INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 15

Summary and Outlook

Pages 366–380

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," 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. 15, pp. 366–380.

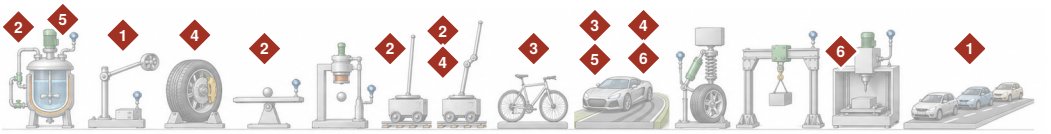
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 15

Summary and Outlook



Open gaps 1 margins 2 regions 3 scheduling 4 estimation 5 equilibria 6 models

Notice that the drifting car carries four of the six gaps, and that the gap of recoverable regions joins the reactor, the ball and beam and both carts. Illustration, not to scale.

IN PLAIN TERMS **What Thirteen Unsteady Machines Taught**

The main idea. Each chapter of this book held a machine that would fall, spin or run away on its own, or that kept swinging after it was moved. A controller had to watch it and push it back, many times a second.

The challenge. How far each machine could be pushed before it was lost was set by the machine, its motor and its sensors, not only by the controller's settings. Past that edge the loss came suddenly, and the usual checks did not always see it coming.

Why it matters. Measuring that edge on the part that fails tells a safe design from one that only looks safe.

Thirteen chapters held thirteen plants that would not hold themselves, or that rang when they were moved. Chapter 1 asked two questions of each plant. What makes it hard, and what holds it by how much? Table 1.2 answered the first and named the structure and the margin of each chapter without numbers. This chapter adds the numbers and the structures that failed, draws a short guide for choosing a stabilizing design, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

15.1 What the Chapters Showed

Tables 15.1 and 15.2 compare the chapters. Each row names the difficulty, the structure that held the plant, a structure or setting that failed, and the margin with the table, figure or section of the chapter that reports it. The first table holds the plants whose gain changes with the operating point and the plants with one unstable mode. The second holds the plants with two unstable modes or a speed that moves the poles, the oscillating plants and the line of cars. Figure 15.1 sets one like pair of numbers from nine of the rows side by side.

The columns of the two tables make five points.

The margin ends at a cliff. Most of the measured margins are thin, and the loss past them is sudden. One kelvin of setpoint separated an overshoot of 12.18 K from one of 69.07 K in Chapter 2. One millisecond of sample interval separated a held double pendulum from a fallen one. At 0 % comfort the leveling loop of Chapter 11 decays at 11 N/mm and grows at 12 N/mm. The line of cars of Chapter 14 held at a reaction delay of 0.5 s and collided at 0.8 s. A stabilizing loop has no gentle region of poor performance between good and lost, so the distance to the edge is the number to report.

The usual margin did not always see the edge. The aeropendulum's PID kept a phase margin of 49 to 55° up to 105° and went unstable at 105.8°, because the mode that failed was a slow one the margin does not describe. The bicycle's schedule was stable by a continuous analysis and grew at up

Table 15.1 What Held Each Plant and by How Much, Chapters 2 to 7

Chapter	Difficulty	What held it	What failed	Margin measured
2, stirred reactor	Every temperature from 335.654 to 379.611 K is open-loop unstable, the 340 K setpoint included.	Reverse-acting PI from a step-test fit, $K_c = 4$, $\tau_I = 1.333$ min. Predictive control on the linearization.	A setpoint step 1 K larger. A lower gain of 1.6, which raised the peak and left a relaxation cycle.	An 11 K step overshoots 12.18 K, a 12 K step 69.07 K and trips (Table 2.5). Largest safe step 11.5 under PI, 13.5 under MPC (Section 2.8).
3, propeller arm	Gravity's stiffness changes sign at horizontal. At 120° the arm diverges at 4.06/s.	Gravity feedforward, a gain schedule, an integral zone and a derivative filter, all four together.	One PID tuned at 30°. Its slow mode is unstable from 105.8° with a phase margin of 49 to 55°.	Program settled in 2.26 s with all four, never without the schedule or the feedforward (Table 3.6). Phase margin down to 19.1° at 150° (Figure 3.2).
4, braked wheel	Slip held at the friction peak, where the gain vanishes. Equilibrium lost at $\lambda_c = 0.1661$.	PI on slip. A sliding-mode law with its boundary layer sized to the brake lag.	A locked wheel, 47.9 m. A rule-based cycle, locking in 201 of 348 runs.	Stop 33.0 m against 31.8 m ideal (Table 4.6). Phase margin 40.6 at 27 m/s, 15.2° at 3 m/s (Table 4.4). Locks in 60 of 300 runs against 80 of 336 (Table 4.9).
5, ball and beam	Two integrators and a servo lag, no damping.	A filtered derivative. A PID on a profile with feedforward.	Proportional alone, each swing 1.37 to 1.42 times the last. An unfiltered derivative, ball flung 66 to 94 mm.	Phase margin 48.5° (Table 5.6), 14.4° lost to the camera's delay (Section 5.6). Relay race 0 arrivals against 15 (Table 5.8).
6, levitated ball	One real pole at 31.3 rad/s, a 10 ms coil, a 3 A ceiling.	A derivative time longer than the coil's, with an equilibrium feedforward.	Integral action, at any setting. A gain above the magnetic stiffness of 88.9 A/m with too little derivative.	Critical gain 107.0 A/m, found between 105 and 110. Critical derivative time 10.30 ms (Section 6.6). A 1 mm drop at 28 mm uses 93 % of the ceiling (Section 6.7).
7, cart and pole	An unstable root of 5.58/s, one e-fold every 179 ms, a wrong-way zero, a rail.	A cascade, angle inside cart position. An LQR on the full state.	The angle loop alone, at the bumper in 4.11 s. The cascade on a 150 °/s push.	Gain window closed at about 105 ms of sensing lag (Figure 7.2). Largest push kept off the bumpers 130 for the cascade, 229 °/s for the LQR (Table 7.6).

to 0.157/s once the 10 ms hold was counted. On the line of cars the linear peak rose only from 1.017 to 1.084 over the range of driver delay in which the measured ratio reached 1.29 and the cars collided (Table 14.5). Table 1.1 showed the same disagreement on one pendulum. The margin worth quoting is the one measured on the mode and the path that fail.

Derivative information holds the plant, and its source limits it. The ball under the magnet and the ball on the beam were held by derivative action that had to act faster than a lag. The levitated ball needs a derivative time above the coil's 10 ms, and no integral setting replaces it. The same derivative differentiates noise. Unfiltered, it flung the ball on the beam 66 to 94 mm off its setpoint. On the double pendulum, rates differenced from encoder counts made the force chatter by 7.17 N, and a Kalman filter on the same counts left 0.43 N. The velocity a stabilizing loop uses is a design choice as much as its gain.

Table 15.2 What Held Each Plant and by How Much, Chapters 8 to 14

Chapter	Difficulty	What held it	What failed	Margin measured
8, double pendulum	Two unstable modes, the fast one an e-fold every 90 ms, and one motor.	An LQR on a Kalman filter.	A loop on each link. Its gains may sum to 12.26 N/rad, the lower link needs 58.6 N/rad. Differenced rates, 7.17 N of chatter against 0.43.	Longest stable sample interval 12.3 ms at a 30 ms drive lag (Figure 8.2). After a switch, 12 ms holds and 13 ms falls (Table 8.5).
9, bicycle	Self-stable only inside a speed window. Countersteering is a wrong-way zero.	A lean PD scheduled on speed. An LQR scheduled on speed.	Gains tuned at 2 m/s, unstable above 3.45 m/s once the 10 ms hold counts. Any of 123 designs at 0.3 m/s.	The PD schedule leaves a band from 2.81 to 4.29 m/s growing at up to 0.157/s (Figure 9.2). LQR roots below −2.76/s from 0.5 to 10 m/s (Section 9.5).
10, drifting car	Roots of 1.95 and 0.21/s at an equilibrium with the rear tire saturated.	A two-input LQR on steer and drive, scheduled on a friction estimate.	A grip controller, spun at 3.61 s. The LQR scheduled for the dry pad, spun on a wet patch at 10.06 s.	Untouched, 5° off in 2.98 s (Section 10.6). A 10° kick back within 2° in 0.92 s (Table 10.4).
11, quarter car	One force on two masses fixes the body at 11.25 Hz, near the 11.9 Hz wheel hop.	A skyhook and groundhook blend on an active strut. Road preview.	A semi-active strut set for comfort, off the road 60 times in 5 s. Any strut on a pothole edge, 1579 m/s^2 asked of 98.1 m/s^2.	Hump rise 36.7 against 47.7 mm, 5.9 with preview (Table 11.4). Leveling gain decays at 11, grows at 12 N/mm (Section 11.5).
12, gantry crane	A swing with a damping ratio of 0.0579 and a frequency set by the cable. One drive.	Two or three impulses shaping the reference of a PID position loop.	A shaper set for 4 m on an 8 m cable. A saturated drive, 7.7° with shaping on or off.	Residual swing 7.87 unshaped, 0.30° with three impulses (Table 12.4). Mistuned, 44.6 % of unshaped, 43.7 % predicted (Figure 12.2).
13, end mill	A regenerative delay turns an 800 Hz tool mode into chatter.	An adaptive feed for the load, a chatter watch that moves the spindle speed.	The feed loop alone on a worn tool, load within 1.6 % and 72 mm of marks.	Limit 2.10 mm, 1.5 mm quiet, 3 mm chatters at 823 Hz. Pocket 12.63 against 16.36 s with 26 mm marked (Table 13.4).
14, line of cars	Each car stable, the line amplifying. A radar latency and a driveline lag.	A time gap above $\sqrt{2/k_p}$, 1.05 s. The predecessor's acceleration over a radio.	Constant spacing, best of 90 gain pairs 1.050. Radar only at 0.6 s, peak 1.306, car 10 down to 36.3 km/h.	Peak 1.000 at 1920 vehicles an hour by radar, at 3692 by radio (Table 14.5). Radio boundary 0.737 s at 400 ms and 20 % loss (Figure 14.2).

Structure decides before tuning on underactuated plants. The cart's angle loop alone balanced the pole and drove the cart into a bumper. On the double pendulum no pair of loops can hold the plant at any tuning, because the sum the cart allows is smaller than the gain the lower link needs. The drift's grip controller spins the car at every request of 10° or more. A full-state regulator held all three, and on the cart it kept a push of 229 °/s off the bumpers where the cascade managed 130. When one actuator serves two outputs, the first question is which structure can hold the plant at all.

Some limits belong to the plant, and only a change to the plant or its command moves them. No strut follows a pothole edge that asks 1579 m/s^2 of a wheel with 98.1 m/s^2 available from its spring and weight. No tuning stops a flicked ball on a beam limited to 3°. No feed setting moves the stability lobe of Chapter 13, while the spindle speed does. The crane's swing was removed by

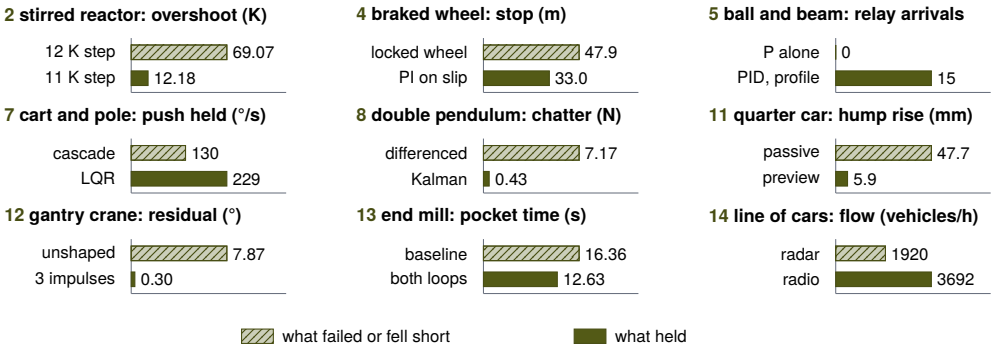


Figure 15.1 One like pair of numbers from nine rows of Tables 15.1 and 15.2, digit for digit, each panel on its own scale. The hatched bar is the setting or structure that failed or fell short, the solid bar the one that held. Higher is better on the ball and beam, the cart and pole and the line of cars, lower on the other six. The aeropendulum, the levitated ball, the bicycle and the drifting car have no like pair in their rows, because each reports an edge or a single margin. Notice that every better number came with a change of structure, command or setting, not with a retuned gain alone.

shaping the command, not by feedback, and a saturated drive gave the same 7.7° with shaping on or off. These are the limits of Section 1.4.7 in their physical form. The engineer's lever is the command, the operating point or the equipment.

15.2 Choosing a Method

The chapters suggest an order of questions for a plant that must be held or kept from ringing. Each names the chapter where its answer decided the result.

1. *Is the operating point stable, and is it the only one?* Find the steady states and the eigenvalues at the chosen one before choosing a controller (Section 1.4.2). The reactor's setpoint sits on an unstable branch, and a stable-looking hot state was an unstable focus.

2. *Can the actuator turn the unstable mode back?* Compare the largest disturbance with $u_{\max}/p$ and check $p\,\tau_d < 1$ (Section 1.4.7). The levitated ball at 28 mm spent 93 % of its current ceiling on a 1 mm drop.

3. *Is a lag inside the loop longer than the derivative time?* Then no gain helps (Equation (1.9)). Shorten the lag, lengthen the derivative time, or filter less, as Chapter 6 and Chapter 5 show.

4. *Does the gain move or change sign over the range?* Use feedforward of the holding input and a schedule, or feedback linearization (Chapter 3). At a peak where the gain vanishes, regulate below it with a margin (Chapter 4).

5. *Is the plant underactuated or nonminimum phase?* Try a cascade with the inner loop faster than $2p$ and the outer slower than $z/2$. If no loop pair satisfies both, as on the double pendulum, design on the full state with an observer.

6. *Does a speed or a friction move the poles?* Schedule on it and check the sampled loop at every point of the schedule, not only the continuous one (Chapter 9, Chapter 10).

7. *Is the problem a resonance the command excites?* Shape the command when the frequency is known to within the shaper's tolerance, which the crane's 3.3 % error already reached. Use preview when the disturbance can be seen coming (Chapter 11).

8. *Is the limit a delay inside the process, such as the regenerative cut?* Move the operating point out of the unstable region (Chapter 13).

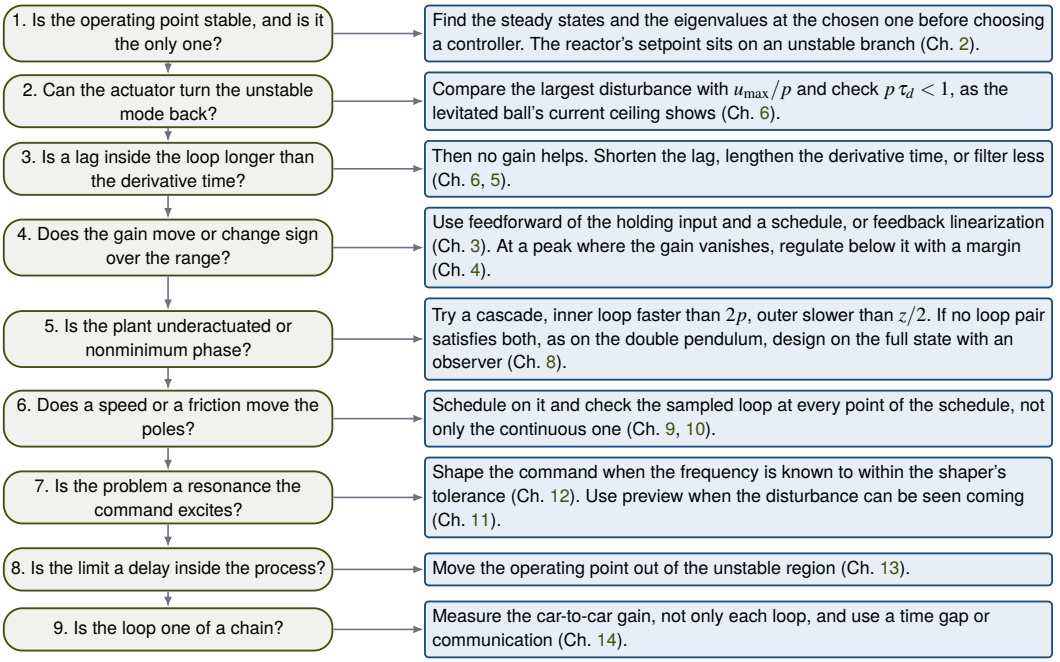


Figure 15.2 The nine questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the operating point and the actuator's reach are checked before any structure is chosen: the reactor's setpoint sits on an unstable branch, and the levitated ball spent most of its current ceiling on one small drop.

9. *Is the loop one of a chain?* Measure the car-to-car gain, not only each loop, and use a time gap or communication (Chapter 14).

A predictive controller earned its place where a limit bound during the transient. On the reactor it raised the largest safe setpoint step from 11.5 to 13.5 K (Section 2.8). On the cart it stopped a kick at 0.768 m where the regulator crossed the line to 0.809 m (Table 7.7), and on the crane it quieted the standard move 2.16 s sooner than three impulses (Table 12.5). Away from the limits it added little or lost. On the levitated ball the same law clipped without look-ahead held the headline step at 0.325 mm s against 0.319 mm s (Table 6.6). On the line of cars the retuned classical law reached 0.975 against 0.980 (Table 14.6). The double pendulum's MPC fell in five cases its regulator held (Table 8.6). Volume 9 treats the choice in full. Figure 15.2 draws the nine questions in order.

15.3 Where Practice Is Heading

Stabilization beyond the handling limit. A controller has drifted a modified car along a path of varying curvature, speed and sideslip, derived without reference to a nearby equilibrium. It used the yaw acceleration to stabilize the sideslip, and a test run without wheel-speed control showed that the dynamics of the saturated rear tire must be accounted for [1]. Chapter 10 held its car about one equilibrium at a time with a schedule. A review of suspension control finds semi-active and active suspensions implemented by industry and road preview becoming viable for production vehicles [2], the structure that cut the hump rise of Chapter 11 from 36.7 to 5.9 mm. In the editor's view the vehicle chapters of this volume describe functions that are moving from test tracks into series production, and the work moves with them to proving the margin over the whole envelope of road and tire.

String stability leaves the simulator. On a circuit with no bottleneck, a uniform flow of cars broke into a jam that traveled backward like a wave once the density passed a critical value [3]. Later experiments on a circular track with more than 20 vehicles damped such waves by controlling the velocity of a single vehicle, and the authors suggest that a few mobile actuators, less than 5 % of the traffic, could control the flow [4]. A review of platoon control shows that the definitions and analysis methods of string stability differ across the literature and sets out the equivalences among them [5]. The ring of 22 drivers in Chapter 14, whose first car stopped 98.7 s after a uniform start, is that experiment in a browser.

Chatter is still the limit of the machine. A review of chatter research since the stability laws of the 1950s calls chatter the main obstacle to raising the productivity and quality of machined parts, and it covers process damping and cutters with variable pitch designed to suppress it [6]. The simulator of Chapter 13 has neither, which is why its lobes are pessimistic at low speed. Command shaping for flexible systems, the method of Chapter 12, has a published history of more than fifty years [7].

Sensors that come closer to the state. Two chapters were limited by how the state reached the controller. The levitated ball is read by an infrared beam that it partly blocks, a sensor separate from the coil, and the camera on the ball and beam cost that loop 14.4° of phase margin. A review of self-sensing magnetic bearings, which estimate the gap from the coil's own inductance, finds them no longer chiefly a research concept but in commercial use in turbomolecular pumps and elevator guideways [8]. A survey of event cameras reports a latency of the order of microseconds and less motion blur than a frame camera, and finds that new methods are needed to process their output [9]. In the editor's view, a sensor that removes a lag lets a stabilizing loop carry more derivative action before noise or delay stops it.

Electric actuators at the brake and the handlebar. In aircraft braking, electromechanically actuated brakes have found wide acceptance as an alternative to hydraulic ones, and advanced antiskid controllers use them [10]. The United States light-vehicle brake standard requires that where an electric vehicle's regenerative braking is part of the service brake, the antilock system control it [11]. A steering motor added to a consumer electric bicycle stabilized the riderless bicycle down to just below 4 km/h, and with the assistance on, 26 riders on a treadmill at 6 km/h recovered from handlebar pushes significantly more often [12]. The slip loop of Chapter 4 and the scheduled steer of Chapter 9 are the laws such actuators carry.

Feedback from the load, not only from the drive. An inertial unit mounted on the load closed an anti-sway feedback loop on a five-ton industrial overhead crane. It damped the first swing mode while the setpoint stood at zero, and once the hoist moved it kept that mode from being excited, at the cost of slower tracking of the trolley's velocity setpoint [13]. The crane of Chapter 12 removes its swing by shaping the command alone, which works only while the cable length is known to within the shaper's tolerance. Figure 15.3 sets these developments and those above in three threads.

Optimization and learning wrap a stabilizing core. Fast solvers for embedded optimal control are available as open-source software [19]. A review of learning-based predictive control finds most work learning the prediction model from recorded data, and it describes predictive control used to add constraint satisfaction to controllers that were learned [20]. Control barrier functions verify and enforce safety properties in optimization-based controllers [21]. In the editor's view the plants of this volume will keep a linear stabilizing law at their core, a PD, a cascade or an LQR, with optimization above it to respect limits. The chapters point the same way, since the predictive controllers paid where a limit bound and not elsewhere.

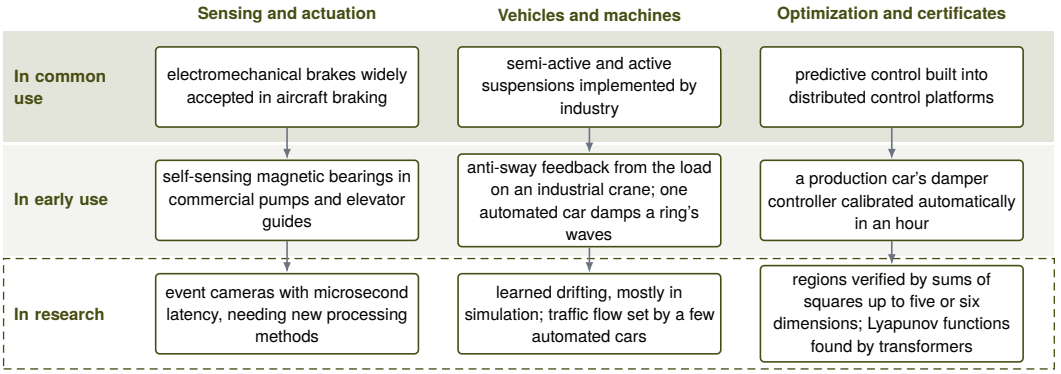


Figure 15.3 Developments in three threads. The top lane holds what the cited source calls widely accepted, implemented by industry or a major change in industrial practice, the middle lane what it reports in commercial use in a few applications or tested on production or industrial equipment, and the bottom lane what it reports from simulation or names as a need, in the editor's reading. Sources by thread: sensing and actuation [8–10], vehicles and machines [2, 4, 13, 14] and optimization and certificates [15–18]. No date is drawn because the sources give none for these lanes. Notice that each thread ends in research that must still prove a margin on hardware.

Certificates and designs proposed by machines. Sequence-to-sequence transformers trained on synthetic samples have found Lyapunov functions for polynomial systems more often than algorithmic solvers and people, and have found new ones for nonpolynomial systems [18]. Agents built on language models tune controllers iteratively against stated requirements on stability, performance and robustness [22]. In the editor's view a certificate or a stabilizing gain proposed by software is a claim that a computation can check, and this volume shows which computation. Check the eigenvalues of the sampled loop, the margin of the mode that fails and the response from the edge of the region it claims.

> ⧖ HISTORY **2010 and 2022: From Local Laws to Regions With Proofs**
>
> The history boxes of Chapter 1 end with a survey of 2000 that left open how fast a schedule of local designs may move. In 2010 a method built a tree of trajectories, each held by a linear quadratic regulator, and computed the region each one recovers by convex optimization. Its authors noted that the weakness of the older picture of funnels had been the difficulty of computing, or estimating by trial and error, the region of each local controller [17]. A review of 2022 found the region of attraction used in learning-based control as a means to guarantee safety, the set of states from which the controller drives the system back [23]. In the editor's reading the thirteen chapters stand at that point. Most of their stabilizing laws are linear about one point, and the regions they report were found by runs of the simulator, not by a proof.

15.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [24]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [25]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

15.4.1 How This Volume Was Made

Each of the thirteen apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run, and a referee report was written on each paper. At least five of the thirteen reports are self-reviews by the agent that wrote the paper. Defects went into a log for each app with a severity, a cause, a suggested fix and a status. For this book, one agent per chapter reran every number an exercise uses. The editor set the scope of each round and what each agent could change. The record holds the following errors and the checks that caught them.

The defect logs. The thirteen logs hold 144 entries with a severity. Of these, 32 are wrong behavior of an app (a wrong output, a wrong lesson or a controller reading what it should not), 85 are documentation, 22 are cosmetic or display, and 5 concern completion rules or the interface. The cases below were chosen for this volume's topic. Each is a claim about stability, an equilibrium or a margin.

A stable state that was not. The reactor's help text said that at a jacket temperature of 300 K the reactor could ignite to a second, stable state at 369.7 K. While writing the paper, an agent computed the Jacobian there and found eigenvalues 1.357 ± 1.540j, an unstable focus. Released from 369.7 K, or from anywhere up to 380 K, the run returned to 324.475 K every time. The same log records a footnote that named the nearest root of the energy balance as the temperature the reactor would settle at, without examining its stability, and a tip that advised lowering the gain against a runaway. Measured, the lower gain raised the peak from 451.44 to 469.88 K.

A stability window computed badly. The first gain window of the cart and pole was computed by power iteration. It produced a floor that wandered from 18.3 to 18.9 N/rad and three ceilings the plant cannot have. A first grid of nonlinear runs, read too early, labeled slowly settling loops as rocking. In review the boundaries were recomputed from the spectral radius of the sampled loop, checked against exact eigenvalues, and the window was set against 40 nonlinear runs, which agreed at 39. The same review found that numbers lifted from the app's guide were wrong, among them a gain of 15 N/rad called "below the floor" of 11.3 N/rad.

A design claim that the plant did not support. Comments in the suspension app said that a stiff proportional leveling gain made the body unstable. In the self-review the agent tested it. At the default comfort setting no gain up to the slider's 50 N/mm destabilized the body. At 0 % comfort it did, between 11 and 12 N/mm, through a mechanism the review then derived. The first draft had also given both passive modes a damping ratio near 0.32 from uncoupled formulas, where the eigenvalues give 0.25 and 0.33.

An equilibrium on the wrong branch. The drift controller designs its LQR about an equilibrium that a solver finds for the requested speed and sideslip. While writing the paper, the agent swept every target a learner can set. On the dry pad 89 of 1,485 returned a yaw rate of the wrong sign, a car turning against its own drift, and the regulator would have been designed about that point. None of the 9,090 targets the shipped cases reach was affected. The fix accepts only a root with the right sign and adds a test over 5,940 targets.

Controllers that read the plant. Three logs record a controller that took values a real controller could not have. The levitation PID computed its holding current from the run's own ball. A 10 % heavier ball scored an IAE of 0.108 mm s when it came through the disturbance slider and 0.008 mm s when it came through the parameter, which the controller read. The braking controller's low-speed exit read the true road. On identical samples, a change from dry asphalt to ice moved its held torque from 1336.9 to 200.0 N m. The fix moved every run that reaches the exit, 10 of the 13 scenarios, and removed wheel locks, among them 0.15 and 0.23 s under the PI loop. The reactor's controllers started from the run's balances, and 10 of 11 edits made before Run moved the controller's start while no run showed it. On the double pendulum, a test meant to catch such reads in the predictive controller missed seven of eight versions doctored to read the truth, and it was rewritten until it caught every doctored version tried. These were found by audits that asked what each controller reads, and by tests built to fail.

A headline chosen after the run. The levitation paper's comparison step was chosen from twelve candidates after both controllers had run on them. The referee asked whether the headline was fair. The paper now says how the step was chosen and carries the counterfactuals. The predictive law clipped without look-ahead held the same step, so state feedback, not the optimization, won it, and the chapter states the result that way.

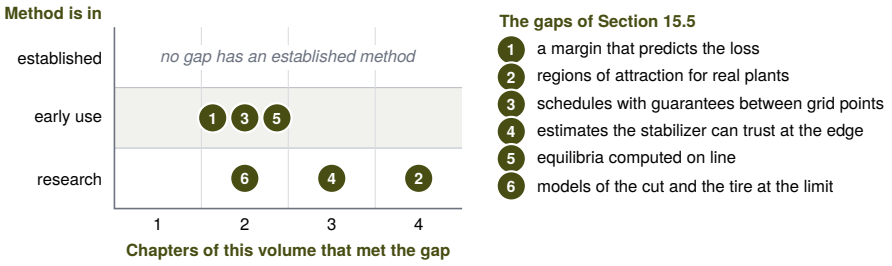


Figure 15.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 3 and 14), four for gap 2 (7, 8, 2 and 5), two for gap 3 (9 and 10), three for gap 4 (4, 10 and 8), two for gap 5 (10 and 2) and two for gap 6 (13 and 10). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap met by the most chapters, the region a controller can recover, is still in research.

The book round. The thirteen chapter records list 215 rows of numeric verification, each a stated value beside its reproduction. No reported value drifted beyond a rounding (a minimum gain printed 1.70, since corrected to 1.69) and three counts of the apps' own checks, which have grown since the papers. The bicycle agent found a briefing step that a student could follow faithfully and learn the opposite lesson from. The "hands-off" run kept the loop's last torque and fell in 0.55 s. It then appeared to fall sooner than the wrong-signed loop's 0.61 s, where a true hands-off run takes 1.66 s. The chapter's exercise tells the reader which button to press, and the briefing was later corrected. The book round's defect list held 10 such app or paper items for this volume. A later repair pass fixed eight, confirmed one as printed and kept one, a count of checks, as a historical value.

The pattern is the one Section 15.4 describes. Every error above was caught by a check that could fail, such as eigenvalues computed at a claimed equilibrium, a sweep over every target, a nonlinear run set against a linear window, or a hidden value moved to see whether the controller moved. For a stabilizing loop the check matters more, because a held plant hides what its controller knows. Four working rules follow. Compute the eigenvalues of every state the text calls stable. Sweep the whole range a user can reach. Read what a controller reads. Let a person decide what changes.

15.5 Gaps in Nonlinear Stabilization

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 15.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 15.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *A margin that predicts the loss.* The aeropendulum lost a slow mode behind a healthy phase margin, and the line of cars collided at a linear peak of 1.031. A textbook notes that the gain and phase margins must both be given to show that a loop stays clear of the critical point, and it defines the stability margin as the shortest distance to it [26]. A review of platoon control finds that the definitions and analysis methods of string stability differ across the literature [5]. The gap is hard because each of these numbers describes one linear loop at one point, while the chapters lost a slow mode, a sampled loop or a chain of saturating drivers. A single robustness number that ranks designs by the disturbance, delay or parameter change that actually loses a nonlinear, sampled, saturating plant is not in routine use.

2. *Regions of attraction for real plants.* The cart and pole's largest recoverable push, the double pendulum's catchable band and the reactor's largest safe step were all found by search on a

simulator, and on the ball and beam no tuning keeps a flicked ball on a beam limited to 3°. The Open challenge box after this list takes the gap further.

3. *Scheduling with guarantees between the grid points.* The bicycle's schedule left a band of growth between design speeds, and the drift regulator scheduled for dry tires spun on a wet patch. A survey of gain scheduling found the relaxation of the slow variation restriction still very much an open question, judged the robustness analysis of gain-scheduled controllers very immature, and named soundly based interpolation between operating points and help in choosing the scheduling variable as pressing needs [27]. The gap is hard because each design point is linear and stable while the loop as it moves is neither. Guarantees for the scheduled loop as the parameter moves, checked on the sampled loop between the grid points, are thin for plants like these.

4. *Estimating what the stabilizer needs.* The braking law needs the friction peak, the drift schedule a friction estimate, and the double pendulum rates that the encoder does not give. A study of antilock braking reports that uncertainty in the road and vehicle states is often neglected or handled apart from the controller, which leads to sub-optimal braking, and that the friction peak is harder to find on snow and ice [28]. For aircraft braking, friction estimation is called an extremely challenging and still open task [10]. The gap is hard because the quantity is visible only near the edge the loop must avoid. An estimator that is fast enough to stabilize with and honest about its error at the edge of the operating region is an open design problem.

5. *Equilibria computed on line.* The reactor's ignited state at the published jacket temperature is an unstable focus, not a second place to operate, and at low speed and small sideslip the drift's equations admit a mirror-image root that turns the wrong way, which its solver rejects by the sign of the yaw rate. A drifting controller on a passenger car took its equilibria from maps computed offline, with an initial guess that steers the solver to the drifting equilibrium it wants [29]. The gap is hard because a solver that converges says nothing about which branch it found or whether it is stable. Controllers that compute their own operating point need a routine check of branch, sign and stability before they act on it.

6. *Models of the machine at the cut and the tire at the limit.* The milling model has no process damping, and the drift tire's friction does not fall as it slides. A review of chatter finds that stability predictions still suffer from uncertain measurements, time-varying dynamics and a worn tool, and asks for physics-based prediction made off line to be joined with on-line adaptive learning and tuning [6]. A drift controller tested on a sports car names tire model fidelity and friction estimation as what still limits it [14]. The gap is hard because the limits that matter most in this volume sit where models are least certain. Identifying them safely from the machine is still research.

OPEN CHALLENGE Certifying the Region a Controller Can Recover

Gap 2 runs through four chapters. Three found a recoverable push, band or step by search, and in the fourth no tuning catches a flick. Verifying regions by sums of squares has succeeded on systems of up to five dimensions, and checking by sampling costs less but loses the guarantee [17]. With a bounded input, a saturated feedback's basin of attraction lies inside the set of states from which some admissible input can reach the origin [30]. Guarantees for learned controllers rest on assumptions hard to verify before a robot runs [23]. In the editor's view the region is hard to certify because it belongs to the whole saturating loop. Progress would be a certified region that a run from its edge confirms.

15.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [23]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [25]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [31]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [32]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [33]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- Each machine could be pushed only so far, and past that edge it was lost at once: the distance to the edge is the number to report.
- The usual margin did not always see the edge. The margin worth quoting is the one measured on the mode and the path that fail.
- On underactuated plants the structure decided before the tuning, and the source of the derivative information limited what the loop could hold.
- Every error in the record was caught by a check that could fail, such as eigenvalues at a claimed equilibrium or a sweep over every target a user can reach.

References

[1] J. Y. Goh, T. Goel, and J. C. Gerdes, "Toward automated vehicle control beyond the stability limits: Drifting along a general path," *Journal of Dynamic Systems, Measurement, and Control*, vol. 142, no. 2, p. 021004, 2020.

[2] J. Theunissen, A. Tota, P. Gruber, M. Dhaens, and A. Sorniotti, "Preview-based techniques for vehicle suspension control: A state-of-the-art review," *Annual Reviews in Control*, vol. 51, pp. 206–235, 2021.

[3] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, "Traffic jams without bottlenecks: Experimental evidence for the physical mechanism of the formation of a jam," *New Journal of Physics*, vol. 10, no. 3, p. 033001, 2008.

[4] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D. B. Work, "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transportation Research Part C: Emerging Technologies*, vol. 89, pp. 205–221, 2018.

[5] S. Feng, Y. Zhang, S. E. Li, Z. Cao, H. X. Liu, and L. Li, "String stability for vehicular platoon control: Definitions and analysis methods," *Annual Reviews in Control*, vol. 47, pp. 81–97, 2019.

[6] Y. Altintas, G. Stepan, E. Budak, T. Schmitz, and Z. M. Kilic, "Chatter stability of machining operations," *Journal of Manufacturing Science and Engineering*, vol. 142, no. 11, p. 110801, 2020.

[7] W. Singhose, "Command shaping for flexible systems: A review of the first 50 years," *International Journal of Precision Engineering and Manufacturing*, vol. 10, no. 4, pp. 153–168, 2009.

[8] 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

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

[10] F. Crocetti, G. Costante, M. L. Fravolini, and P. Valigi, "Tire-road friction estimation and uncertainty assessment to improve electric aircraft braking system," in *29th Mediterranean Conference on Control and Automation (MED)*, 2021, pp. 330–335.

[11] National Highway Traffic Safety Administration, "Standard no. 135; light vehicle brake systems, 49 CFR 571.135," Code of Federal Regulations, Title 49, Chapter V, edition of October 1, 2024, U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title49-vol6/pdf/CFR-2024-title49-vol6-sec571-135.pdf

[12] M. T. Haitjema, L. Alizadehsaravi, and J. K. Moore, "Automatic bicycle balance assistance reduces probability of falling at low speeds when subjected to handlebar perturbations," *IATSS Research*, vol. 50, no. 1, pp. 701–711, 2026.

[13] V. Helma and M. Goubej, "Active anti-sway crane control using partial state feedback from inertial sensor," in *2021 23rd International Conference on Process Control (PC)*. IEEE, 2021, pp. 137–142.

[14] F. Werner, T. Heintzenberg, M. Lienkamp, and J. Betz, "Drifting in the future: Stabilizing path following drifting on high-latency vehicle systems," in *2026 IEEE International Conference on Robotics and Automation (ICRA)*, 2026, pp. 12 763–12 770.

[15] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[16] G. Savaia, Y. Sohn, S. Formentin, G. Panzani, M. Corno, and S. M. Savaresi, "Experimental automatic calibration of a semi-active suspension controller via Bayesian optimization," *Control Engineering Practice*, vol. 112, p. 104826, 2021.

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

[18] A. Alfarano, F. Charton, and A. Hayat, "Global Lyapunov functions: A long-standing open problem in mathematics, with symbolic transformers," arXiv preprint arXiv:2410.08304, 2024.

[19] R. Verschueren, G. Frison, D. Kouzoupis, J. Frey, N. van Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados: A modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, no. 1, pp. 147–183, 2022.

[20] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020.

[21] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[22] X. Guo, D. Keivan, U. Syed, L. Qin, H. Zhang, G. Dullerud, P. Seiler, and B. Hu, "ControlAgent: Automating control system design via novel integration of LLM agents and domain expertise," arXiv preprint arXiv:2410.19811, 2024.

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

[24] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[25] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

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

[27] D. J. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control*, vol. 73, no. 11, pp. 1001–1025, 2000.

[28] B. Sullivan, J. Jiang, G. Mavros, and W.-H. Chen, "An exploration-exploitation approach to anti-lock brake systems," arXiv preprint 2306.14730 [eess.SY], 2023. [Online]. Available: https://arxiv.org/abs/2306.14730

[29] S. Meijer, A. Bertipaglia, and B. Shyrokau, "A nonlinear model predictive control for automated drifting with a standard passenger vehicle," in *2024 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, 2024, pp. 284–289.

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

[31] International Electrotechnical Commission, "Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels," IEC 62443-3-3:2013, Edition 1.0, 2013.

[32] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[33] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

15.6 Exercises

15.1 Discussion · 8 min

An agent hands you a PID tuning for a new propeller arm with the report "phase margin 50 to 55° at every angle from 15 to 105°, so the loop is robust over the range." Using Chapter 3 and Table 1.1, say why that report is not enough, and list what you would ask the agent to compute or run before the tuning is used.

15.2 Discussion · 7 min

On the double pendulum of Chapter 8 a loop on each link beside a cart loop cannot hold the plant at any tuning. A colleague proposes to let an agent search the six PID gains overnight on the simulator. Explain why the search cannot succeed, what result it would return, and how you would redirect the work.

15.3 Discussion · 8 min

Section 15.1 says that some limits belong to the plant. Take two of the following and, for each, name the limit, say why no feedback controller of the chapter removes it, and say what does: the pothole edge of Chapter 11, the flicked ball of Chapter 5, the stability lobe of Chapter 13, the crane swing of Chapter 12 with a saturated drive. Then name one plant from your own field with a limit of the same kind.

15.4 Discussion · 8 min

An agent proposes replacing the LQR of a balancing machine with a predictive controller, citing a case where the predictive controller won. Using the results gathered in Section 15.2, design a comparison that would convince you, and say what result would make you keep the LQR.

15.5 Discussion · 6 min

Section 15.4.1 reports that the levitation PID computed its holding current from the run's own ball, and that no shipped run showed it. Explain why a controller that reads the plant's true parameters is harder to detect on an unstable plant than on a stable one, and describe a probe that would detect it in an agent-written controller.

15.6 Discussion · 7 min

The drift solver of Chapter 10 once returned equilibria on the wrong branch for 89 of 1,485 targets, and the reactor's help text once called an unstable focus a stable operating point. Write a short checklist that a person or an agent should apply to any computed operating point before a stabilizing controller is designed about it, and say which item would have caught each of the two errors.