunbek sharpened that statement in 1939 by separating the materials that evade it from the materials that cannot, so diamagnets and superconductors float passively while ordinary ferromagnetics do not [10]. The practical consequence is direct. A steel ball can only be held if something varies the field, and that something is a feedback controller.

> ⧖ HISTORY **A Servo Holds a Rotor in the Air, From the Late 1930s**
>
> A paper of 1937 showed that a vertical ferromagnetic needle can be held in equilibrium by magnetic forces alone if the field is varied, and a 6 g rotor held that way spun at about 1200 revolutions per second [11]. The electromagnetic support developed in the late 1930s was a one-dimensional system: a servo controlled one direction, and only the shape of the field held the others, the single axis this chapter keeps. By 1954 that support had been developed for more than a decade and a half, and an electronically controlled solenoid held a high-speed rotor in vacuum as a frictionless bearing. In 1957 a French aeronautical research office first showed that a wind tunnel model can be held by magnetic suspension, with no support to disturb the flow [11].

Industry adopted the result. Active magnetic bearings replaced oil films in turbomolecular pumps, machine-tool spindles, flywheels and gas compressors, trading contact and lubrication for a control system that has to compute continuously for the machine to stay assembled. The design literature treats the levitation loop, its power stage and the rotordynamics it supports as a single artifact [8]. Transport adopted it in parallel. Electromagnetic-suspension vehicles regulate the same attractive force across a gap of a few millimeters beneath a guideway rail, while electrodynamic suspension becomes passively stable above a threshold speed. Both families are surveyed together [9], and an

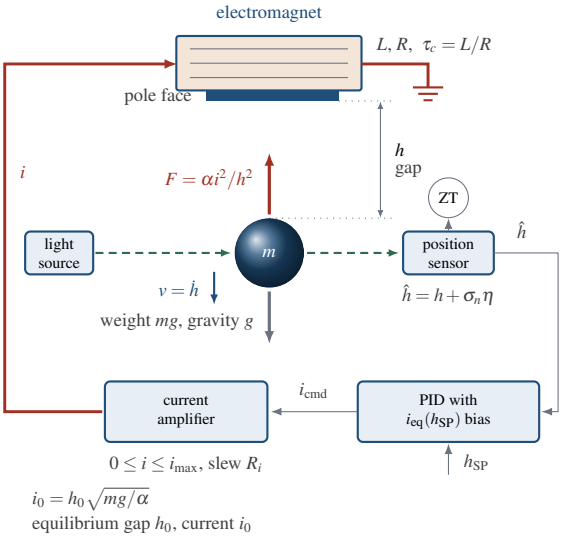


Figure 6.1 Magnetic suspension rig with the electromagnet, the steel ball, the light-beam position sensor and the current path. The gap h is measured downward from the pole face, so the gap rate $v = \dot{h}$ is positive while the ball escapes. Coil current i produces the attractive force $\alpha i^2/h^2$ against the weight mg. The amplifier holds the current between zero and $i_{\max}$ at a slew rate no greater than R_i, and the coil contributes the lag $\tau_c = L/R$. Transmitter ZT reports $\hat{h}$ to a PID controller biased by the equilibrium current i_0 at the setpoint gap h_{SP}.

assessment of maglev economics places the control problem well below the guideway in the cost of a deployment [12].

The single-degree-of-freedom ball and coil became the standard object of study because it isolates the difficulty from everything else. The electromagnet design and the force law for this exact geometry were derived for teaching use [3]. Linear state feedback was compared against exact feedback linearization on one apparatus [13], squaring the current command was shown to cancel most of the nonlinearity [14], a nonlinear controller was built from the same starting point [15], and adaptation was added for an uncertain payload [16]. Underneath the specific designs sits the general accounting for what a right-half-plane pole costs, in the sensitivity integrals of Freudenberg and Looze [17] and in the argument that unstable plants must be respected rather than merely tuned [18]. The simulator of this chapter makes that accounting visible on a plant a student can operate.

6.3 Suspension Model

The simulated rig is a 20 g steel sphere below a single electromagnet driven by a current amplifier, with a light beam across the gap serving as the position measurement. Distance is measured downward from the pole face, so the gap h grows as the ball falls and a positive gap rate $v = \dot{h}$ means the ball is escaping. The controlled variable is the gap reported by the position sensor and the manipulated variable is the current command. There is one loop, one actuator and no inner cascade, so every mechanism in this chapter acts inside a single feedback path. Figure 6.1 draws the rig with each symbol placed where the equations use it.

Table 6.1 sorts the signals into states, the controlled and manipulated variables, the two disturbances and the computed outputs. Two latched events bound the excursion. Below $h_{\min} = 2\,\mathrm{mm}$ the ball is against the pole face and the simulator raises a magnet-contact fault, and above $h_{\max} = 60\,\mathrm{mm}$ it raises a loss-of-capture fault, because beyond that radius no available current produces enough force. Both events annotate the trend, and both clear if the loop recovers the ball, which it can do

Table 6.1 Signals and Their Roles. Nominal Values Are Those of the Default 20 mm Scenario at $t = 0$.

Symbol	Role	Nominal	Range
h	state, true gap	20 mm	2 to 60 mm
$v = \dot{h}$	state, gap rate	0	—
i	state, coil current	1.778 A	0 to $i_{\max}$
$\hat{h}$	CV, measured gap	20 mm	—
h_{SP}	setpoint	20 mm	8 to 34 mm
i_{cmd}	MV, current command	1.778 A	0 to $i_{\max}$ applied
Δm	DV, mass error	0 %	−20 to +20 %
σ_n	DV, sensor noise 1σ	0 mm	0 to 0.5 mm
F	output, magnetic pull	0.196 N	—
mg	output, weight	0.196 N	—
i_{eq}	output, balancing current	1.778 A	—
p	output, unstable pole	31.3 rad/s	—

from contact and cannot do from loss of capture. Three states are integrated in SI units with time in seconds.

🕮 **Pause.** Try Exercise 6.1 before reading on.

6.3.1 Force Balance

The magnet attracts the ball upward with a force proportional to the square of the current and inversely proportional to the square of the gap, while gravity acts downward in the direction of increasing h, as shown in Equation (6.1).

$$m_{\mathrm{eff}}\ddot{h} = m_{\mathrm{eff}}\,g - \frac{\alpha\, i^2}{h^2}, \qquad m_{\mathrm{eff}} = m\left(1 + \tfrac{\Delta m}{100}\right) \tag{6.1}$$

where m is the nominal ball mass in kg, Δm is the mass error in percent set as a disturbance, g is gravitational acceleration in m/s^2 and α is the force coefficient in N m^2/A^2. Both sides of Equation (6.1) carry units of N. The coefficient α lumps coil turns, pole geometry and permeability into one constant, which is the standard reduced form for this configuration [3]. Inside the derivative evaluation the gap is floored at 1×10^{-4} m and the squared gap at 1×10^{-9} m^2, so an integrator excursion cannot divide by zero. Setting $\ddot{h} = 0$ gives the current that balances the weight at any gap, as shown in Equation (6.2).

$$i_{\mathrm{eq}}(h) = h\sqrt{\frac{mg}{\alpha}} = K_{\min}\,h, \qquad K_{\min} \equiv \sqrt{\frac{mg}{\alpha}} \tag{6.2}$$

The constant $K_{\min} = 88.9\,\mathrm{A/m}$ is not a curiosity of the arithmetic. Section 6.6 shows that it is exactly the proportional gain below which no controller of the implemented form stabilizes the ball, at any gap whatsoever. A mass disturbance scales the balancing current by $\sqrt{1 + \Delta m/100}$, and the scaled value is reported alongside the measured gap.

6.3.2 Amplifier, Coil and Sensor

The current cannot step. A first-order current source with a rate limit and a hard ceiling stands between the controller and the coil, as shown in Equations (6.3) and (6.4).

$$\tau_c\,\dot{i} = i_{\mathrm{app}} - i, \qquad 0 \le i \le i_{\max} \tag{6.3}$$

$$i_{\mathrm{app}} = \mathrm{slew}\left(i_{\mathrm{cmd}},\, i_{\mathrm{app}}^{-},\, R_i,\, \Delta t_c\right) \tag{6.4}$$

where τ_c is the coil time constant in s, i_{app} is the current actually applied to the coil in A, i_{app}^{-} is its value at the previous control execution, R_i is the slew limit in A/s and Δt_c is the control period. The

function slew$(\cdot)$ moves the previous applied value toward the command by at most $R_i \Delta t_c$. Physically $\tau_c = L/R$ for a coil of inductance L and resistance R, and the ratio is the adjustable constant rather than the two separately. Setting $\tau_c = 0$ replaces Equation (6.3) by $i = i_{\text{app}}$, and that option exists because τ_c turns out to decide how much derivative action the loop needs. The controller never reads a state. It reads the sensor output of Equation (6.5).

$$\hat{h} = h + \sigma_n \eta(t) \tag{6.5}$$

where σ_n is the noise standard deviation in m and η is a first-order-filtered Gaussian sequence of correlation time $\tau_n = 4\,\text{ms}$, scaled so that its stationary standard deviation is unity. The sequence is drawn from a seeded generator, so a noisy scenario replays identically on any machine, as Section 1.6.1 describes.

6.3.3 Hard Events and Initial Condition

Contact and escape are discontinuous and are applied after the integration step rather than inside the derivative function. If $h \le h_{\min}$ the gap is clipped there and any velocity toward the magnet is zeroed, which is a perfectly inelastic stop against the pole face. If $h \ge h_{\max}$ the gap is clipped and any velocity away from the magnet is zeroed. A negative current is clipped to zero. The contact flag clears above $1.05\,h_{\min}$ and the loss-of-capture flag clears below $0.98\,h_{\max}$, and that hysteresis is what stops the fault banner from flickering at the boundary.

Every scenario starts at the algebraic equilibrium of the setpoint in force at the start of the run, with $h(0) = h_{\text{SP}}$, $v(0) = 0$ and $i(0) = \min\left(i_{\text{eq}}(h_{\text{SP}}),\, i_{\max}\right)$, so a case whose balancing current exceeds the ceiling starts with the coil already at its limit. Two scenarios override $h(0)$ by $\pm 1\,\text{mm}$. No iteration is required because Equation (6.2) is closed form, and that is itself instructive, because the steady state of this plant is trivial to compute and useless on its own.

> **+ DID YOU KNOW? Gaps in Service Are Far Smaller**
>
> An attractive-suspension vehicle surveyed in 1992 kept a clearance of typically 7 to 10 mm to its guideway and levitated even at rest [19]. A compressor on active magnetic bearings had a typical radial clearance of 0.3 mm, and its auxiliary ball bearing 0.15 mm [20]. The chapter's ball hangs 20 mm below its pole face.

Table 6.2 carries every constant a second implementation would need. The proportional gain is entered in A/mm and used internally in A/m, a factor of 1000 that has to be respected when the stability conditions of Section 6.6 are compared with an entered value. Three disturbances are available in any scenario, a mass error spanning $\pm 20\,\%$, a sensor noise standard deviation reaching 0.5 mm, and three one-shot events that displace the ball 1 mm away from the magnet, add 10 % mass, or apply a 0.02 m/s velocity impulse away from the magnet. The impulse is the most selective of the three, because it changes no position and no force, so proportional and integral action see nothing at the instant it lands and only the derivative term reacts.

6.3.4 Assumptions

Only the mechanisms that alter the tuning survive in the model. Every assumption is given below with the operating range in which it is sound and the situation that breaks it.

1. *One translational degree of freedom.* The ball moves only along the coil axis. This is accurate for a centered ball in a rig with a pole face wide relative to the sphere, and it is exactly what a single-axis teaching rig is built to approximate. It fails for any question about lateral guidance, rotation about the axis, or the coupled roll and heave modes that dominate a suspended vehicle [9], and none of those can be studied here.

2. *Inverse-square force law with a lumped coefficient.* The relation $F = \alpha i^2 / h^2$ holds while the core is unsaturated and the gap is small compared with the pole-face diameter, which covers

Table 6.2 Parameters, Controller Settings and Numerics. Every Plant Parameter Is Operator-Adjustable Over the Range Shown Except the Contact Gap $h_{\min}$.

Symbol	Meaning	Value	Unit	Range
Plant				
m	ball mass	0.020	kg	0.010-0.040
α	force coefficient	2.4832×10^{-5}	$\mathrm{N\,m^2/A^2}$	$1\text{-}5\times10^{-5}$
g	gravity	9.81	$\mathrm{m/s^2}$	1.6-24.8
τ_c	coil time constant	0.010	s	0-0.060
$i_{\max}$	current ceiling	3.0	A	1.0-6.0
R_i	current slew limit	400	A/s	20-2000
$h_{\min}$	contact gap	0.002	m	fixed
$h_{\max}$	capture radius	0.060	m	0.040-0.120
τ_n	noise correlation time	0.004	s	fixed
Controller				
K_c	proportional gain	0.400	A/mm	0-1.5
τ_I	integral time	0.50	s	0.05-20
τ_D	derivative time	0.020	s	0-0.08
τ_f	derivative filter	0.001	s	0-0.02
χ_{FF}	equilibrium feedforward	on	—	on/off
Derived at the 20 mm operating point				
$K_{\min}$	$\sqrt{mg/\alpha}$	88.9	A/m	—
i_0	balancing current	1.7778	A	—
p	unstable pole $\sqrt{2g/h_0}$	31.32	rad/s	—
$1/p$	divergence time constant	31.9	ms	—
B	control effectiveness $2g/i_0$	11.04	$\mathrm{m/(s^2\,A)}$	—
Numerics				
Δt	integration step (RK4)	5×10^{-4}	s	—
Δt_c	controller sample period	1×10^{-3}	s	—
—	historian rate	50	Hz	—

the 2 to 60 mm range the simulator uses with α fixed. It fails at large gap, where flux fringing makes the true exponent smaller than two, and at high current, where core saturation makes the force grow more slowly than i^2. A design study would need the flux-path model rather than the lumped constant [3].

3. *Point mass and rigid ball.* The sphere is a point mass with no induced-dipole correction, no eddy currents and no remanent magnetization, so the force depends on the present current alone. This is adequate at the millimeter-per-second speeds the loop produces in normal operation. It fails during a fast escape, where eddy-current drag in a solid steel sphere would add a velocity-dependent term that the model does not carry, and it fails for a hysteretic material, where the force would depend on the current history.

4. *First-order current amplifier.* The power stage is a current source with time constant τ_c, a slew limit R_i and a ceiling $i_{\max}$. This represents a current-mode amplifier working well inside its voltage headroom, which is the normal operating region. It fails when the supply voltage limits first, because back electromotive force from a moving ball and the gap-dependence of the coil inductance both change the achievable $\mathrm{d}i/\mathrm{d}t$, and the inductance dependence matters most over a wide gap range.

5. *Ideal gap sensor apart from added noise.* The measurement is the true gap plus a colored noise of correlation time 4 ms. The noise model is realistic for a photodetector chain, and the correlation time is short relative to the 31.9 ms divergence constant, so the noise is genuinely

broadband as far as the loop is concerned. The sensor has no bandwidth limit, no quantization, no calibration drift and no dependence of gain on gap, so any question about a sensor whose scale factor changes with distance is outside the model.

6. *Perfectly inelastic pole face.* Contact clips the gap and zeros the approach velocity. The result is the correct qualitative outcome, a ball that stays on the magnet until the current is reduced. It fails as a description of the impact itself, because a real sphere rebounds, deforms locally and can chip a coating, and no restitution coefficient is available to change that.

7. *Loss of capture declared at a fixed radius.* Escape is declared at $h_{\max}$ rather than derived from the field. At the default 3 A ceiling the true limit is $i_{\max}/K_{\min} = 33.8\,\text{mm}$, so the 60 mm radius is conservative in the sense that the ball is already unrecoverable well before it is reached. The fixed radius is wrong whenever the ceiling is changed, and the ceiling is adjustable, so the declaration point does not track the physics.

8. *Constant m, g and α during a run.* Mass changes only through the mass-error disturbance and the two mass events. The coefficient α does not drift with coil temperature, although coil resistance in a real rig rises with dissipation and pulls the current down at a fixed drive voltage. Over the tens of seconds a scenario runs, thermal drift is negligible, and over a laboratory session on hardware it is not.

9. *Fixed-step integration and a fixed control period.* Classical fourth-order Runge-Kutta on a 0.5 ms step resolves an instability whose time constant is 31.9 ms with a margin of more than sixty steps per *e*-folding, and the 1 ms control period is comfortable against the same figure. Both become marginal near the 8 mm setpoint limit, where the unstable pole reaches 49.5 rad/s and the coil contributes another at 100 rad/s. Reducing the gap below about 5 mm would require a smaller step.

10. *Continuous-time stability analysis.* The three conditions derived in Section 6.6 come from a linearized continuous-time loop with no derivative filter and no sampling. They are accurate to within one setting increment against the nonlinear simulator, as Section 6.6 reports. They are not a substitute for a sampled-data treatment of a real kilohertz loop, and they say nothing about robustness margins.

11. *Omitted entirely.* Air drag, acoustic and structural resonance of the mounting, magnetic hysteresis, stray fields, amplifier deadband and current measurement noise are all absent. Leaving them out does not alter the feedback structure this application is built to expose.

6.4 Linearization and the Unstable Pole

Writing $h = h_0 + \delta h$ and $i = i_0 + \delta i$ with $i_0 = i_{\text{eq}}(h_0)$ and expanding Equation (6.1) to first order collapses both partial derivatives onto the equilibrium itself, because $\alpha i_0^2/h_0^2 = mg$, as shown in Equation (6.6).

$$\delta\ddot{h} = \underbrace{\frac{2g}{h_0}}_{p^2}\,\delta h \;-\; \underbrace{\frac{2g}{i_0}}_{B}\,\delta i \tag{6.6}$$

where $p = \sqrt{2g/h_0}$ in rad/s is the unstable pole and $B = 2g/i_0$ in m/(s^2 A) is the control effectiveness. This is the linearization of Section 1.4.3 taken at one gap. The open-loop poles are $\pm p$, and by Equation (6.2) the ratio of the two coefficients is a constant of the hardware rather than of the operating point, as shown in Equation (6.7).

$$\frac{p^2}{B} = \frac{2g/h_0}{2g/i_0} = \frac{i_0}{h_0} = K_{\min} \tag{6.7}$$

Table 6.3 Operating-Point Dependence. Closed-Loop Poles Are Roots of (6.16) at the Default Tuning $K_c = 400\,\mathrm{A/m}$, $\tau_I = 0.5\,\mathrm{s}$, $\tau_D = 20\,\mathrm{ms}$.

	11 mm	20 mm	28 mm
i_0 (A)	0.978	1.778	2.489
p (rad/s)	42.2	31.3	26.5
$1/p$ (ms)	23.7	31.9	37.8
B (m/(s^2 A))	20.1	11.0	7.88
dominant pair	-23.5 ± 105.3j	-18.5 ± 70.8j	-14.4 ± 56.2j
damping ratio ζ	0.22	0.25	0.25

Including the amplifier lag, the linearized transfer function from current command to gap is given by Equation (6.8).

$$\frac{\delta h(s)}{\delta i_{\mathrm{cmd}}(s)} = \frac{-B}{(s-p)(s+p)(\tau_c s+1)} \tag{6.8}$$

The negative sign is the physical statement that more current means a smaller gap, and it is the reason the controller runs a reversed error convention.

Pause. Try Exercise 6.2 before reading on.

Table 6.3 lists both quantities at the three gaps the application ships as scenarios. They move in opposite directions. Further from the magnet the instability is slower, which is an advantage, but B falls as well, because a larger equilibrium current means each additional ampere buys proportionally less acceleration. Their ratio does not move at all, by Equation (6.7). One tuning therefore produces three different closed-loop responses across the setpoint range, with the damping ratio nearly unchanged while the natural frequency falls by almost a factor of two.

6.4.1 Manual Operation at the Equilibrium Current

Manual operation sets the coil current directly, in 0.01 A increments, and a helper, the Set current to equilibrium button, evaluates Equation (6.2) at the present gap and applies it. The applied current is not approximately right. It is right to the last digit double precision carries.

Running that configuration untouched reports loss of capture at $t = 3.06\,\mathrm{s}$, and what grew over that interval was floating-point round-off. The computed net acceleration at the exact current is $1.8 \times 10^{-15}\,\mathrm{m/s^2}$, one unit in the last place of g, and the gap does not change in its last bit until $t = 1.95\,\mathrm{s}$. The scenario's stored manual current, 1.777 762 116 9 A, carries eleven significant digits and falls 1.3×10^{-11} A short of the exact value, which shifts the equilibrium by $\delta i / K_{\mathrm{min}} = 1.4 \times 10^{-13}\,\mathrm{m}$, and the same untouched run loses capture at $t = 0.90\,\mathrm{s}$. Covering 20 mm to the 60 mm capture radius from that offset at $p = 31.3\,\mathrm{rad/s}$ takes 0.86 s in the linearization. A deliberate displacement shortens the interval. A $+1$ mm offset loses capture at $t = 0.17\,\mathrm{s}$ and a -1 mm offset reaches the pole face at $t = 0.10\,\mathrm{s}$, faster because the instability is faster on the near side. A human visual-motor loop closes in roughly 250 ms, which is eight e-foldings, or a factor of 2500 by the time a correction is applied. Manual levitation is not difficult. It is arithmetically excluded. Feedback has more time than a hand but not much more. Switched to automatic operation during the fall from the $+1$ mm offset, the loop still catches the ball if the switch comes by $t = 0.07\,\mathrm{s}$, at a 24 mm gap.

> **TRY IT IN THE APP** **2 · +1 mm Displacement (Dropped)**
>
> Load the scenario, which starts in Manual with the ball 1 mm farther from the magnet than its balance point, press Run and watch the gap trend: capture is lost at $t = 0.17\,\mathrm{s}$. Reload it, run, and switch to Auto by $t = 0.07\,\mathrm{s}$, at a 24 mm gap: the loop still catches the ball.

6.5 Controller Design

One digital PID executes every $\Delta t_c = 1\,\text{ms}$, twice as slowly as the physics step. It is the velocity-form PID of Section 1.6.2 [7], with the derivative taken on the measurement so that moving the setpoint produces no derivative kick. The error is written with the measurement first, because the process gain in Equation (6.8) is negative and a larger gap must be answered with a larger current, as shown in Equations (6.9) to (6.12).

$$e_k = \hat{h}_k - h_{\text{SP}} \tag{6.9}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1}) \tag{6.10}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{6.11}$$

$$\Delta u_k^D = K_c\,\tau_D\,\frac{\hat{h}_k^f - 2\hat{h}_{k-1}^f + \hat{h}_{k-2}^f}{\Delta t_c} \tag{6.12}$$

where e_k is the error in m and $\hat{h}_k^f$ is the filtered measurement in m. That filter is an exact zero-order-hold first-order lag of time constant τ_f, $\hat{h}_k^f = a\,\hat{h}_{k-1}^f + (1-a)\hat{h}_k$ with $a = \exp(-\Delta t_c/\tau_f)$, so it is stable for any $\tau_f \geq 0$ because a is never negative. The three increments are combined with the equilibrium bias and the demand is clipped and rate limited, as shown in Equation (6.13).

$$\begin{aligned} z_k &= z_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= i_{\text{eq}}(h_{\text{SP}})\,\chi_{\text{FF}} + z_k + I_{k-1} \\ u_k &= \text{slew}(\text{clamp}(c_k, 0, i_{\max}), u_{k-1}, R_i, \Delta t_c) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{6.13}$$

where u_k is the applied current command in A, $\chi_{\text{FF}} \in \{0, 1\}$ selects the equilibrium feedforward, z_k and I_k are in A, and $T_t = \tau_I$ is the tracking time of the back-calculation in s. These are Equations (1.13) to (1.15) of Section 1.6.2, with the running sum z_k in place of S_k, the demanded current c_k in place of v_k and the equilibrium current as the bias u_b. While the command is clipped or rate limited, the term $(\Delta t_c/T_t)(u_k - c_k)$ drives I_k toward the value that reproduces the applied current, removing $1/500$ of the excess per scan at the default $\tau_I = 0.5\,\text{s}$, so only the integral is corrected and the running sum z_k is never truncated [7]. The tracking time matters on this plant. A displacement step reaches the derivative path as a step in the measurement, and the 1 mm drop of Section 6.7 asks for 6.4 to 7.9 A on the scan at which it lands. A recursion that absorbed the whole clipped excess in one scan would load I_k with about $-5\,\text{A}$ and lose the ball after the drop at 28 mm. The feedforward term is unity times $i_{\text{eq}}(h_{\text{SP}})$ unless it is switched off, in which case the loop has to build the whole equilibrium current out of proportional and integral action.

In manual operation I_k is set so that $i_{\text{eq}}(h_{\text{SP}})\chi_{\text{FF}} + z_k + I_k$ equals the operator's current exactly, and both transfers are bumpless, as Section 1.6.2 describes. Output limits are 0 to $i_{\max}$ with the slew limit R_i applied afterwards, and the same slew acts on the manual command because a real amplifier cannot step either. The loop is direct acting in the ISA sense, since (6.9) writes the error as $\hat{h}_k - h_{\text{SP}}$ and a positive K_c therefore raises the current as the gap grows, which is the action $a = -1$ of Equation (1.12). The sample period is fixed at 1 ms regardless of the speed setting, and there is no output bias other than the feedforward. The model contains no transport delay, and the application implements no cascade, no gain scheduling and no feedforward beyond the algebraic bias in Equation (6.13), because those structures do not apply to a single-actuator suspension.

Three model layers are kept apart throughout. The simulation model is Equations (6.1) to (6.5), nonlinear and integrated at 0.5 ms. The analysis model is the linearization (6.6) with the amplifier lag, used to derive stability boundaries and, once discretized, as the prediction model of the explicit

controller of Section 6.5.1, and nowhere else. The PID's internal model is the single algebraic relation (6.2) evaluated at the setpoint, so the PID knows nothing of p, B or τ_c, and its feedforward uses the nominal mass, which makes it wrong by design whenever a mass disturbance is applied. The feedforward and the starting command are evaluated with the declared m, α and g of Table 6.2, fixed when the application loads, so a ball, magnet or gravity changed as a parameter also reaches the PID only through the measured gap.

6.5.1 Explicit Model Predictive Control

A second controller can take the coil in automatic operation in place of the PID. It is a model predictive controller (MPC) whose quadratic program is solved offline for every state it can meet. The predictive controllers of Section 1.6.3 solve their program at every decision, which a kilohertz loop cannot afford, and Volume 2, Chapter 12 uses a table of the same kind on a machine-tool axis. When the linear term and the right-hand sides of a quadratic program depend affinely on a parameter, its minimizer is piecewise affine in that parameter over polyhedral critical regions, one for each set of active constraints that is optimal somewhere, so the constrained linear quadratic regulator (LQR) can be stored as a table of regions and laws and evaluated by lookup [21], at a memory and search cost that grows quickly with the horizon and the number of constraints [22]. The PID stays the default, and only scenario 9 selects the MPC.

The prediction model is Equation (6.6) with the coil lag of Equation (6.3), linearized at $h_0 = 20\,\mathrm{mm}$ with the declared constants of Table 6.2 ($p^2 = 981.0/\mathrm{s}^2$, $B = 11.04\,\mathrm{m}/(\mathrm{s}^2\,\mathrm{A})$, $1/\tau_c = 100/\mathrm{s}$). Its state is $x = (h - h_{\mathrm{SP}},\, v,\, i - i_{\mathrm{ss}})$ and its input the move $\delta u = i_{\mathrm{cmd}} - i_{\mathrm{ss}}$, both measured from a target $(h_{\mathrm{SP}}, 0, i_{\mathrm{ss}})$, and an exact zero-order hold over $T_s = \Delta t_c = 1\,\mathrm{ms}$ gives $x_{k+1} = \Phi x_k + \Gamma\,\delta u_k$. The model is written from the declared physics rather than identified, because no open-loop step test exists for a plant that leaves its equilibrium in 32 ms, and central differences of the simulator's own derivative function reproduce $2g/h_0$ and $2g/i_0$ to relative errors of 5×10^{-11} and 7×10^{-11}. Its slopes remain those of 20 mm at every setpoint.

Every millisecond the controller plans three moves of the commanded current, one per period, and applies the first. The objective weighs the gap over the first two predicted steps, prices the state after the third with the LQR's cost-to-go and charges each move's distance from the target current, subject to the ceiling and the slew on all three moves, as shown in Equation (6.14).

$$\begin{aligned} \min_{\delta u_0, \delta u_1, \delta u_2} \; & \sum_{k=1}^{2} Q_h x_{1,k}^2 + x_3^T P x_3 + R \sum_{k=0}^{2} \delta u_k^2 \\ \text{subject to } & x_{k+1} = \Phi x_k + \Gamma\,\delta u_k, \quad x_0 = \hat{x} \\ & 0 \le i_{\mathrm{ss}} + \delta u_k \le i_{\max} \\ & |\delta u_k - \delta u_{k-1}| \le R_i T_s, \quad \delta u_{-1} = i_{\mathrm{app}}^{-} - i_{\mathrm{ss}} \end{aligned} \tag{6.14}$$

where $x_{1,k}$ is the predicted gap deviation, $\hat{x}$ is the estimated state, and P solves the discrete algebraic Riccati equation for Φ, Γ, $\mathrm{diag}(Q_h, 0, 0)$ and R, the terminal weight of the standard stabilizing construction [23], used here without a terminal constraint. Bryson's rule (Section 1.4.5) on the 1 mm band and on the 1.222 A between i_0 and the declared 3 A ceiling gives $Q_h = 1 \times 10^6/\mathrm{m}^2$ and $R = 0.669/\mathrm{A}^2$. With no constraint active the first move is the LQR's, $\delta u = 1.356\,\mathrm{A/mm}\,x_1 + 23.88\,\mathrm{A\,s/m}\,x_2 - 1.497\,x_3$, with closed-loop poles at $-62.5 \pm 80.3\mathrm{j}$ and −130.6/s, a rate gain 2.98 times the PID's $K_c\tau_D = 8\,\mathrm{A\,s/m}$, and feedback on the estimated coil current. Every constraint is hard, the ceiling and the slew are read as the PID reads them, and only inputs are constrained, so the program is feasible at every state.

In the parameter $\theta = (\hat{x},\, i_{\mathrm{app}}^{-} - i_{\mathrm{ss}},\, -i_{\mathrm{ss}},\, i_{\max} - i_{\mathrm{ss}},\, R_i T_s,\, 1)$ the target, the ceiling and the slew enter only the right-hand sides of the twelve inequality rows, so one table serves every setting a learner selects. Of the 173 sets of at most three independent active rows, each with an affine law from its optimality conditions, a set is kept when its critical region holds a ball of positive radius in the

stated domain. That domain is the application's own range of setpoints, ceilings and slews, gaps from contact to the capture radius, rates within ±1 m/s, coil currents of 0 to 6 A and disturbance estimates within ±3 m/s^2, about 30 % of g. It holds 91 regions, of which 63 occur at the declared 3 A and 400 A/s. Gaps to 120 mm and rates to ±5 m/s add one sliver whose center lies in a stored region, and boxes 10^4 times wider give 101 to 105 regions. One to five moves give 5, 23, 91, 283 and 710 regions, and the design admitted fewer than 100, which fixed the horizon at three.

A decision tests the regions in order, the unconstrained one first, each by the sign of its multipliers and by the twelve rows, which define the region, so a region found returns the exact optimum. A parameter that no region accepts is a miss. The PID, which tracks the applied current at every execution, then takes the coil without a bump, the record is marked, and after 20 found decisions in a row the table returns over a 20 ms blend. The lookup agrees with an online solve of the same program by the dual active-set solver of the series to 4.6×10^{-14} A on 100,000 uniform parameters (98,719 with a row active) and to 1.0×10^{-10} A within 10^{-6} of 3000 region boundaries, and 203,983 closed-loop decisions on every scenario and seven hard cases never missed. In Node.js a lookup has a median of 0.72 µs (0.93 µs in the slowest region) against 7.6 µs for the online solve. In the Chrome browser, on the application's frame loop at one and two times real time, the lookup adds 0.16 to 0.31 % of one core and the online solve 1.1 to 2.7 %, both measured on a host shared with other computations.

The MPC reads the measured gap, the setpoint, its applied current and the ceiling and slew. A constant-gain extended Kalman filter (Section 1.4.6) estimates the gap, the rate, the coil current and a constant downward acceleration $\hat{d}$ that stands for every unmodeled force. It predicts with Equations (6.1) and (6.3), the declared constants and $g + \hat{d}$, in two fourth-order Runge-Kutta substeps per period, and it corrects with the steady Kalman gain $(0.187,\ 19.25/\mathrm{s},\ 0,\ 901.9/\mathrm{s}^2)$ of the 20 mm linearization augmented with d, designed for 0.05 mm of measurement noise and 0.05 m/s^2 per step on d, which leaves error poles at $-51.6 \pm 83.8\mathrm{j}$, -103.3 and -100/s. The estimate shifts the target to the declared equilibrium of Equation (6.15), the target-shift form of offset-free predictive control [24].

$$i_{\mathrm{ss}} = h_{\mathrm{SP}} \sqrt{\frac{m\,(g + \hat{d})}{\alpha}} \tag{6.15}$$

On scenario 4 the estimate settles at $\hat{d} = 0.981\,\mathrm{m/s}^2$, the declared model's value for the 10 % heavier ball, with the gap 5.6×10^{-14} mm from its setpoint, while the same controller without the shift ends 0.193 mm off. Volume 9, Chapter 4 removes the offset of a poorer fuel with a Kalman filter that likewise estimates a constant disturbance, there on the fuel. The filter has no pole face, so a measured gap more than 5 mm from the prediction, twice the largest innovation of 38 runs without contact or loss of capture (2.44 mm), restarts it at rest at the measured gap with $\hat{d} = 0$.

The plant integrates Equations (6.1) to (6.5) with the run's parameters, while the MPC holds only the declared ones, so a ball, magnet, gravity or coil lag changed by a learner is model mismatch for the MPC, as it is feedforward error for the PID. The estimator runs in every mode and structure and writes nothing the PID reads, so every scenario and suggested tuning under the PID records the same run, to the last bit, as before the MPC was added.

6.6 Stability Conditions for the Closed Loop

Substituting the ideal PID transfer function $K_c(1 + 1/\tau_I s + \tau_D s)$ into Equation (6.8) and clearing denominators gives the closed-loop characteristic polynomial of Equation (6.16).

$$\tau_c \tau_I s^4 + \tau_I s^3 + \tau_I \left(B K_c \tau_D - p^2 \tau_c\right) s^2 + \tau_I \left(B K_c - p^2\right) s + B K_c = 0 \tag{6.16}$$

Applying the Routh-Hurwitz test and simplifying with $p^2/B = K_{\min}$ produces three conditions, each of which reads as a sentence about the hardware.

The first condition is that the proportional gain must exceed the magnetic stiffness, as shown in Equation (6.17).

$$K_c > \frac{p^2}{B} = K_{\min} = \sqrt{\frac{mg}{\alpha}} = 88.9\,\mathrm{A/m} \tag{6.17}$$

Below this value the loop cannot reverse the sign of the position feedback, because the magnet pulls harder with displacement than the controller pushes back. By Equation (6.7) the bound is identical at every gap, which is why the application can print it as one fixed number, 0.089 A/mm, labeled necessary but not sufficient.

The second condition is that the derivative time must exceed the coil time constant, as shown in Equation (6.18).

$$\tau_D > \tau_c \tag{6.18}$$

This is the sharpest result available from the analysis and the least obvious one. The Routh element that decides it evaluates to $\tau_I B K_c(\tau_D - \tau_c)$, so the proportional gain cancels completely and no amount of gain buys stability if the derivative lead is slower than the lag it has to overcome. With $\tau_c = 10\,\mathrm{ms}$ a derivative time of 6 ms fails at $K_c = 150\,\mathrm{A/m}$ and fails equally at 1000 A/m. It is the lag limit of Equation (1.9) in Section 1.4.7, written for the coil, and Chapter 5 finds the same condition with a servo's lag in place of the coil's.

The third condition is that integral action must be slow enough to be paid for out of the margin the first two conditions earn, as shown in Equation (6.19).

$$\left(BK_c - p^2\right)\left(\tau_D - \tau_c\right) > \frac{1}{\tau_I} \tag{6.19}$$

The two margins on the left form a budget that reset action spends. The default tuning has $BK_c - p^2 = 3434/\mathrm{s}^2$ and $\tau_D - \tau_c = 10\,\mathrm{ms}$, a product of 34.3/s against a requirement of 2/s, so the default sits well inside the region rather than on its edge.

6.6.1 Proportional-Only and Integral-Only Limits

Setting $\tau_c = 0$ and $\tau_D = 0$ in Equation (6.16) leaves $s^2 + (BK_c - p^2) = 0$, a pair of poles on the imaginary axis. Proportional action alone, however large, converts a runaway into an undamped oscillation and goes no further. It never levitates anything.

Keeping reset action but still without derivative action or coil lag gives $\tau_I s^3 + 0 \cdot s^2 + \tau_I(BK_c - p^2)s + BK_c = 0$. The s^2 coefficient is identically zero, so the Routh-Hurwitz necessary condition fails for every K_c and every τ_I. Integral action cannot stabilize this plant at any setting whatsoever. The ranking of the three terms is therefore fixed by the characteristic polynomial rather than by tuning preference, and it is the reverse of the ranking a first course usually establishes.

6.6.2 Checking the Boundaries Against the Simulator

Conditions (6.17) to (6.19) come from a linearization with no derivative filter and no sampling, so they are worth testing against the nonlinear model. Loading the graded scenario, which applies a $+5\,\%$ mass step at $t = 3\,\mathrm{s}$, and asking whether the ball is still captured at $t = 9\,\mathrm{s}$ brackets both testable boundaries within one setting increment. With $\tau_I = 0.5\,\mathrm{s}$ and $\tau_D = 20\,\mathrm{ms}$, Equation (6.19) predicts a critical gain of 107.0 A/m. The ball is lost at 95 and 105 A/m, survives at 110 A/m while spending 13.3 A of current travel against 1.1 A at 120 A/m, and is held cleanly at the higher gain. With $K_c = 150\,\mathrm{A/m}$ and $\tau_I = 5\,\mathrm{s}$, the same condition predicts a critical derivative time of 10.30 ms, and the ball is lost at 9 and 10 ms and held at 11 and 12 ms. Both sets of outcomes are plotted against the predicted boundaries in Figure 6.2.

The residual discrepancy is itself instructive. The capture test cannot place a boundary more finely than the slider increments, and at 110 A/m the ball survives in an oscillation whose swing is still growing, from 4.4 mm peak to peak at $t = 4\,\mathrm{s}$ to 5.9 mm at $t = 38\,\mathrm{s}$. Runs off the slider grid, which a learner cannot make, place the simulator's own boundaries slightly above the predicted ones, between

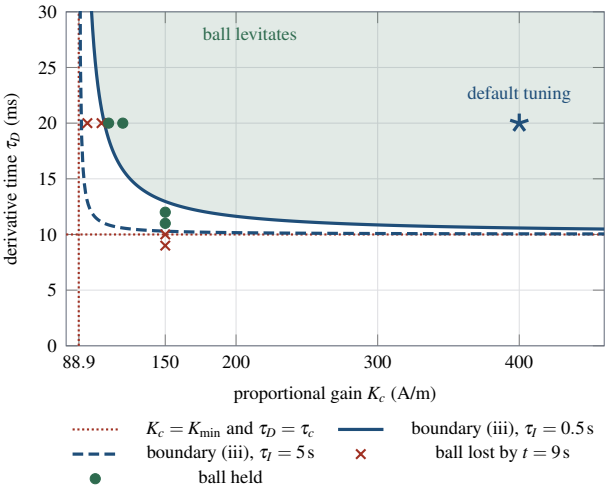


Figure 6.2 Predicted stability region and measured outcomes at the 20 mm gap. The shaded region satisfies all three conditions at $\tau_I = 0.5$ s. The four points at $\tau_D = 20$ ms were run at $\tau_I = 0.5$ s and bracket the predicted critical gain of 107 A/m. The four points at $K_c = 150$ A/m were run at $\tau_I = 5$ s and bracket the predicted critical derivative time of 10.30 ms, so they are judged against the dashed boundary. Crosses mark runs that lost the ball and filled circles runs that held it. The dotted lines are the asymptotes $K_c = K_{\min}$ and $\tau_D = \tau_c$, and the solid curve is the boundary at $\tau_I = 0.5$ s.

110 and 110.5 A/m and between 10.9 and 11.0 ms, because the 1 ms derivative filter and the 1 ms sample period each add a phase lag that Equation (6.18) does not know about, and both lags act in the one path that is holding the ball up.

6.7 Upper Bounds on the Proportional Gain

Conditions (6.17) to (6.19) are all lower bounds. If they were the whole story every gain would be set to its highest available value. Two mechanisms stop that, and the application separates them so each can be met alone.

6.7.1 Amplifier Ceiling

The static current requirement grows linearly with gap by Equation (6.2), so the default 3 A ceiling makes gaps beyond $i_{\max}/K_{\min} = 33.8$ mm unreachable in principle, and the highest setpoint offered, 34 mm, cannot be held at all. The ceiling consumes the dynamic range long before it consumes the static one. Table 6.4 reports the response to a 1 mm displacement away from the magnet, applied as a one-shot event, at three gaps with the tuning unchanged. Peak coil current is 1.27 A at 11 mm and 2.78 A at 28 mm for the identical disturbance, which is 42 % and 93 % of the ceiling, and at 28 mm the command itself sits on the 3 A ceiling for three scans. The increment the loop adds is almost the same at all three gaps, 0.29 A at 11 mm, 0.30 A at 20 mm and 0.29 A at 28 mm, so what consumes the headroom is the equilibrium current itself, which rises linearly with the gap. The gap at which the instability is slowest, 28 mm with $p = 26.5$ rad/s, is also the gap at which the amplifier has the least margin.

One scenario removes the ambiguity by lowering the ceiling to 2.4 A while asking for 28 mm, where $i_{\text{eq}} = 2.49$ A. No equilibrium exists at all, since the highest gap the reduced ceiling can hold is 27.0 mm. The command pins at the limit from the first scan, the ball leaves, and loss of capture is declared at $t = 0.18$ s with a gap error of 32 mm. Lowering the setpoint to 26 mm before the run starts holds the ball on 2.31 A, because the ball then starts at the new gap. Restoring the 3 A ceiling at 28 mm lets the coil start on the 2.49 A equilibrium current, and the ball is held. An actuator that

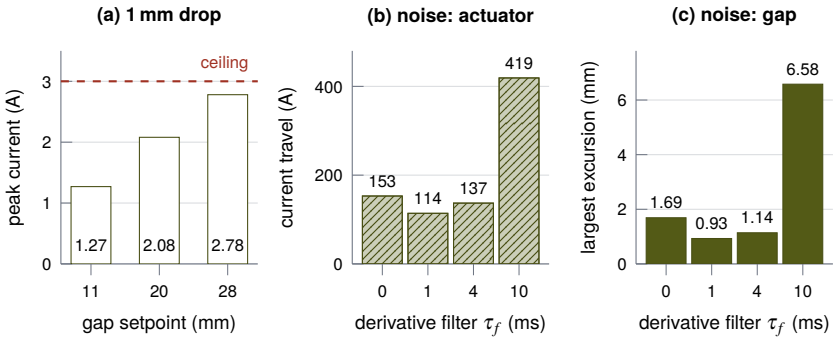


Figure 6.3 The two blocks of Table 6.4. (a) Peak coil current after the same 1 mm displacement at three gaps (white bars), against the 3 A ceiling (dashed line). (b) Current travel (hatched bars) and (c) largest gap excursion (solid bars) over 9.3 s of the noisy-sensor scenario for four derivative filters. Notice that the drop costs almost the same increment at every gap, so the headroom shrinks with the balancing current, and that both the actuator and the gap are best at the 1 ms filter, with 10 ms the worst on both counts.

physically cannot reach the demand is not a tuning problem, and that lesson transfers directly to valves, thrusters and inverters. It is the actuator headroom of Section 1.4.7 in its static form, and Volume 2, Chapter 3 meets it as a dilution that puts the required feed concentration outside the actuator range.

> **TRY IT IN THE APP** **5 · Current Saturation**
>
> Load the scenario and compare the `i_eq at SP` badge, 2.49 A, with the 2.4 A ceiling. Press Run: the command pins at the limit from the first scan and capture is lost at $t = 0.18$ s. Lower Gap setpoint to 26 mm, press Reset and run again: the ball is held on 2.31 A.

6.7.2 Sensor Noise

Derivative action is compulsory by Equation (6.18), and derivative action differentiates noise. The derivative path carries a gain of $K_c\tau_D = 8\,\mathrm{A\,s/m}$ at the default tuning, so a measurement wobbling by 0.15 mm with a 4 ms correlation time injects tenths of an ampere of pure chatter, a standard deviation of 0.44 A in the command at the default filter. The ball filters most of that out of the gap record, which wanders by 0.33 mm (one standard deviation), and the accumulated actuator movement is the reported quantity in which the chatter shows most.

The lower block of Table 6.4 sweeps the derivative filter at fixed noise. Total current travel over 9.3 s falls from 153 A unfiltered to 114 A at the default $\tau_f = 1$ ms, rises again to 137 A at 4 ms, and reaches 419 A at 10 ms, where the largest gap excursion grows from 0.93 mm to 6.58 mm as well. The same ordering holds for eight noise seeds, with travel of 107 to 124 A at 1 ms and 127 to 149 A at 4 ms, and Figure 6.3 draws both blocks of the table. That curve carries the design argument. A filter long enough to matter is a lag in exactly the path that is holding the ball up, and Equation (6.18) says that path has no margin to spare. The optimum is bounded below by the noise and above by the coil, and the usual rule of thumb $\tau_f \approx \tau_D/10$ puts it at 2 ms on this plant, within a factor of two of the 1 ms sample period, which is why physical suspensions run their loops at 10 kHz rather than at 1 kHz.

6.8 Results

Ten set-ups each change one quantity and name what they changed. The first three concern the existence of an equilibrium and are run in manual operation (the first opens in automatic operation and is switched by the learner). The next five run in automatic operation and concern the bounds on

Table 6.4 Measured Performance. Upper Block: a 1 mm Displacement Away From the Magnet at $t = 1$ s With the Default Tuning, Read at $t = 4.7$ s, by Which Time Every Entry Has Converged. Peak Coil Current Is Tracked at the 1 kHz Controller Rate. Lower Block: the Noisy-Sensor Scenario at $\sigma_n = 0.15$ mm, Read at $t = 9.3$ s.

1 mm displacement at three gaps			
h_{SP} (mm)	11	20	28
IAE (mm s)	0.030	0.033	0.049
Current travel (A)	1.46	1.25	1.44
Peak current (A)	1.27	2.08	2.78
Peak over ceiling (%)	42	69	93
Derivative filter under 0.15 mm sensor noise			
τ_f (ms)	0	1	4	10
IAE (mm s)	3.7	2.3	2.8	25.3
Current travel (A)	153	114	137	419
Largest gap excursion (mm)	1.69	0.93	1.14	6.58

Table 6.5 Scenario Set. Modes Are Initial Values and Either Mode May Be Selected at Any Time.

#	Set-up	Mode	Mechanism isolated
1	exact equilibrium, 20 mm	auto	equilibrium is not stability
2	+1 mm displacement	manual	runaway away from the magnet
3	−1 mm displacement	manual	faster runaway on the near side
4	+10 % ball mass	auto	feedforward error, reset removes it
5	$i_{max} = 2.4$ A, SP 28 mm	auto	demand outside actuator capability
6	$\sigma_n = 0.15$ mm	auto	derivative noise amplification
7	SP 28 mm	auto	slower pole, weaker effectiveness
8	SP 11 mm	auto	faster pole, stronger effectiveness
9	explicit MPC, SP 20 to 26 mm at $t = 1$ s	auto	step into the current ceiling, MPC against PID
10	+5 % mass at $t = 3$ s, $K_c = 150$ A/m, $\tau_I = 5$ s, $\tau_D = 6$ ms	auto	graded, violates (6.18)

the gain, the ninth sets the explicit MPC against the PID, and the last is graded. Table 6.5 lists them with the mechanism each one isolates. Every scenario leaves the remaining settings at their defaults, so the effect of the single change is unambiguous.

> ✦ DID YOU KNOW? **A Longer Filter Moves the Coil More**
>
> From the simulator: under 0.15 mm of sensor noise, the coil's current travel over 9.3 s is 153 A with no derivative filter, 114 A at 1 ms, 137 A at 4 ms and 419 A at 10 ms, where the largest gap excursion grows to 6.58 mm (Table 6.4).

The graded challenge is shown in Figure 6.4. It ships $K_c = 150$ A/m, $\tau_I = 5$ s and $\tau_D = 6$ ms, so Equation (6.17) holds while Equation (6.18) is violated by 4 ms, and the loop nonetheless holds a perfectly steady gap for three seconds because nothing has disturbed it. The mass step then excites the $+3.18 \pm 25.0$ rad/s pair predicted by Equation (6.16), the oscillation grows by a factor of e every 314 ms (320 ms between the second and third peaks of the record), the current command first reaches the 3 A ceiling on the third swing at the scan at $t = 3.64$ s, and the ball is gone at $t = 3.94$ s. Raising K_c is the response the first eight scenarios encourage, and it changes nothing here, because Equation (6.18) does not contain K_c. Every gain from 300 to 1500 A/m loses the ball between 3.56 and 3.73 s. The graded target is a gap held within ± 1 mm of setpoint from the mass step to $t = 12$ s with no time at a current limit and a peak below 3 A, and a suggested tuning can be

Table 6.6 Setpoint Step From 20 to 26 mm at $t = 1$ s (Scenario 9), Measured Over $t = 1$ to 3 s

Controller	IAE (mm s)	$n_{3\,\mathrm{A}}/n_{0\,\mathrm{A}}$	i_{pk} (A)	Outcome
PID, default tuning	62.84	1964/3	3.000	lost, 1.225 s
PID, LQR gains	64.80	1968/20	3.000	lost, 1.144 s
I-PD	2.98	0/0	2.308	held
MPC, online solve	0.319	69/12	2.998	held
MPC, lookup	0.319	69/12	2.998	held
Counterfactual runs				
PID, $i_{\max} = 6$ A	0.405	0/3	3.542	held
PID, 50 ms setpoint ramp	0.337	39/0	2.989	held
MPC law clipped, no look-ahead	0.325	70/13	2.999	held
MPC online, ten moves	0.257	55/12	2.993	held

IAE of the true gap at every 0.5 ms step from the event, finer than the 50 Hz scorecard of Table 6.4. $n_{3\,\mathrm{A}}$ and $n_{0\,\mathrm{A}}$ count the 1 ms decisions with the command at the ceiling and at zero, and i_{pk} is the peak coil current. The clipped law is the unconstrained law clipped into the box and the slew on each move. It and the ten-move MPC run outside the application.

revealed for comparison. The challenge is judged on the case as given, so a changed model parameter or disturbance voids an attempt. Read at $t = 12$ s by the challenge's own rule, the shipped tuning leaves a worst gap error of 40.00 mm after the step, 67.0 % of the run at a current limit and a peak of 3.00 A. The suggested tuning, $K_c = 400\,\mathrm{A/m}$, $\tau_I = 0.5$ s and $\tau_D = 20$ ms, keeps the worst error to 0.18 mm, 0.82 mm inside the band, with no time at a limit and a peak of 1.87 A.

Pause. Try Exercise 6.3 before reading on.

6.8.1 Explicit MPC Against the PID

Scenario 9 steps the setpoint from 20 to 26 mm at $t = 1$ s under the MPC, where the declared equilibrium current is 2.31 A of the 3 A ceiling, and it is credited when the gap is within 0.5 mm of 26 mm at $t = 3$ s and the step has then been run under the PID after Reset. Every row of Tables 6.6 and 6.7 uses one seed, sensor, amplifier, plant and duration, the PID at its default tuning (the challenge at both of its tunings) and the MPC at scenario 9's settings, with nothing retuned per case. The step was chosen from twelve candidate cases after both controllers had run on them, so the neighboring steps are reported with it. Table 6.6 adds the online solve, the PID retuned once by the MPC's own rule (the LQR's gap and rate gains as $K_c = 1.356\,\mathrm{A/mm}$ and $\tau_D = 17.6$ ms), the PID with its proportional term on the measurement (I-PD), the classical remedy for a setpoint step into a limit, and four counterfactual runs.

Under the MPC the command falls to 0 A by 1.005 s, stays there to 1.016 s and sits on the ceiling from 1.029 to 1.097 s, while the ball sinks to 28.52 mm and is within 0.1 mm of 26 mm from 1.167 s, an IAE of 0.319 mm s, as shown in Figure 6.5. In the window of the figure the table visits 14 regions and a constraint shapes 92 of its 351 plans, and the lookup and the online solve give the same current at every decision to 2.7×10^{-15} A with no miss. Under the PID, K_c times the 6 mm error removes 2.4 A from the new 2.31 A feedforward, and the command sits at 0 A from 1.005 to 1.007 s. It then climbs more slowly than the MPC's and reaches the ceiling at 1.037 s, 8 ms later, with 2.08 A in the coil against the MPC's 2.39 A, and capture is lost at 1.225 s. With the LQR's gains the PID's kick is 8.1 A and the ball is lost at 1.144 s.

Each run that removes one ingredient holds the ball. The same PID holds the step at 0.405 mm s with the ceiling raised to 6 A, and at 0.337 mm s with the step taken as a 50 ms ramp of setpoint moves, while a 25 ms ramp gives 1.356 mm s, a 10 ms ramp still loses the ball, and the LQR's gains on the 50 ms ramp give 0.191 mm s. The MPC's own law clipped into the box and the slew, with no look-ahead, holds at 0.325 mm s against the table's 0.319. The headline result is therefore the PID's proportional kick into the 3 A ceiling, met on the MPC side by state feedback with three times the PID's rate gain, and not the table's look-ahead, which decides only at the edge. A step to 26.5 mm is held at 0.637 mm s by the table and 0.689 by the clipped law, a step from 20 to 10 mm is held by the

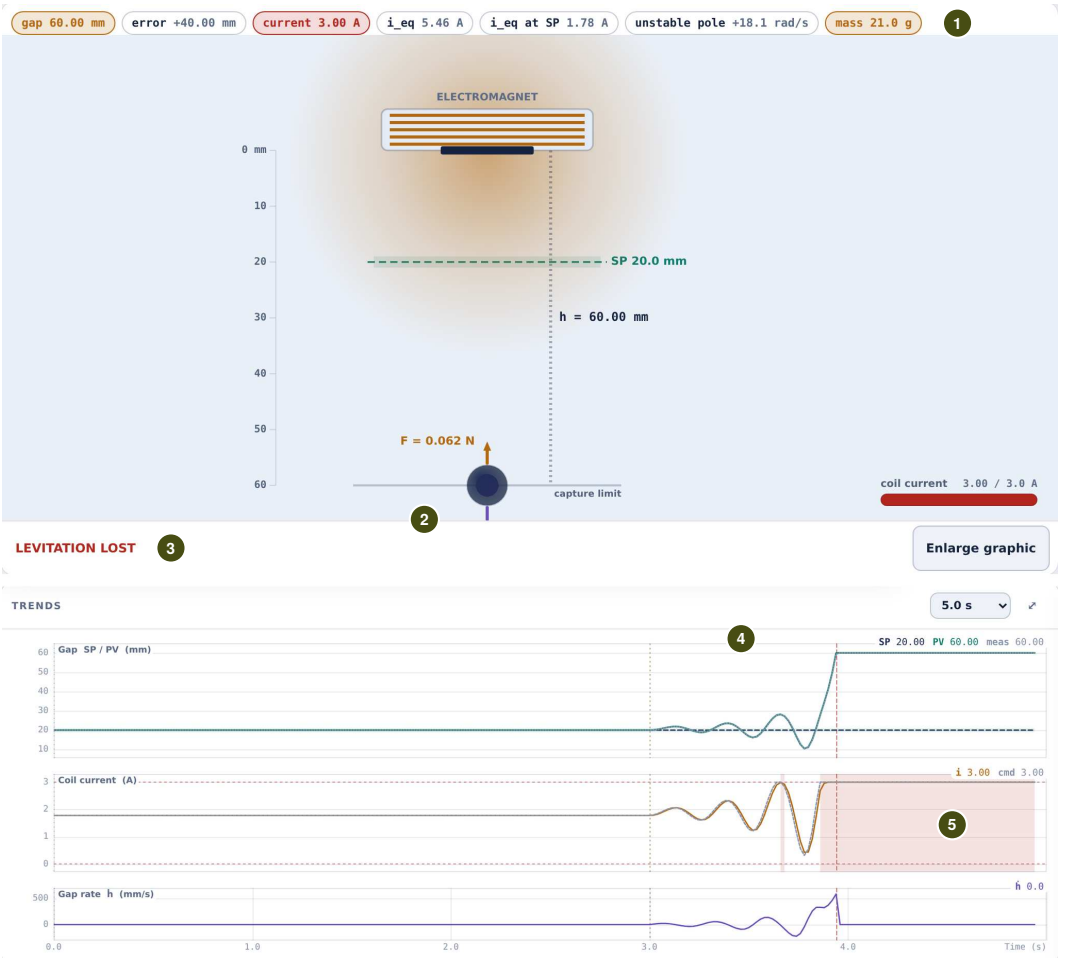


Figure 6.4 The graded scenario at $t = 4.95\,\mathrm{s}$, two seconds after a loop that had held a steady gap for three seconds diverged, in a composite of the screen: the plant view above the trends, with the header that prints the time, the control panel and the scorecard cut. The shipped tuning $K_c = 0.150\,\mathrm{A/mm}$, $\tau_I = 5\,\mathrm{s}$, $\tau_D = 6\,\mathrm{ms}$ satisfies the gain condition (6.17) and violates $\tau_D > \tau_c$ by 4 ms. (1) The badge row: the gap error reads 40 mm, and at the 60 mm radius holding station would need 5.46 A against a 3 A amplifier, with the open-loop pole slowed to $+18.1\,\mathrm{rad/s}$. (2) The ball at the capture limit, where a loss-of-capture fault is latched, and (3) its banner. (4) The gap strip is flat at the 20 mm setpoint until the $+5\,\%$ mass step at $t = 3\,\mathrm{s}$, after which the divergent $+3.18 \pm 25.0\,\mathrm{rad/s}$ mode grows visibly over three cycles and the ball departs on the fourth. (5) The current strip shows the same envelope reaching the dashed 3 A ceiling, with the shaded intervals marking the time spent against it. The gap-rate strip below it spikes as the ball departs, to the 0.59 m/s of the run's record, which the strip does not print, and the cut scorecard reports an accumulated current travel of 9.39 A and 19.8 % of the run at the ceiling.

table while the clipped law touches the magnet at 1.069 s and the PID at 1.077 s, and a step to 27 mm defeats all three. The ten-move online MPC, which the region budget excludes, holds that step at 1.06 mm s, and the I-PD holds every step from 8 to 27 mm, at 2.98 mm s on this one, nine times the table's IAE.

On the mass steps, the drops and the noise of Table 6.7 the table and the clipped law give identical

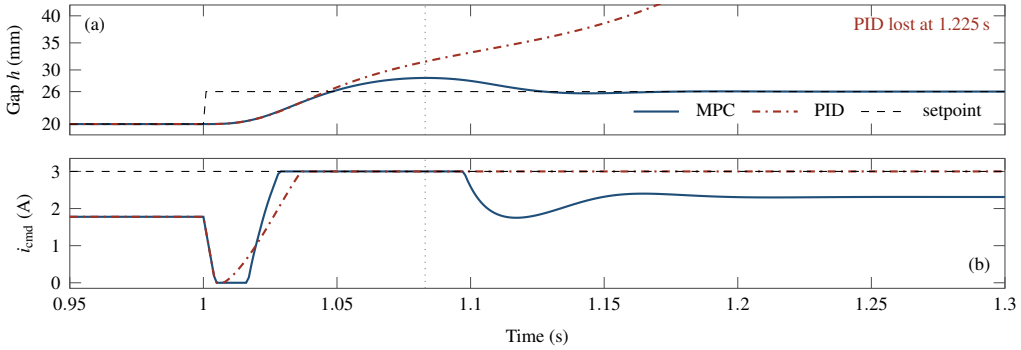


Figure 6.5 Scenario 9 under the explicit MPC and, after Reset, under the PID, from 0.95 to 1.30 s with the setpoint step at 1 s. (a) Gap. (b) Current command against the 3 A ceiling. The MPC is solid, the PID dash-dotted, and the setpoint and the ceiling are thin dashed lines. At the dotted line, $t = 1.083$ s, the MPC's ball is at its deepest, 28.52 mm, and the PID's at 31.53 mm and falling, both with the command on the ceiling.

Table 6.7 Disturbances and Losses, the PID at Its Default Tuning Against the Explicit MPC. IAE From the Event and Largest Gap Error

	IAE (mm s)		e_{max} (mm)	
Case	PID	MPC	PID	MPC
Challenge, +5 % mass[a]	0.055	0.006	0.19	0.15
Challenge, shipped PID tuning[b]	326.6	0.006	40.00	0.15
Scenario 4, +10 % mass	0.108	0.014	0.38	0.29
+10 % mass at 28 mm	0.152	0.019	0.55	0.34
1 mm drop at 20 mm	0.045	0.053	1.00	1.13
1 mm drop at 28 mm	0.059	0.044	1.00	1.00
Scenario 6, noise[c]	2.35	2.02	0.94	1.00
Drop, $R_i = 100$ A/s[b]	0.052	61.14	1.01	40.00
Drop, $m = 0.010$ kg	0.074	8.89	0.86	9.89
SP 28 to 8 mm[d]	2.09	1.49	20.00	20.00

IAE as in Table 6.6, so the PID's 0.045 after the drop at 20 mm is 0.033 on the scorecard of Table 6.4. [a]PID at the suggested tuning, 0.013 with the MPC's target current as its feedforward. [b]Ball lost, by the PID at 3.941 s and by the MPC at 1.640 s. [c]Over 9.3 s, current travel 114 against 212 A. [d]Both touch the magnet at 1.095 s and recover.

rows, so there the comparison is the LQR with its estimator against the PID. On the challenge's 5 % heavier ball the MPC reaches 0.006 mm s against 0.055 for the PID at the suggested tuning and 0.013 for the PID given the MPC's target as its feedforward, so the estimate does most of that work, and with the shipped tuning the PID loses the ball at 3.941 s. The MPC loses four cases. After a 1 mm drop at 20 mm its IAE is 0.053 against 0.045 mm s, with the command 10 ms on the ceiling and 8 ms at zero. Under the noise of scenario 6 its current travel is 212 against 114 A and its time on the ceiling 120 against 24 ms, for a lower IAE, a cost of the LQR's gains that the clipped law shares. After a 1 mm drop the PID holds at every slew tried down to 30 A/s, while the table holds at 120 A/s (0.288 mm s) and loses the ball at 100 and below, as the clipped law does, so that loss belongs to gains designed for the declared 400 A/s. A learner's 10 g ball is model mismatch, and after a drop it swings by 9.89 mm against the PID's 0.86 mm.

No scenario or hard case produces a miss, so 50 ms of misses were forced on scenario 9 with the noise off. At rest and on the ceiling the PID takes the coil within 25 mA of the MPC's command and the ball is held (0.319 to 0.394 mm s). Forced at 1.005 or 1.020 s the hand-over is bumpless, with first steps of 26 and 125 mA against a 400 mA slew step, but the PID, which loses this step by itself, loses the ball. The fallback is only as good as the shipped structure.

6.9 Numerical Implementation

The physics, the controller and the historian advance on separate periods, at 0.5 ms for classical fourth-order Runge-Kutta over the three states, 1 ms for the controller, and 50 samples per simulated second for the record. The engine of Section 1.6.1 keeps the computed trajectory independent of the display, so two machines produce identical exports. Within one step the controller and its sensor read run on every second step, and the hard-event correction follows the integration. The step size is chosen against the fastest element in the model rather than the slowest, because at the 8 mm setpoint limit the unstable pole reaches 49.5 rad/s while the coil contributes another at 100 rad/s. Among the checks that run with the application's test suite, four test the plant and the PID, on the algebraic equilibrium, on a coil current that never exceeds $i_{\max}$ in any scenario, on the sign of the force response to a current increase, and on the sign of the controller response to a growing gap. Further checks cover the feedforward's use of the declared constants and the explicit MPC.

The manual mode supplies the equilibrium current on request, so the existence of an equilibrium can be separated from its stability, and each run exports at 50 Hz for offline analysis.

6.10 Technology in Practice

What Measures It

The chapter's sensor is a light beam across the gap, read as the true gap plus colored noise. An industrial magnetic bearing measures its rotor with displacement probes instead, and a design study of a high-speed spindle lists what is fixed when they are chosen: measurement range, bandwidth, resolution, temperature rating and sensitivity to the magnetic field. The bandwidth caps the frequency up to which the position is measured accurately, because attenuation and phase lag grow as vibration approaches it, and the resolution caps the accuracy the spindle can reach. That spindle chose eddy-current probes whose response lies above its tooth-passing frequency [25]. A probe's lag would add to the filter and sample lags that already lift the simulator's stability boundaries above the predicted ones (Section 6.6).

> **TECHNOLOGY** **Sensor: An Infrared Beam the Ball Partly Blocks**
>
> One teaching rig finds its ball with an infrared beam that starts at a light-emitting diode and ends at a photodiode. The ball partly blocks the beam, so the light that reaches the detector, and the amplitude of its signal, follows the ball's position. The diode is pulsed 10,000 times a second, and an infrared beam rather than a visible one reduces disturbances from room light. With the ball levitating, the rig's filtered position voltage still carried about ±0.05 V of noise on 1.6 V [26]. The chapter's sensor is that beam with colored noise added and no bandwidth limit, drift or gap-dependent gain.

What Moves It

In an industrial bearing, sensors monitor the rotor's position, the signal is compared with a reference, and the controller commands power amplifiers that feed the electromagnets. Each degree of freedom needs two electromagnets, because they can only attract [20]. The chapter's single magnet has gravity as its second. Auxiliary bearings with a smaller clearance than the magnets carry the rotor at rest and let it decelerate safely if the main and redundant features fail together [20]. A high-friction contact with such a bearing can grow into a destructive backward whirl, with contact forces as high as 300 times the rotor's weight [25]. The chapter's pole face is a perfectly inelastic stop instead.

Where the Controller Runs

On one industrial compressor an analog PID loop set the bearing's stiffness and damping: the damping by passive elements of the compensation circuit, the gain by a potentiometer [20]. The teaching rig's analog proportional-derivative controller is built from operational amplifiers, and digital control is offered beside it as a quick way to test whether a set of components can levitate at all [26]. On an attractive-suspension vehicle individually controlled direct-current magnets in the undercarriage hold

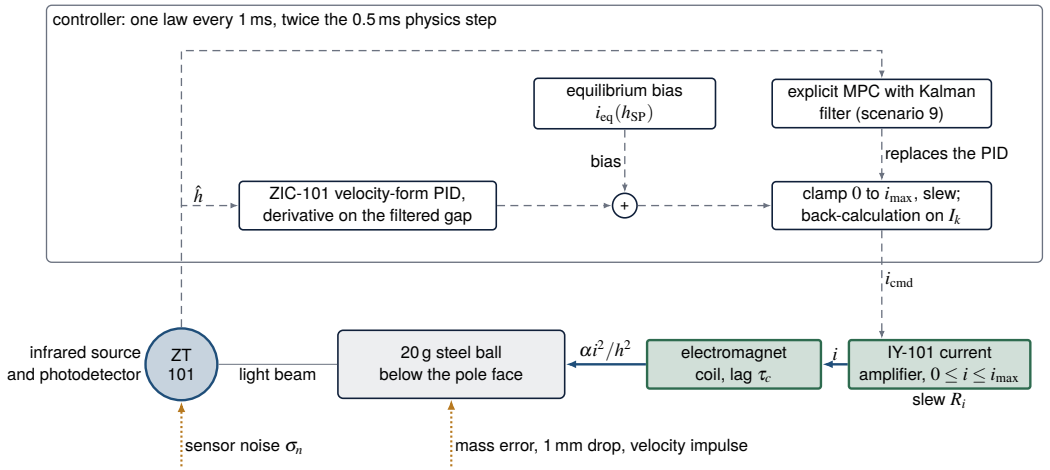


Figure 6.6 The suspension loop on real hardware. An infrared beam partly blocked by the ball, ZT-101, reports the gap $\hat{h}$ to the velocity-form PID ZIC-101, which takes its derivative on the filtered measurement. The equilibrium current at the setpoint enters as a bias, the demand is clamped to the amplifier's range and slew, and back-calculation corrects the integral alone. The current amplifier IY-101 drives the coil, whose lag τ_c the derivative term has to outrun. In scenario 9 the explicit MPC with its Kalman filter replaces the PID. Notice that only the gap is measured: under the PID the ball's speed reaches the loop through the derivative term alone.

the gap, fed through pulse-width-modulated power supplies, and the survey lists complex control systems for suspension and guidance among the drawbacks [19]. The chapter's PID runs every 1 ms, twice the physics step, in velocity form with the balancing current as its bias and back-calculation on the integral alone. Figure 6.6 draws the loop on hardware.

> **TECHNOLOGY** **Final Element: A Power Stage and a Coil That Lags**
>
> On the same rig operational amplifiers drive a power transistor that feeds the electromagnet from a 30 V, 1 A supply, and the ball levitated on 470 mA at an air gap of 7 mm. The coil's time constant, its inductance over its resistance, sets how fast the control signal can change the field in the gap, and a loop on the coil current would improve that response, which the rig did without [26]. Industrial bearings use switching amplifiers, which moved from two output states to three because the amplifier is then more efficient and the bearing has lower eddy-current losses and less acoustic noise [27]. The chapter's amplifier is a current loop of that kind: a lag τ_c, a slew limit and a ceiling.

What Is Modernizing It

Self-sensing removes the probe: the bearing's own coils report the gap through their inductance, and the method has reached turbomolecular pumps and elevator guideways [27]. Design frameworks now carry a bearing spindle from its requirements through the magnets, the backup bearings, the instruments and the control hardware to its assembly, because design knowledge had stayed fragmented across isolated prototypes [25]. The explicit predictive controller of Section 6.5.1, a stored table looked up in place of an online solve, is the chapter's own example of fitting optimization into a kilohertz loop.

IN PRACTICE **Proportional First, Then Derivative**

The teaching rig's analog controller is set on two potentiometers. The proportional gain is set first, with the derivative gain low, while the ball is held at the desired place, and too little gain cannot lift it. With enough proportional gain but too little derivative gain the ball may fly into the electromagnet and bounce off, and more derivative gain stops it. A levitating ball that vibrates also asks for more derivative gain. Settings found before the magnet and the power transistor warm up may need a final adjustment [26]. Condition (6.18) explains the order: the derivative term is what holds the ball up.

6.11 Limitations

This simulator teaches a loop and does not size a magnet. Its single degree of freedom excludes the guidance and lateral-stability problems that dominate real suspension vehicles [9], and its lumped force coefficient excludes the magnetic design questions that determine α [3]. The controllers are a fixed-structure PID and a three-move explicit MPC on one linearization, so the feedback-linearization and adaptive designs this plant motivated [13–16] cannot be demonstrated, although Table 6.3 showing one tuning producing three different closed-loop responses is exactly the evidence that makes those designs look necessary. Every disturbance except the sensor noise is applied on command rather than arising as stochastic plant behavior. A fixed integration step and a fixed control period give bit-identical repeats, which no field controller delivers. Section 6.3.4 states where each modeling assumption breaks down, and the three that bind hardest are the constant force coefficient at large gap, the absence of voltage saturation in the amplifier, and the continuous-time basis of the stability conditions.

The explicit MPC's model keeps the slopes of the 20 mm design gap at every setpoint, so the $1/h^2$ curvature enters only through the target of Equation (6.15). Its estimator treats the colored sensor noise as white, has no pole face or capture radius, and knows only the declared ball, magnet, gravity and coil lag, so a changed plant is absorbed at steady state but not in a fast transient, as the 10 g ball shows. Three 1 ms moves are the longest horizon whose table stays under 100 regions, too short to anticipate a slow amplifier, and a state outside the tabulated domain hands the coil to a PID that may not hold the ball.

OPEN CHALLENGE **Sensing the Gap With the Coil Itself**

A self-sensing magnetic bearing drops the position probe and estimates the gap from its coils' voltage and current, because a coil's inductance changes with the gap. A review found it in commercial use, with hurdles left for demanding products. Its robustness hinges on the size of the switching ripple, which three-state amplifiers cut, sometimes about tenfold, so a robust self-sensed bearing is likely to use somewhat more electrical power than one with a probe. Eddy currents and saturation of the iron also blur the coil's sensitivity to the gap [27]. Progress would be self-sensing that is both efficient and robust.

6.12 Discussion

Removing contact removes wear, lubrication, friction loss and a maintenance schedule, which is why active magnetic bearings hold rotors in turbomolecular pumps, gas compressors and flywheels, and why those rotors turn faster than any bearing that touches them permits [8]. The price is that the machine has no passive resting state. A power failure in an oil-film bearing is an interruption, while a power failure in a magnetic bearing drops the rotor onto its auxiliary bearings, and much of the design discipline of that field follows from the asymmetry. The same argument scales to electromagnetic-suspension vehicles, where the gap under a guideway rail is a few millimeters and the loop that maintains it carries the vehicle [9, 12].

The transferable lesson is broader than magnets. A right-half-plane pole at p imposes a floor rather than a preference, because the sensitivity integral is fixed at πp instead of zero, so the closed-loop bandwidth must exceed the instability and the disturbance amplification that appears elsewhere in frequency cannot be tuned away but only relocated [17]. Such plants must be respected rather than

merely tuned, because a control failure on them is qualitatively different from a poorly damped response [18]. Aircraft built with relaxed static stability, reactors held on an unstable branch for yield, the elongated plasma of a tokamak (Volume 18, Chapter 12) and this suspension all share that structure, and a stability margin on such a plant bounds the time available to act rather than the quality of a response.

6.13 Conclusion

An electromagnet and a steel ball make a plant with an equilibrium at every gap and stability at none, and that property reorders what a PID loop is understood to be for. The simulation of this chapter integrates three states under an inverse-square force with a first-order current amplifier, a 3 A ceiling, a 400 A/s slew limit and a gap sensor with a 4 ms correlation time, and closes a kilohertz digital PID in velocity form with an equilibrium feedforward bias and back-calculation anti-windup on the integral alone. Linearization about a 20 mm gap gives an open-loop pole at $+31.3\,\mathrm{rad/s}$, and Routh analysis of the resulting fourth-order loop gives three inequalities that decide whether the ball floats. The simulator brackets the predicted critical gain of 107.0 A/m between 105 and 110 A/m and the predicted critical derivative time of 10.30 ms between 10 and 11 ms. The order of those inequalities is the teaching point. Proportional gain above the magnetic stiffness of 88.9 A/m only converts a runaway into an undamped oscillation, integral action cannot stabilize the plant under any setting whatsoever, and derivative action longer than the 10 ms coil time constant is what holds the ball up, bounded from above by the amplifier ceiling and by the sensor noise it differentiates. An explicit MPC of 91 regions holds a step into the 3 A ceiling that the PID loses (0.319 mm s against a lost ball), although the same law clipped without look-ahead holds it at 0.325 mm s, so the PID's proportional kick, not the optimization, decides that case. A fuller model would add a gap-dependent coil inductance with a voltage limit, an eddy-current drag term, a second translational axis for lateral stability, and a sampled-data version of the three conditions. The suspension is an accessible browser platform for studying a loop that manufactures its own operating point.

KEY POINTS

- A steel ball below an electromagnet has a balancing current at every gap and is stable at none, so even the exact current loses the ball.
- Proportional gain alone gives an undamped oscillation at best, and integral action cannot stabilize the plant.
- Derivative action longer than the 10 ms coil lag holds the ball up at any gain.
- The amplifier ceiling and the sensor noise that derivative action amplifies bound the gain from above.

Where else this applies. Volume 18, Chapter 12 holds a plasma with the same open-loop instability.

References

[1] T. H. Wong, "Design of a magnetic levitation control system—an undergraduate project," *IEEE Transactions on Education*, vol. E-29, no. 4, pp. 196–200, 1986.

[2] P. S. Shiakolas and D. Piyabongkarn, "Development of a real-time digital control system with a hardware-in-the-loop magnetic levitation device for reinforcement of controls education," *IEEE Transactions on Education*, vol. 46, no. 1, pp. 79–87, 2003.

[3] W. G. Hurley and W. H. Wolfle, "Electromagnetic design of a magnetic suspension system," *IEEE Transactions on Education*, vol. 40, no. 2, pp. 124–130, 1997.

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

[5] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[6] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

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

[8] E. H. Maslen and G. Schweitzer, Eds., *Magnetic Bearings: Theory, Design, and Application to Rotating Machinery*. Berlin, Heidelberg: Springer, 2009.

[9] H.-W. Lee, K.-C. Kim, and J. Lee, "Review of maglev train technologies," *IEEE Transactions on Magnetics*, vol. 42, no. 7, pp. 1917–1925, 2006.

[10] W. Braunbek, "Freischwebende Körper im elektrischen und magnetischen feld," *Zeitschrift für Physik*, vol. 112, no. 11-12, pp. 753–763, 1939.

[11] M. H. Tuttle, R. A. Kilgore, and R. P. Boyden, "Magnetic suspension and balance systems: A selected, annotated bibliography," NASA Langley Research Center, Hampton, Virginia, NASA Technical Memorandum TM-84661, Jul. 1983. [Online]. Available: https://ntrs.nasa.gov/citations/20080047362

[12] R. D. Thornton, "Efficient and affordable maglev opportunities in the United States," *Proceedings of the IEEE*, vol. 97, no. 11, pp. 1901–1921, 2009.

[13] W. Barie and J. Chiasson, "Linear and nonlinear state-space controllers for magnetic levitation," *International Journal of Systems Science*, vol. 27, no. 11, pp. 1153–1163, 1996.

[14] D. L. Trumper, S. M. Olson, and P. K. Subrahmanyan, "Linearizing control of magnetic suspension systems," *IEEE Transactions on Control Systems Technology*, vol. 5, no. 4, pp. 427–438, 1997.

[15] A. El Hajjaji and M. Ouladsine, "Modeling and nonlinear control of magnetic levitation systems," *IEEE Transactions on Industrial Electronics*, vol. 48, no. 4, pp. 831–838, 2001.

[16] Z.-J. Yang and M. Tateishi, "Adaptive robust nonlinear control of a magnetic levitation system," *Automatica*, vol. 37, no. 7, pp. 1125–1131, 2001.

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

[18] G. Stein, "Respect the unstable," *IEEE Control Systems Magazine*, vol. 23, no. 4, pp. 12–25, 2003.

[19] J. L. He, D. M. Rote, and H. T. Coffey, "Survey of foreign maglev systems," Argonne National Laboratory, Center for Transportation Research, Argonne, Illinois, Tech. Rep. ANL/ESD-17, Jul. 1992. [Online]. Available: https://www.osti.gov/servlets/purl/10134413

[20] R. G. Kirk, J. F. Hustak, and K. A. Schoeneck, "Active magnetic bearings applied to industrial compressors," in *Magnetic Suspension Technology Workshop*, ser. NASA Conference Publication, no. 3202. Hampton, Virginia: NASA Langley Research Center, 1993. [Online]. Available: https://ntrs.nasa.gov/citations/19930018381

[21] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[22] A. Alessio and A. Bemporad, "A survey on explicit model predictive control," in *Nonlinear Model Predictive Control*, ser. Lecture Notes in Control and Information Sciences, L. Magni, D. M. Raimondo, and F. Allgöwer, Eds. Berlin, Heidelberg: Springer, 2009, pp. 345–369.

[23] 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.

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

[25] K. S. Ahmed and B. Bediz, "Design framework and manufacturing of an active magnetic bearing spindle for micro-milling applications," arXiv:2603.00169v2 [eess.SY], 2026.

[26] S. Barker and R. C. Matusiak, "A technology approach to magnetic levitation," in *2000 ASEE Annual Conference & Exposition*, St. Louis, Missouri, 2000.

[27] E. H. Maslen, "Self-sensing for active magnetic bearings: Overview and status," in *Proceedings of ISMB10, International Symposium on Magnetic Bearings*, 2006. [Online]. Available: https://www.magneticbearings.org/app/uploads/publications/ismb10/ismb10_submission_key03.pdf

6.14 Exercises

6.1 Concept check · 1 min

The suspension has one loop. Which pair names what the controller reads and what it adjusts?

(a) It reads the coil current and adjusts the gap setpoint
(b) It reads the measured gap and adjusts the coil current command
(c) It reads the ball mass and adjusts the force coefficient α
(d) It reads the magnetic force and adjusts the amplifier ceiling

6.2 Concept check · 2 min

Scenario 1 starts the ball at 20 mm with the balancing current. The reader switches to Manual before Run and touches nothing. What happens?

(a) The ball stays at 20 mm, because the net force on it is zero
(b) The ball drifts to a nearby gap where the magnetic force again balances
(c) The ball leaves within a second or so, because any error grows by e every 32 ms
(d) The ball swings about 20 mm, because the magnet pulls back like a spring

6.3 Concept check · 2 min

The challenge ships $K_c = 0.150\,\mathrm{A/mm}$, $\tau_I = 5\,\mathrm{s}$ and $\tau_D = 0.006\,\mathrm{s}$, with the coil time constant $\tau_c = 10\,\mathrm{ms}$, and the ball is lost after the 5 % mass step. Which single change keeps the ball?

(a) Raise K_c to 1.000 A/mm, far above the floor of 0.089 A/mm
(b) Shorten τ_I to 0.50 s so the integral removes the sag sooner
(c) Raise the derivative filter to 0.010 s to quiet the derivative term
(d) Raise τ_D to 0.020 s, above the coil time constant

6.4 Theory and calculation · 6 min

Use the declared ball and magnet: $m = 0.020\,\mathrm{kg}$, $\alpha = 2.4832 \times 10^{-5}\,\mathrm{N\,m^2/A^2}$ and $g = 9.81\,\mathrm{m/s^2}$.

(a) Find $K_{\min} = \sqrt{mg/\alpha}$ and the balancing current at gaps of 20 and 28 mm. *Checkpoint:* 88.9 A/m, 1.778 A and 2.489 A.

(b) Find the unstable pole p, its time constant $1/p$ and the control effectiveness B at both gaps, and show that p^2/B is the same at both. *Checkpoint:* 31.3 rad/s, 31.9 ms, 11.04 m/(s^2 A) at 20 mm, and 26.5 rad/s, 37.8 ms, 7.88 m/(s^2 A) at 28 mm. $p^2/B = 88.9\,\mathrm{A/m}$ at both.

(c) At the stored manual current the ball starts 1.4×10^{-13} m from its equilibrium at rest, so the linear model gives $\delta h = \delta h_0 \cosh pt$. How long does it take to cover the 40 mm to the capture radius at 20 mm, and by what factor does an error that is already growing, $\delta h \propto \mathrm{e}^{pt}$, grow during a 250 ms human reaction? *Checkpoint:* 0.86 s, and about 2500 ($pt = 7.8$).

6.5 Theory and calculation · 8 min

The closed loop of the PID on the linearized plant has the characteristic polynomial (6.16). Work at 20 mm with $p^2 = 981/\mathrm{s^2}$, $B = 11.04\,\mathrm{m/(s^2\,A)}$ and $\tau_c = 10\,\mathrm{ms}$.

(a) Set $\tau_c = 0$ and $\tau_D = 0$. Write the polynomial for proportional action alone, with $K_c = 400\,\mathrm{A/m}$, and for PI action. What does each loop do? *Checkpoint:* $s^2 + (BK_c - p^2) = 0$, an undamped oscillation at 58.6 rad/s. The PI polynomial has no s^2 term and is unstable for every setting.

(b) With $\tau_I = 0.5\,\mathrm{s}$ and $\tau_D = 20\,\mathrm{ms}$, find the smallest K_c that satisfies condition (6.19), and the size of the budget at the default $K_c = 400\,\mathrm{A/m}$. *Checkpoint:* 107.0 A/m, and 34.3/s against the 2/s needed.

(c) With $K_c = 150\,\mathrm{A/m}$ and $\tau_I = 5\,\mathrm{s}$, find the smallest derivative time. *Checkpoint:* 10.30 ms.

6.6 Theory and calculation · 6 min

The amplifier holds the current between 0 and $i_{\max} = 3.0\,\mathrm{A}$ and slews at most $R_i = 400\,\mathrm{A/s}$. Use $K_{\min} = 88.9\,\mathrm{A/m}$ and the default tuning $K_c = 400\,\mathrm{A/m}$, $\tau_I = 0.5\,\mathrm{s}$, $\tau_D = 20\,\mathrm{ms}$, $\tau_f = \Delta t_c = 1\,\mathrm{ms}$.

(a) How far below its setpoint can the ball sit at rest and still be lifted by the full current, at 20 and at 28 mm? What is the highest gap a 2.4 A amplifier can hold? *Checkpoint:* 13.75 mm and 5.75 mm (both to 33.75 mm), and 27.0 mm.

(b) A 1 mm drop lands between two scans. On the next scan the filtered measurement moves by $(1 - \mathrm{e}^{-1})\,1\,\mathrm{mm}$. Find the current the PID demands, c_k, at 20 and 28 mm, starting from the balancing current with $I = 0$. *Checkpoint:* 7.23 A and 7.95 A.

(c) The slew limits the applied current to $i_0 + R_i \Delta t_c$. Find the excess $c_k - u_k$ and the change it makes in I_k through Equation (6.13), against a recursion that absorbs the whole excess in one scan. *Checkpoint:* 5.06 A at both gaps, and $-0.0101\,\mathrm{A}$ against $-5.06\,\mathrm{A}$.

6.7 App challenge · 7 min

Load scenario 2, +1 mm displacement (dropped), which opens in Manual with the balancing current for 20 mm and the ball 1 mm lower. Set the speed to 0.01, so 0.01 s of simulated time passes per second, and use the default seed (the scenario has no noise). Read times and gaps from the gap trend.

(a) Run until LEVITATION LOST is marked and read the time of the mark. Then load scenario 3, −1 mm displacement (pulled up), and do the same for MAGNET CONTACT. *Checkpoint:* 0.17 s and 0.10 s (±0.01 s).

(b) Reload scenario 2 and press Run. Switch to Auto when the clock reads between 0.04 and 0.06 s, and let it run another 1 s. Read the gap at the switch and the Peak current tile. *Checkpoint:* The ball is caught. Gap 21.8 to 23.1 mm, peak 2.4 to 2.9 A.

(c) Reload and repeat with the switch at 0.09 s. *Checkpoint:* The ball is lost. The gap at the switch is about 27 mm.

6.8 App challenge · 7 min

Use the default tuning, the default seed and speed 1. In each run press Run, fire Drop ball 1 mm at any time between 0.5 and 1.5 s, and pause at about 4.7 s. Read the Peak current and MV travel tiles.

(a) Scenario 8, Low levitation position (11 mm). *Checkpoint:* Peak 1.27 A, MV travel 1.46 A.

(b) Scenario 1, Equilibrium without stability, which holds 20 mm in Auto, and scenario 7, High levitation position (28 mm). *Checkpoint:* 2.08 A and 1.25 A, then 2.78 A and 1.44 A.

(c) Using the balancing currents $i_0 = K_{\min} h_{\mathrm{SP}}$, find each peak as a share of the 3 A ceiling and the current the loop added above i_0. What uses up the headroom? *Checkpoint:* 42, 69 and 93 %. Added: 0.29, 0.30 and 0.29 A.

6.9 App challenge · 8 min

Load the challenge scenario, Challenge, reject a 5% mass increase, and set the speed to 2 (a 12 s run takes 6 s) and use the default seed. Make each setting before Run, run to 12 s and read the challenge's result line (worst gap error, time saturated, peak current). Reload before each part.

(a) Run as shipped. Then set Proportional gain to 1.000 A/mm and run again. Read the time of LEVITATION LOST in each run. *Checkpoint:* Lost at 3.94 s, then lost at 3.68 s.

(b) Set only Derivative time to 0.020 s. *Checkpoint:* Held: 1.01 mm, 0.0 %, 1.94 A. The target fails by 0.01 mm.

(c) Set Derivative time to 0.020 s and Integral time to 0.50 s. Then press Reveal suggested tuning, Reset and run once more. *Checkpoint:* 0.92 mm, 0.0 %, 1.94 A, met. Suggested: 0.18 mm, 0.0 %, 1.87 A.

6.10 Application area · 4 min

A turbomolecular pump runs its rotor on active magnetic bearings.

(a) Compare what a power failure does to an oil-film bearing and to a magnetic bearing, and name the equipment the magnetic bearing needs as a result. *Checkpoint:* An interruption against a dropped rotor, and auxiliary (touchdown) bearings.

(b) The chapter's loop runs at 1 kHz, and the best derivative filter is near $\tau_D/10 = 2\,\text{ms}$. Why do physical suspensions run their loops at 10 kHz? *Checkpoint:* The 1 ms sample period is within a factor of two of the filter, and both lags sit in the path that holds the ball.

6.11 Application area · 4 min

The explicit MPC of the chapter stores 91 regions for three 1 ms moves. One to five moves give 5, 23, 91, 283 and 710 regions. In Node.js a lookup takes at most 0.93 µs and an online solve of the same program a median of 7.6 µs.

(a) A bearing controller must decide within 100 µs on a processor 50 times slower than that host. Which evaluation fits? *Checkpoint:* The lookup, at about 47 µs. The online solve needs about 380 µs.

(b) What does the short horizon cost on this plant? Give two cases from the chapter. *Checkpoint:* A slow amplifier (lost at 100 A/s) and the step to 27 mm, which a ten-move online MPC holds.

6.12 Review the agent · 5 min

For scenario 6, Noisy position sensor ($\sigma_n = 0.15\,\text{mm}$), an AI assistant writes: *"The coil current chatters by 0.44 A because the derivative term amplifies the noise. Raise the derivative filter from 0.001 s to $\tau_D/2 = 0.010\,\text{s}$. That cuts the noise bandwidth tenfold and the MV travel with it, and since the derivative only adds damping, the gap will barely change."* Find the flaw, and say how you would check it in the app.