INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 4 Override, Split-Range, and Constraint Control

CHAPTER 15

Summary and Outlook

Pages 373–386

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Override, Split-Range, and Constraint Control*, John D. Hedengren, Ed., vol. 4 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 15, pp. 373–386.

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/v04-override-constraint/

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

CHAPTER 15

Summary and Outlook



Open gaps 1 actuator sizing 2 release of overrides 3 margins from prices 4 usable stock 5 unmeasured limits 6 selectors and optimizers

Notice that the gaps of release and of moving limits meet on the oven and the shield, while the gap of usable stock joins the office, the fan car and the vault. Illustration, not to scale.

IN PLAIN TERMS **What Thirteen Plants Taught About Limits**

The main idea. Every chapter of this book ran something close to a limit it must not cross, from a jet engine to a company. Some designs let a safety rule take over the controls as a limit came near. Others shared the work among several handles, the cheap ones first.

The challenge. A safety rule works only if its handle moves fast and far enough, and only if it lets go when the danger has passed. One that stays on costs output. One that comes late costs the machine.

Why it matters. Testing each safety rule away from the normal case, and pricing what it gives up, separated the designs that worked from those that failed.

Chapter 1 asked two questions of a plant at its limits. When several variables must stay inside a limit and one actuator moves them all, which variable gets the actuator? When several actuators move one variable, which one should move, and for how long? Thirteen chapters answered them on machines, process lines, vehicles, a company and a central bank. This chapter sets the answers side by side, draws a short guide for choosing a structure, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

15.1 What the Chapters Showed

Table 1.4 listed what the selectors of each chapter protect, how its actuators are shared and which stock depletes. Tables 15.1 and 15.2 return to that table with the measured result, the first for one main actuator and several limits, the second for more actuators than controlled variables. Each row names the limit or the sharing problem, a structure that met it, one that did not, and the result with the table or section that reports it.

The columns of the two tables make five points.

A protective loop is bounded by its actuator before its tuning. On the compressor of Chapter 2 the controller acted correctly with a slow valve and the machine surged anyway, to −105.4 % of margin. With a valve too small it surged to −129.0 %, and no gain or integral time changed either outcome. The trigeneration plant of Chapter 9 emptied its uninterruptible supply in every one of 54 settings once demand exceeded the turbine and the grid, the fan car of Chapter 12 found no controller among 864 that landed a 12 s hover softly, and the CO_2 cycle of Chapter 10 made 1188 kW of a requested 1270 kW with every layer at its stop (Section 10.6.4). Each chapter said so, as Section 1.4.5

Table 15.1 What Each Selector Protected, Chapters 2 to 8

Chapter	Limit	What worked	What failed	Measured result
2, compressor	Surge line, guarded by one recycle valve against a fast throttle closure.	A PI override resting on its closed clamp, a fast valve, and the throttle signal as a minimum opening through a high select.	Any tuning on a valve too slow to stroke or too small to pass the flow.	Minimum margin −8.6 standard, +3.3 fast, −105.4 slow, −129.0 undersized, in percent (Table 2.5). High select +16.16 against −8.61 (Table 2.6).
3, turbofan	Six caps and two floors on one fuel valve. A slam outruns the spools.	An 8 % acceleration schedule with integrator tracking on every limiter.	No schedule. In the challenge, the schedule or tracking alone.	Challenge passes at 5.86 s, 9.4 % margin, 938 °C. Shipped 8.49 s and 996 °C, schedule alone 1019 °C (Table 3.5).
4, combine	Engine, rotor torque and loss limit one ground speed. Moisture is unseen.	Five demands in a low selector with external reset, a 90 % engine limit, preview, a loss supervisor and trims.	The overrides off at the shipped feed. A 95 % engine limit.	Against the rain 10.77 t, worst pass 1.71 %, 1.45 s overload, against 9.03 t, two plugs and 74.70 s. At 95 %, 31.50 s (Table 4.5).
5, extruder	Torque, die pressure and dryer air on a melt heated by its own shear. A wrong flour certificate.	Overrides as clamps, a density trim restarted at each lot, a throughput supervisor at 420 kg/h.	The supervisor alone. The trim carried from one lot into the next.	Accepted 290.2 against 186.3 kg shipped and 172.5 kg for the supervisor alone. No lot restart: torque 99.2 % (Table 5.8).
6, coater	A 25 % LEL cap that binds in whichever zone holds the drying front.	An override with three handles in order, exhaust then speed or zone heat, under a push to the first binding limit.	The push without the override. The push without the solids feedforward.	3076 m^2/h at 21.9 % LEL passes. Without override 3400 m^2/h at 28.1 % LEL, without feedforward 2075 m^2/h (Table 6.6).
7, drill	Torque, mud ratio or thrust holds the rate, by soil. The pressure limit falls with the cover.	A rate optimizer on the binding limit, a pressure override on the thrust as backstop, and guidance.	Fixed clay settings in sand. More mud flow alone on the exit climb.	Exit climb: breakout at 149.8, and at 154.8 m with more flow. Override 44.2, optimizer 76.6 m/h (Table 7.5). Crossing 0.60 m off in 2.13 h, no return (Table 7.6).
8, startup	A cash floor, and a credit line whose ceiling moves with trailing EBITDA.	A runway override split over three ranges ahead of the growth loops.	The board by hand. Higgins's rule. A cash-flow MPC.	Lowest cash $0.79 against −$2.03 million, insolvent at 1.88 yr (Table 8.5). Challenge 20.7 % a year, $272.5 million. Higgins fills 79.5 % (Table 8.7).

recommends: a saturated protective actuator is a sizing statement. Figure 15.1 sets one like pair of numbers from each row of the two tables side by side.

Anti-windup decides when the override takes over. The turbofan of Chapter 3 passed its challenge only with integrator tracking on every limiter. With the tightened schedule alone the exhaust limiter's integral had wound up while the engine was cold, and the indicated temperature reached 1019 °C against a 950 °C limit. The tunnel shield of Chapter 11 starts each of its five advance limits from the request actually selected, so none winds up while another holds the advance. The CO_2 cycle hands the bypass controller's unsaturated demand upward, because the clamped position reports the same error for a small and a large shortfall (Section 10.5.3). External reset is part of the protection, not a refinement of it (Section 1.4.4).

Table 15.2 How Spare or Limited Actuators Were Shared, Chapters 9 to 14

Chapter	Sharing problem	What worked	What failed	Measured result
9, trigeneration	Cheap absorption cooling answers in minutes, electric chillers in seconds, a store in between.	A fast electric trim and a VPC that hands the duty back to absorption, with a stronger store discharge.	Absorption alone. The store gain alone. Any tuning once capacity is lost.	IAE 951.41 against 5.16 K s, reserve regained at 806 s (Table 9.7). Challenge 0 against 600 s out of band (Table 9.8).
10, CO_2 cycle	Bypass fast and wasteful, firing slow and capped, compressor pressure the only headroom.	Three mid-ranging layers, each handing the unsaturated demand upward, with a 3 % resting bypass.	Firing alone, or bypass and firing. A resting bypass of 0 or 8 %.	Step up: in band at 6.8 min, IAE 453 against 16751 (Table 10.7). Day: 0 gives 116.3 min out of band, 8 discards 509 kW h (Table 10.10).
11, tunnel shield	The screw holds the chamber, ring builds stop it, and five limits bound the advance.	The least of five limits with external reset, a screw-room limit, a bentonite hold and a chainage feedforward.	A fixed advance. No hold. A fixed pressure setpoint under the river.	Drive 46.9 h and 0.0 mm against 57.6 h and 245.3 mm. Fixed advance: torque 90.04 %. No hold: 21.7 mm (Table 11.5).
12, fan car	Fans, tires and cold-gas thrusters at three prices. The tank runs out by pressure.	Fixed priority: fans, tires at the slip limit, thrust metered to 3.4 g. A pressure abort.	Thrust wide open. A 12 s hover without the abort.	Launch 0.83 s on 29.1 kg against 0.60 s on 40.5 kg (Table 12.5). A 12 s hover lands at 2.83 m/s with 22.6 % left, the abort at 0.94 m/s (Table 12.6).
13, quadruped	Feet that push but cannot pull, inside friction cones that change with each step.	Weighted least squares inside the true cone, the actual leg limits and the sensed contacts.	A cone planned at 0.6 on 0.15. Nominal limits for a weak leg. Scheduled feet.	Slip 15.6 against 1.07 cm. Weak leg 8.43° against 3.22°. Scheduled feet fall. Course finished at 9.25 s against a fall (Table 13.6).
14, currency peg	A vault with a floor, a drain that grows as it empties, two instruments for three limits.	A capital-controls override checked every day, outside the weekly calendar.	Every setting without controls. Any controller against a monetized deficit.	Override: 4.97 months of imports, 6.72 %, controls 0.384 yr. Shipped 2.45 months (Table 14.6). Deficit: attack at 2.252 yr under every controller (Table 14.5).

Every override has a measured price in throughput. The coater's override held every zone under the cap through a fan loss at about a fifth less output, a mean of 27.6 m/min (Table 6.4). On the drill's exit climb the override alone made 44.2 m/h where the optimizer, which predicts the pressure, made 76.6 m/h and stayed clear. The startup's 24-month runway target kept $1.50 million at the lowest and left $42.3 million of founders' value at year 5, against $74.0 million at 12 months (Table 8.5). Section 1.4.2 called this the back-off, and each chapter measured it.

A resting position or a priority is a purchase. The CO_2 cycle's resting bypass bought tracking with fuel, at about 3.8 kW h for each minute out of band recovered, and only 3 %, of the three positions tried, passed both clauses of the dispatch day. The fan car's metered launch spent 28 % less gas than the wide-open one for 0.23 s. The quadruped's assumed friction coefficient traded speed of recovery against slip, and the trigeneration plant held 0.5 MW of electric cooling as the reserve to which its valve-position controller returns the fast actuator. None of these numbers comes from a tuning rule. Each is a price.

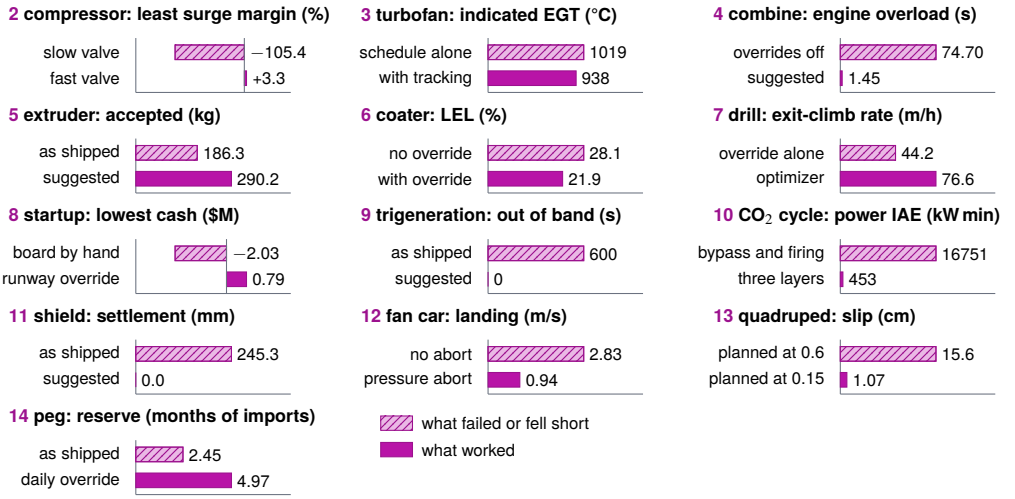


Figure 15.1 One like pair of numbers from each row, as the "Measured result" column of Tables 15.1 and 15.2 prints them. Each panel has its own scale, and the compressor and startup panels run below zero, with a line at zero and the negative value printed to its right. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. On the drill both rates kept clear of the pressure limit, and the override alone was the slower. Notice that the closest pairs, on the turbofan and the coater, still straddle their limits, 950 °C of indicated exhaust temperature and 25 % LEL: a small difference in the number was the difference between passing and failing.

A stock that depletes turns time into a limit. The peg of Chapter 14 held only with a protective layer that ran every day. Checked at the weekly decisions instead, the same override lifted late and kept the controls in force for 0.521 year, past the six-month limit. Against a monetized deficit no controller moved the date of the attack, which credit growth set at 2.252 yr. The fan car's tank ran out by pressure with 22.6 % of its nitrogen still inside, so the abort that read the tank pressure landed the car softly while the reserve on the mass gauge never fired. Where a protection worked, it read the quantity that actually runs out, and acted faster than the policy it protected (Section 1.4.9).

15.2 Choosing a Structure

The chapters suggest an order of questions for a plant at its limits. Each names a chapter where its answer decided the result.

1. *Which limits are active at the best operating point, and what does each unit of back-off cost?* Price the margin first, in grain, kilograms or meters an hour (Chapter 4, Chapter 5).
2. *Can the protective actuator act fast enough and far enough?* Check the stroke time and capacity against the fastest credible upset before any tuning (Chapter 2).
3. *One actuator and several limits?* Put the demands in a low or high selector, floors after caps, with external reset from the selected output (Chapter 3, Chapter 4). Then test that each override lets go when its limit clears (Chapter 6).
4. *Several actuators for one variable?* Split range shares one demand among actuators of similar cost. A fast, costly actuator beside a slow, cheap one calls for mid-ranging on the unclamped demand, with the resting position set as a price (Chapter 9, Chapter 10). Many actuators with limits call for allocation, by priority where one actuator is free and another is spent (Chapter 12) and by weighted least squares where the limits change each period (Chapter 13).

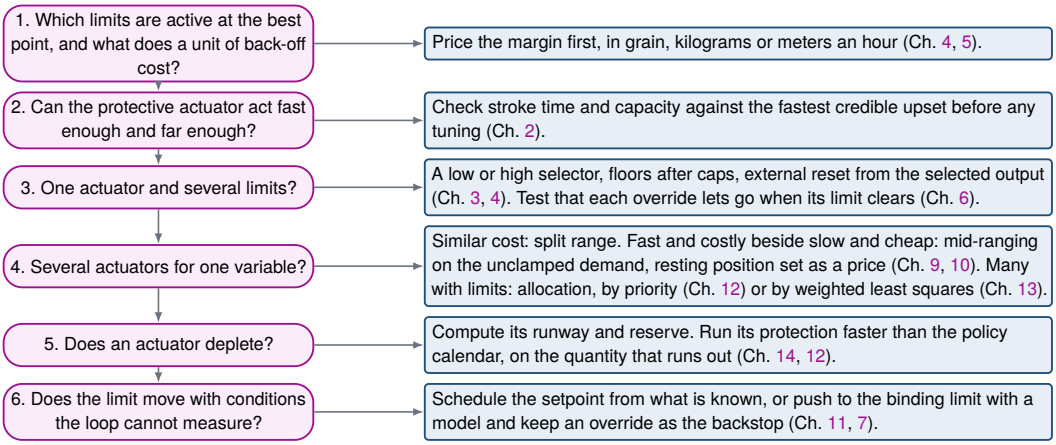


Figure 15.2 The six questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the price of the margin and the reach of the actuator come before any selector is drawn: a structure cannot protect a limit that its actuator cannot reach.

5. *Does an actuator deplete?* Compute its runway and reserve, and run its protection faster than the policy calendar, on the quantity that actually runs out (Chapter 14, Chapter 12).
6. *Does the limit move with conditions the loop cannot measure?* Schedule the setpoint from what is known, as the shield's chainage feedforward does from the geology, or push to the binding limit with a model and keep an override as the backstop (Chapter 11, Chapter 7).

Figure 15.2 draws the six questions in order.

A predictive controller handles limits and prices in one optimization, and it paid where it coordinated handles that the selectors moved in sequence. Over the CO_2 cycle's dispatch day it left 781 kW min of power error against 8425 for the shipped hierarchy and 1050 for a retuned one (Table 10.11). Through a wet flour lot it accepted 212.9 kg with nothing off specification, against 169.4 kg for the loops (Table 5.11). It did not always win. On the compressor the PI override given the throttle signal through a high select held +16.16 % of margin against the predictive controller's +15.00 % (Table 2.6). On the combine the MPC failed the challenge with 48.60 s of engine overload where the selector network passed with 1.45 s (Section 4.6.8). Much of the extruder's gain came from its estimate: the loops given the MPC's estimate accepted 186.0 kg against the MPC's own bias form at 177.6 kg (Table 5.9). Volume 9 treats that choice in full.

15.3 Where Practice Is Heading

Simple elements are being given a theory. A 2023 review collects the standard elements of advanced regulatory control, among them selectors, split range, cascade and feedforward, and challenges the view that predictive control can solve every control problem while these elements are ad hoc. It calls for better teaching of the elements and for research on their theory and design methods [1]. Split range itself has been generalized to several independent controllers with one setpoint, of which only one is active at a time, and the same work notes that predictive control needs a detailed dynamic model and does not use one input at a time [2]. Feedback optimizing control, which reaches the economic optimum by choosing what to control rather than by solving an optimization online, has been surveyed as a family of real-time optimization methods [3]. The design and performance of mid-ranging controllers were studied in the 1990s [4], and the resting position of Chapter 10 shows that the question is still economic as much as dynamic. In the editor's

view, the selector networks of this volume will increasingly be designed from the active constraints of the plant's optimum rather than drawn from experience, as Chapter 1 began to do.

Allocation is solved online. Control allocation, the sharing of a demanded force or moment among redundant actuators, has been surveyed as a field of its own [5]. On a quadruped robot, ground reaction forces have been planned by convex predictive control over horizons up to 0.5 s and solved in under 1 ms, with one set of gains and weights across many gaits [6]. The fixed priority of Chapter 12 and the weighted least squares of Chapter 13 are the two ends of that range. Industrial predictive control technology, the other route to the same problem, was surveyed in 2003 [7].

Constraint enforcement becomes a separate layer. Reference and command governors modify the setpoint given to a loop already designed, so that its constraints hold, and their theory and applications have been surveyed [8]. Control barrier functions express safety as the forward invariance of a set and join it with performance objectives in a quadratic program, demonstrated on adaptive cruise control and lane keeping [9]. A 2023 survey presents safety filters that keep state constraints satisfied under the physical limits of the system whatever the controller behind them asks, and observes that missing safety certificates still keep many new designs inside research environments [10]. In the editor's view, this is where the override of this volume is going. A controller written or tuned by an agent, or learned from data, will sit behind a layer that protects the limits, and that layer will have to be simpler, faster and better verified than the controller it protects, as the daily override of Chapter 14 was faster than the weekly policy.

> ⧖ HISTORY **1942 to 1997: From a Fuel Stop to an Estimated Limit**
>
> The timeline of Chapter 1 places a jet-engine governor in 1942. The first jet engine built in the United States that year had a hydromechanical governor, a minimum-flow stop on the fuel valve against flame-out and a maximum-flow schedule against over-temperature. Flight research later brought estimates into the control. A distortion-tolerant control, flight-tested in 1997, estimated the inlet distortion from engine-face pressures and flight-control estimates and trimmed the engine's actuator commands for stability. A Kalman filter that estimated engine parameters in flight informed a self-tuning onboard engine model implemented in a military engine's control unit [11]. In the editor's reading, the fixed caps of Chapter 3 stand at the start of that arc, and the estimated limits of Section 15.3 at its end.

Limits are estimated where no sensor reaches them. In an aircraft engine, thrust cannot be measured, and limited quantities such as turbine inlet temperature and minimum stall margin cannot be measured directly. An onboard engine model with a tracking filter lets a loop close on such unmeasured values, but the health parameters to be estimated generally outnumber the sensors, which can leave significant error in the unmeasured outputs. One such approach was not yet fully implemented in operating engine controls, because of the processing it needs and of certification issues with a model-based sensor estimate [12]. The limiters of Chapter 3 read an indicated exhaust temperature. In the editor's view, an override that acts on an estimated limit will need its back-off sized from the error of the estimate, which is gap 5 of Section 15.5.

Predictive control reaches the protective layer. For industrial centrifugal compressors, nonlinear predictive controllers and a sequential quadratic programming approximation have been formulated for process and anti-surge control together, tested in simulations that emulate a realistic test bench and timed on an industrial programmable logic controller. In those scenarios they outperformed standard industrial control on surge margin, tracking and actuation without significantly more computation [13]. A study of predictive control for a supercritical carbon dioxide power converter argues that one predictive regulator might outperform several PID loops that each need their own tuning and decoupling, because it coordinates the inputs [14]. On the compressor of Chapter 2, the PI

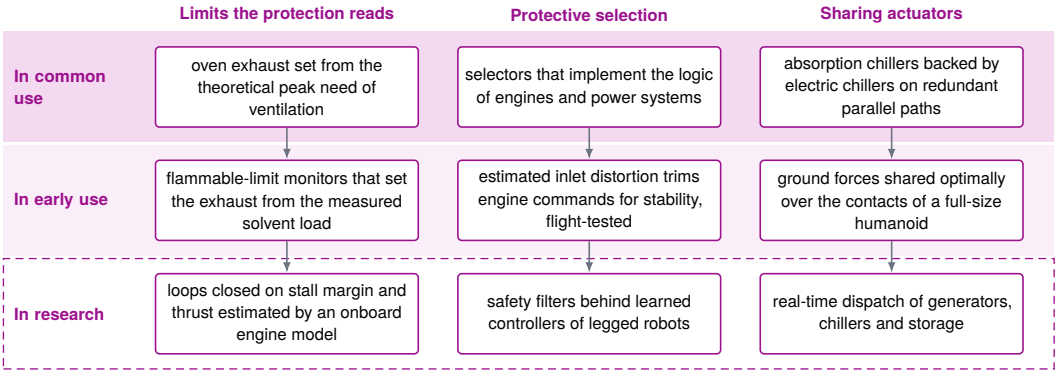


Figure 15.3 Developments in three threads. The top lane holds what the cited source calls standard, favored or common practice, the middle lane what it reports offered as equipment, flight-tested or shown on hardware, and the bottom lane what it reports from research, as a proposal or as a need, in the editor's reading. Sources by thread: limits the protection reads [12, 15], protective selection [11, 16, 17] and sharing actuators [18–20]. Notice that each thread moves from a fixed rule toward a computation that must itself be checked while the plant runs.

override given the throttle signal kept more margin than the predictive controller (Section 15.2). In the editor's view the predictive layer will take over the coordination, while a simple override keeps the last word on the limit.

Ventilation follows the measured solvent load. Ovens that drive off flammable solvents have set their exhaust from the theoretical peak need of ventilation, so that the highest vapor concentration the process can produce stays well below the lower flammable limit. Monitoring equipment that tracks the solvent load in real time and sets the ventilation from the real need keeps a safe ratio with less exhaust air, and linking it to an adjustable-speed drive on the exhaust fan improves the efficiency further [15]. The coater of Chapter 6 reads six zone monitors and trims the exhaust first when a zone nears its cap. In the editor's view a limit measured continuously is what lets the back-off of Section 15.1 shrink. Figure 15.3 sets these developments in three threads.

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 [21]. 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 [22]. 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 with verification notes that trace each number to a run. Each paper received a referee report, and at least four of the thirteen are self-reviews by the agent that wrote the paper. Defects went into a log for each app. For this book one agent per chapter reran the numbers its exercises use. The editor set the scope of each round and what each agent could change, and statements that runs contradicted were left for the editor.

The defect logs. Twelve of the thirteen logs give each entry a severity, 108 entries in all. Of these, 35 are wrong behavior of an app, 56 are documentation, 15 are cosmetic or display, and 2 are of other kinds. Nine of the 35, in seven logs, are a controller or an operator aid that read the plant's true state or the run's hidden settings where a plant would offer only a measurement. The drill's log records no defect and says that every number its paper quotes was reproduced on the shipped app.

A limit with no override. The extruder's dryer air had a hard limit of 150 °C, while the moisture cascade could ask for up to 160 °C. With humid intake air and a low moisture setpoint the air ran at 160 to 161.1 °C for half an hour, and a heater driven by hand reached 175.0 °C. No shipped or suggested run had passed 147.4 °C, so the scenarios could not show it. An audit made weeks after

the paper, when a predictive option was added, found it in runs at the ends of the setting ranges, and a high-limit override that acts in every mode now holds the air under the limit.

An override that did not let go. On the coater a zone that joined the speed-limited set at a low speed was then held at its target inside the hold band and never left, so the limit stayed down for good. One run of case 3 rested at 28.1 m/min, and sweeps of the moment of the setpoint change found 10 of 81 and 22 of 82 runs resting under 31 m/min. A release rule brought both sweeps to none. A second rule fixed a deep cut in which the upstream zone took the speed handle as well and the line settled near 21.5 m/min instead of 27.6 in 3 of 8 seeds, and in 6 of 8 at half the step. A third candidate rule was dropped because it let the line run on cut heat at 44.7 m/min. Both defects showed away from the case as shipped, at another moment of the change or another seed, and neither rule changed the challenge's runs, so the graded case could not have shown them.

A clamp that ignored its limit. The CO_2 cycle's turbine inlet setpoint was clamped at a constant 618 °C, two degrees under the default metal limit. With the learner's limit lowered to 605 °C, scenario 8 spent 452.8 min of 500 above it, and the predictive controller's bound was the same constant. An audit that moved the limit away from its default found it. The clamp now follows the limit the run declares, and no run at the shipped limit moved, which is why no shipped scenario had shown it.

A supervisor that pushed the wrong way. The extruder's supervisor treated a light product like a dense one and lowered the feed for both. Less feed puts more work into each kilogram and makes the pieces lighter still, so on the light side the correction drove the violation it answered. The coordinator measured a one-sided fix in a scratch copy, found that it moved the challenge's headline result and doubled the off-specification product of the light-lot case, from 96.6 to 192.8 kg, and kept the condition as implemented, stated in the chapter.

A score that skipped the violation. The peg's scorecard recorded its extremes only while the peg stood. On the day a rate defense broke the peg at a 15 % overnight rate, the tile still read 5.00 %. Once fixed, the challenge scorer counted the exit day, and seven rows of its table moved. Each of these five errors lay in a protective element or in the measure that judges one, and most showed only away from the cases as shipped.

The book round. The thirteen chapter records list 179 rows of numeric verification. Most chapters showed no drift. The coater's case 3 starting values, once printed as 18.5 and 18.3 % LEL where runs give 18.25 and 18.57, and the shield's settlement after one stop, 1.76 mm (1.78 only once the second ring build has begun), were corrected in paper and chapter, and three chapters quoted counts of app checks that have since grown. The CO_2 cycle's Conclusion disagreed with its own tables in five figures and now uses the table values. The agent writing Chapter 1 flagged a sentence that ranked the quadruped's vertical force above its horizontal force, where the two weights are equal, and the chapter pass reworded it. At the end of that round the editor's list of defects held eleven entries for this volume.

Every error above was caught by a check that could fail: a setting driven to the end of its range, a sweep of timing and seed, a limit moved away from its default, a score compared with the logged states. For constraint control the record adds one rule to those of the shared section. Test every protective element away from the nominal case, and ask of each override when it lets go.

15.5 Gaps in Override and Constraint Control

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow in Section 15.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the

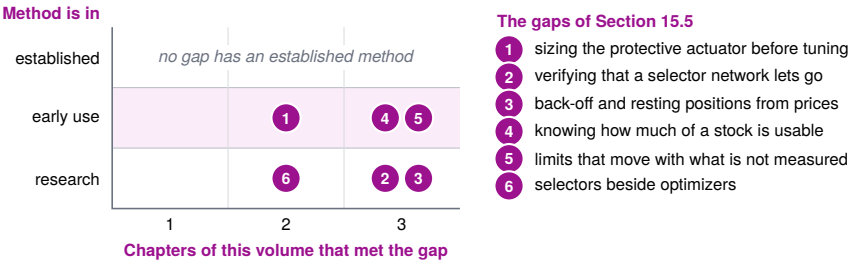


Figure 15.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 2 and 9), three for gap 2 (6, 11 and 3), three for gap 3 (10, 4 and 13), three for gap 4 (12, 8 and 14), three for gap 5 (7, 11 and 6) and two for gap 6 (4 and 10). Up 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 four gaps were met by three chapters each. Of those four, the two in early use, usable stock and moving limits, rest on a rule of thumb and on prediction models that their own sources call variable or inaccurate.

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. *Sizing the protective actuator before tuning it.* The compressor's slow valve and the trigeneration plant's lost capacity were found by simulation, case by case. A review of anti-surge control notes that sensor delay and slow valve response often limit the control algorithm, that faster actuators add cost and complexity, and that strategies shown on test benches or in simulation must still be validated on full-scale compressors with safety mechanisms [23]. The gap is hard because a saturated actuator looks on a trend like a slow tuning, as the compressor's did when no gain changed the outcome. A method that states the stroke time and capacity a protection needs against a stated upset is not standard.

2. *Verifying that a selector network lets go.* The coater's resting points were found only by sweeping the timing of a change over 81 runs, and the shield and the turbofan carry networks of five and eight limits. The Open challenge box after this list takes the gap further.

3. *Setting back-off and resting positions from prices.* The resting bypass, the engine limit of the combine and the planned friction of the quadruped were each chosen by a sweep against a scorecard. A design study of the same kind of cycle chose the throttle valve's pressure drop so that the turbine bypass valve stays around 10 % open during throttling [24], and a report on the dispatch of combined cooling, heating and power found savings of 8 to 25 % that depend on pricing, season, forecast accuracy and grid constraints [20]. The gap is hard because the price of a margin moves with markets and seasons while the margin is set once, at commissioning. Methods that set margins from prices, and move them as prices move, are thin.

4. *Knowing how much of a stock is usable.* The fan car's tank ran out by pressure with nitrogen inside, the startup's credit ceiling moved with its earnings, and the peg's date was set by credit the defense did not touch. For reserves, a central bank task force found the ratio of short-term debt the most popular adequacy indicator but not the only one, and gave three to four months of imports as a rule of thumb [25]. At a later meeting of central bank officials, some participants noted that the adequate level can vary considerably over time, particularly in crises [26]. The gap is hard because what can be used depends on the state in which it is used, a tank pressure, a trailing earnings figure or the confidence of a market. The usable runway is still modeled plant by plant.

5. *Limits that move with what is not measured.* The drill's pressure limit followed the cover and the soil, the shield's target rose under the river, and the coater's binding zone moved with the drying front. A thesis on drilling found that the commonly used Bingham plastic model overestimates the annular pressure and that the commonly used cavity expansion method significantly overestimates the allowable pressure, which can lead to hydraulic fracture [27]. The gap is hard because such a limit is often learned only when it is crossed. A limit estimate that states its own uncertainty, so the back-off can follow it, is not yet routine.

6. *Selectors beside optimizers.* The combine's MPC collected more grain in the challenge and failed it on overload, and the CO_2 cycle's MPC kept limits that depend on inner loops tuned as its model assumes. A review of learned legged locomotion notes that safety filters have been studied extensively in control but applied to locomotion only in a limited way, and that more research could help certify such controllers [17]. The gap is hard because the optimizer and the override see the same limit through different models, and either can be wrong. A rule for which layer owns which limit, and how an override backs up an optimizer whose model is wrong, is still written per plant.

> ⚠ **OPEN CHALLENGE** **An Override Network That Is Proved to Let Go**
>
> Gap 2 runs through the coater, the shield and the turbofan. A fieldbus standard gives the control selector a back-calculation signal, so a controller it de-selects halts integral action [28]. A unified framework treats an override as a substitution of the plant input, which wrongly updates the controllers not chosen and bumps the input, and shows every known linear anti-windup scheme as a special case [29]. In the editor's view the gap is hard because release depends on the order and timing of the limits, which no linear framework covers. Progress would be a check that lists each combination of active limits and shows its exit.

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 [30]. 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 [22]. 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

- A protective loop was bounded by its actuator before its tuning, and external reset decided when an override took over.
- Every override, resting position and priority had a measured price in throughput, fuel or gas.
- A stock that depletes made time a limit, and protections worked where they read what actually runs out.

References

[1] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[2] A. Reyes-Lúa and S. Skogestad, "Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy," *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[3] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem: A review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[4] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[5] T. A. Johansen and T. I. Fossen, "Control allocation: A survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[6] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the MIT Cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–9.

[7] 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.

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

[9] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[10] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-Jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems*, vol. 43, no. 5, pp. 137–177, 2023.

[11] L. C. Jaw and S. Garg, "Propulsion control technology development in the United States: A historical perspective," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2005-213978, 2005, nTRS 20050238468.

[12] S. Garg, "Aircraft turbine engine control research at NASA Glenn Research Center," NASA Glenn Research Center, Cleveland, OH, Technical Memorandum NASA/TM-2013-217821, 2013, nTRS 20130013439.

[13] G. Torrisi, S. Grammatico, A. Cortinovis, M. Mercangoz, M. Morari, and R. S. Smith, "Model predictive approaches for active surge control in centrifugal compressors," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 6, pp. 1947–1960, 2017.

[14] R. Ponciroli and R. Vilim, "Analyses of model-based predictive control for a S-CO_2 Brayton cycle power converter," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep., 2016.

[15] U.S. Department of Energy, "Use lower flammable limit monitoring equipment to improve process oven efficiency," National Renewable Energy Laboratory for the U.S. Department of Energy, Industrial Technologies Program, Golden, CO, Process Heating Tip Sheet #11 DOE/GO-102007-2511, Oct. 2007.

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

[17] S. Ha, J. Lee, M. van de Panne, Z. Xie, W. Yu, and M. Khadiv, "Learning-based legged locomotion: State of the art and future perspectives," *The International Journal of Robotics Research*, vol. 44, no. 8, pp. 1396–1427, 2025.

[18] K. Darrow and B. Hedman, "Opportunities for combined heat and power in data centers," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2009/052, Mar. 2009, prepared by ICF International.

[19] S. Hyon, J. G. Hale, and G. Cheng, "Full-body compliant human–humanoid interaction: Balancing in the presence of unknown external forces," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 884–898, 2007.

[20] S. Samuelsen and J. Brouwer, "Novel controls for time-dependent economic dispatch of combined cooling heating and power (CCHP)," University of California, Irvine, Irvine, CA, Final Report DOE-UCI-01113, 2013.

[21] 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.

[22] 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.

[23] J. Du, Y. Zhang, M. Martínez García, and A. Spencer, "Review on compressor surge monitoring, modeling, and anti-surge control approaches," *Journal of Dynamics, Monitoring and Diagnostics*, vol. 4, no. 4, pp. 292–307, 2025.

[24] A. Moisseytsev and J. J. Sienicki, "Development of the ANL plant dynamics code and control strategies for the supercritical carbon dioxide Brayton cycle and code validation with data from the Sandia small-scale supercritical carbon dioxide Brayton cycle test loop," Argonne National Laboratory, Argonne, IL, USA, Tech. Rep. ANL-ARC-218, 2011.

[25] International Relations Committee Task Force, "The accumulation of foreign reserves," European Central Bank, Frankfurt am Main, ECB Occasional Paper 43, 2006.

[26] M. S. Mohanty, "Market volatility and foreign exchange intervention in EMEs: what has changed? An overview," in *Market Volatility and Foreign Exchange Intervention in EMEs: What Has Changed?*, ser. BIS Papers. Basel: Bank for International Settlements, 2013, no. 73, pp. 1–10. [Online]. Available: https://www.bis.org/publ/bppdf/bispap73.htm

[27] A. Rostami, "Prediction and evaluation of annular pressure in horizontal directional drilling," Ph.D. dissertation, University of Alberta, Edmonton, AB, 2017.

[28] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[29] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[30] 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.

[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

After a surge on the slow-valve case of Chapter 2, an agent proposes raising the anti-surge controller's gain and shortening its integral time, because "the controller reacted too slowly". Using the chapter, explain why the proposal will not prevent the surge, list the checks you would ask for first, and name the changes that did hold the closure.

15.2 Discussion · 8 min

A design review lists integrator tracking as an optional refinement of a selector network. Using Chapter 3 and Chapter 11, argue that the anti-windup arrangement is part of the protection, and state what a reviewer should ask about it on a new selector.

15.3 Discussion · 8 min

A manager asks that the resting bypass of the CO_2 cycle in Chapter 10 be set to zero to save heat. Using that chapter and Chapter 9, answer the manager, and describe how you would set a resting position on a new plant.

15.4 Discussion · 8 min

Using Chapter 14, Chapter 12 and Chapter 8, write three requirements for a protective layer on an actuator that depletes, and support each with a measured result.

15.5 Discussion · 9 min

An agent delivers a new override network for a coating line with a vapor limit, and its run of the graded challenge passes. Using the cases of Section 15.4.1, list the tests you would require before the network is accepted, and say which defect each test would have caught.

15.6 Discussion · 8 min

A project proposes one model predictive controller in place of every selector and override. Using Chapter 2, Chapter 4, Chapter 5 and Chapter 10, argue where it should take over, where the selectors should stay, and how to arrange the two.