INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 5 Delays, Slow Measurements, and State Estimation

CHAPTER 14

Summary and Outlook

Pages 344–357

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Delays, Slow Measurements, and State Estimation*, John D. Hedengren, Ed., vol. 5 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 14, pp. 344–357.

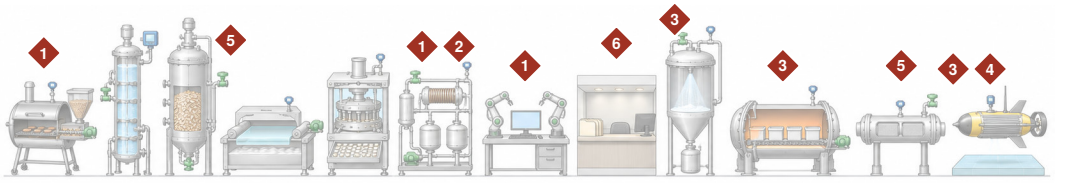
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v05-delays-estimation/

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

CHAPTER 14

Summary and Outlook



Open gaps 1 moving delays 2 corrections per material 3 unobserved directions 4 gate recovery 5 model resolution 6 calendar decisions

Notice that the gap of moving delays joins the grill, the flow synthesis skid and the teleoperator, while the gap of unobserved directions joins the spray dryer, the kiln and the survey vehicle. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Late Measurements Taught**

The main idea. Every chapter of this book ran something that learns about itself late, or never, from a grill to a pension fund and a vehicle under the sea. Some designs waited for the news. Others predicted it with a model and corrected the model when the news came.

The challenge. A prediction is only as good as its idea of how long the news takes. An estimate cannot fix what no instrument sees, and some loss is already on its way before any controller can act.

Why it matters. Asking what each estimate read, and testing it where the model is wrong, separated the designs that worked from those that failed.

Chapter 1 asked two questions of every plant in this volume. What does each delay or missing measurement cost, and what recovers it? Twelve chapters answered them on a grill, process lines, a teleoperator, a pension fund, an exhaust line and a survey vehicle. This chapter sets the answers side by side, draws a short guide for choosing a method, 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.

14.1 What the Chapters Showed

Table 1.4 named the delay or missing measurement of each chapter and the remedy, without results. Tables 14.1 and 14.2 return to it with the measured result. The first holds the chapters whose measurement is right but late, through a fuel line, an analyzer, a vessel, a ribbon, a reactor or a network. The second holds a valuation a year old and four plants whose key variable no instrument reads. Each row names a structure that recovered the cost, one that did not, and the result with the table or section that reports it. Figure 14.1 sets one like pair of numbers from each row side by side.

The columns of the two tables make five points.

A predictor is only as right as its delay. The MPC of Chapter 2 overshot the grill's lid opening by 2.7 °F when it assumed the lid would shut and by 92.8 °F when told it would stay open, with the same opening in both runs. A model auger delay of 0 or 60 s against the true 20 s cost almost nothing, because the heat bias absorbed it, while 120 s doubled the overshoot (Section 2.6). The predictive controller of Chapter 3 omitted the analyzer lag and stopped settling at 1.15 min of lag, where the PI loop still decayed. On the flow synthesis the plug-flow delay, timed from a flow the air lock had

Table 14.1 What Each Delay Cost and What Recovered It, Chapters 2 to 8

Chapter	Delay	What recovered it	What failed	Measured result
2, pellet grill	A 20 s auger, then ignition and burning. A lid that opens.	A lid-aware hold of the feed. An MPC whose filter carries the fuel in transit.	The PI tuned from the step test. The MPC told that the lid stays open.	Overshoot after the lid 9.7, 3.0, 2.7 and 92.8 °F (Table 2.6). Challenge in band 70.1 shipped, 92.5 % with the MPC (Table 2.7).
3, column	An analyzer lag up to 2 min on a column with 0.21 min of apparent dead time.	Detuning to $K_c = 36$, $\tau_I = 5$ min. A predictive controller with a disturbance state.	One tuning carried across the lags. A predictive model that omits the lag past 1.15 min.	Overshoot 7.6 to 30.7 %, settling 13.5 to 23.6 min (Table 3.4). Feed upset IAE 0.026 against 0.061 (Table 3.5).
4, digester	Four hours from heater to blowline. Three instruments on three clocks, one drifting.	H-factor feedforward. The combination held on an estimate reconciled with the laboratory.	Residual alkali held on a new wood. Feedback on the raw drifting analyzer.	Loss \$119.68 against \$7.27 per tonne (Table 4.3). Challenge passed at 0 of 11 seeds shipped and 11 of 11 reconciled (Table 4.5).
5, float line	Thickness read 15.2 min after forming. Stress settled before the polarimeter.	A forming model that makes the gauge loop a Smith predictor. Thickness by mass balance at the camera.	The same tuning with the model off, which hunts with a period of about 54 min.	Rejected 734 by the gauge, 212 by mass balance, 1827 m^2 with no model (Table 5.6). Challenge accepts 1367 against 8745 m^2 (Table 5.7).
6, tableting line	Dose fixed at the feeders twelve minutes before the press. A refilling feeder. A late analyzer.	Refill and humidity feedforward. A tracker that closes the gate on what the analyzer will judge late.	Tracking off with a 40 s analyzer delay. The line as shipped.	Under-dosed tablets accepted 816 without tracking, none with it (Table 6.5). Challenge out-of-limit accepted 194 shipped, none tuned (Table 6.6).
7, flow synthesis	Delay and spread that change with throughput. A drum only the analyzer sees.	Diversion on residence-time quantiles counted in volume. A coil longer than the detection time.	A plug-flow delay timed from a changed flow. A 10 min method at 100 mL/min.	Air lock: 1.53 L diverted, none released, against 177 mL released (Table 7.5). Drum escapes 10 of 12 times (Table 7.6).
8, teleoperator	A round trip up to 2.6 s, with jitter and loss, inside a force loop.	Wave variables, a passivity controller, or model mediation once its model is right.	Full force reflection. Reflection cut to 0.2. Waves and passivity together.	Challenge reversals and peak force 41 and 108.8 N shipped, 15 and 6.42 N waves, 3 and 2.00 N passivity, 12 and 4.08 N mediation (Table 8.6).

changed, released 177 mL that tracking in volume diverted. A delay kept as a volume of material, as Section 1.4.4 recommends, stays right when the flow moves.

Follow the material, not the clock. Every process chapter that met its target tagged material with what was done to it. The digester carried its chips as slugs advanced on the measured flow. The float line tagged each parcel of ribbon with the speed and thickness of its birth. The tableting line's tracker closed the gate on 816 tablets that the late analyzer would have passed, and the kiln's corrected feedforward waited for the affected material to reach each zone. The rule that sized the synthesis coil, detection time against holding time, predicted the escapes: none where the coil held 20.83 min and 10 of 12 where it held 12.50 min against 20.5 min of detection.

Table 14.2 What Each Late Result or Unmeasured State Cost, Chapters 9 to 13

Chapter	Late or missing	What recovered it	What failed	Measured result
9, pension fund	A valuation a year old, one decision a year, a contribution cap.	A predictive plan on a projected estimate, with a funding buffer and a glide path.	The shipped amortization rule. A seven-year deadline under a 12.5 % cap.	The year-old valuation raises the contribution SD from 1.91 to 2.35 points (Table 9.5). Challenge met on 0 of 5 seeds shipped, 5 of 5 suggested (Table 9.7).
10, spray dryer	Moisture from the laboratory 40 min late. A sticky margin never measured.	A psychrometric soft sensor whose laboratory bias is aligned with each sample, and feedforward.	The outlet loop alone at fixed setpoints. The soft sensor without its correction.	Case 4 off specification 0.57 shipped, 1.56 t uncorrected (Table 10.7). Campaign 1.98 to 3.88 t in specification (Table 10.9).
11, kiln	The brick's core never measured, and 30 h of transport.	A core cascade on a two-node soft sensor. A supervisor that predicts the exit quality.	The same cascade on the surface. Zone loops that hold a biased thermocouple perfectly.	Exit ordering 0.9900 core, 0.9834 surface. Good rate 6.07 to 8.14 t/d. Quality 96.6 with the bias against 98.6 (Table 11.4).
12, urea dosing	Stored ammonia never measured. A sensor that reads ammonia as nitrogen oxides.	Storage control on a model-based observer at a 100 % level.	A trim on the tailpipe sensor. Every feedforward ratio tried.	Trim 0.293 g/kWh with a 1301.9 ppm slip peak. Storage 0.423 and 23.5 ppm, passing. Feedforward at 0.6: 1.466 g/kWh (Table 12.5).
13, survey vehicle	Position never measured under water. Acoustic fixes late and sometimes wild.	An extended Kalman filter that applies each fix at its instant behind a gate.	Dead reckoning. Fixes applied at arrival. The USBL alone switched on.	Error RMS 0.56 at the instant, 1.92 m at arrival (Table 13.5). Survey 83.86 to 99.95 % imaged, 79.68 with the USBL alone (Table 13.6).

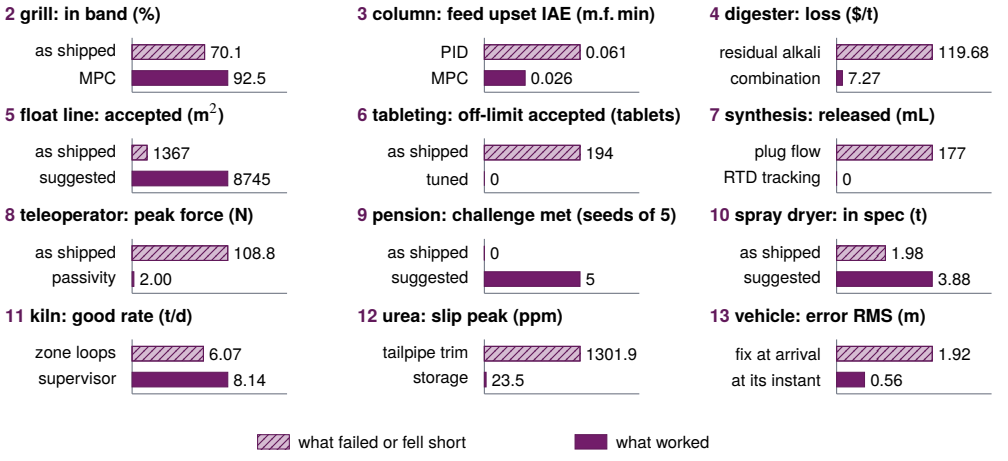


Figure 14.1 One like pair of numbers from each row, as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. Lower is better on every panel except the grill, the float line, the pension fund, the spray dryer and the kiln. Notice that no panel compares two gains: each better number came from a change of structure or of what the controller was told, such as a storage observer, a passive channel or a fix applied at its instant.

A late result should correct a model, not drive the loop. The float line's mass balance saw a bath drift 2.3 min after forming instead of 15.5 min and cut the rejected glass from 734 to 212 m^2. The spray dryer's soft sensor, corrected by results 40 min late, took a campaign from 1.98 to 3.88 t in specification, while a loop on the laboratory alone oscillated with growing amplitude. On the digester the reconciled estimate held kappa within 28.68 to 29.51 under a drifting analyzer that drove the raw loop to 26.52 (Section 4.6). The same chapter adds a caution. Its estimate's gate held back a large new offset for 75 minutes, so on a day without the drift the raw analyzer held kappa within ± 2 at 11 of 11 seeds and the estimate at none. An estimate built for a bias is slow as the input of a loop.

What no measurement sees stays wrong. The kiln's zone loops held a thermocouple that read 15 °C high perfectly and cost the product 2.0 points of quality. The spray dryer's laboratory bias corrected the moisture and left the sticky margin 0.9 K optimistic (Section 10.6). Over the survey vehicle's canyon a current carried it 21.46 m from its estimate and left a holiday 25 m wide that no tuning recovered (Section 13.5.5). In each case the estimate was consistent with every measurement it had. Section 1.4.6 names the cause: a direction the measurements cannot see is corrected by nothing.

Some cost is set before any controller acts. Chips already past the digester's heater at a production step could not be rescued by any heater setting, and their worst slugs ran $+2.46$ to $+4.49$ off kappa (Section 4.4). A failed annealing zone sent 915 to 918 m of glass over its limit under every structure (Table 5.10). The teleoperator's stable reflection gain fell toward 0.21, the hand's stiffness over that of the slave and wall in series (Section 8.5). The pension fund's cap made a seven-year deadline impossible for every amortization period (Section 9.9). Each chapter met such a limit with equipment or structure, a diverter, a passive channel or a buffer, and not with gain.

14.2 Choosing a Method

The chapters suggest an order of questions for a loop that sees its plant late. Each names a chapter where the answer decided the result.

1. *What is the delay made of?* Separate transport from lags that a fit reports as dead time and from the age of a held result, $\theta_e = \theta + t_d + T_s/2$ (Equation (1.8)). The grill's 20 s auger fit as 4.7 or 2.8 min, and the column's 0.21 min contains no delay (Chapter 2, Chapter 3).
2. *Does the delay move with flow?* Keep it as a volume or track the material, and never time it from a flow an upset has changed (Chapter 7, Chapter 4).
3. *Can a model or a faster instrument stand in between results?* Put a soft sensor upstream of the delay and correct its bias with each slow result, compared with the estimate stored for the same material (Equation (1.14), Chapter 5, Chapter 10).
4. *Can the measurements separate what matters?* Test observability before building the estimator, and add the measurement or the maneuver that restores it, as the survey vehicle's opposed lines do for its compass bias (Chapter 13, Chapter 11).
5. *Is part of the upset already past the actuator?* Then hold or divert it, and size the holding equipment from the detection time (Chapter 7, Chapter 6).
6. *Does the loop decide on a calendar from a late statement?* Keep the correction per decision near the bound of Equation (1.9) and add a projection across the lag (Chapter 9).
7. *Does the delay sit inside an energy exchange with a person?* Keep the channel passive first and recover feel second (Chapter 8).

Figure 14.2 draws the seven questions in order.

A predictive controller paid where its model carried something the loops could not: fuel in transit, a disturbance state or a preview. It held the grill's challenge at 92.5 % in band where the best lid-aware

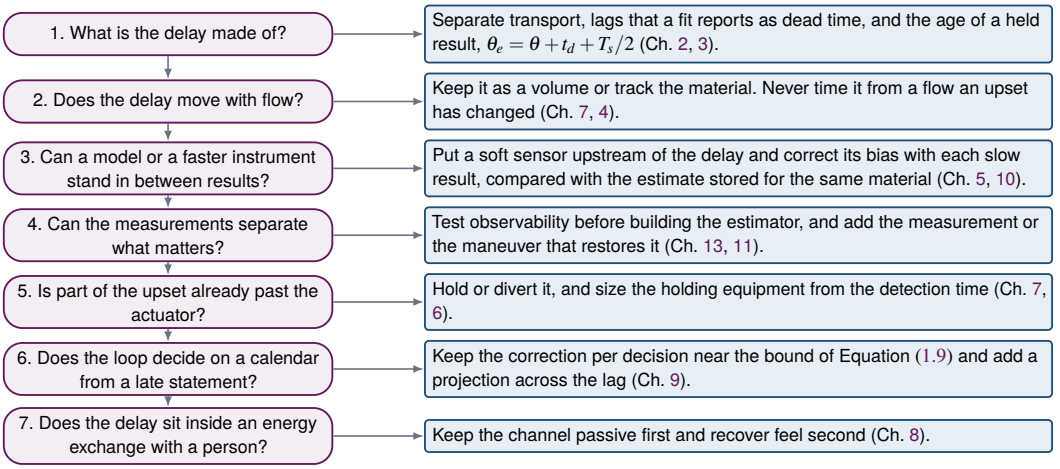


Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that what the delay is made of, and whether it moves with flow, come before any estimator or predictor is chosen: a wrong delay cost the column's predictive controller and the flow synthesis's plug-flow timing.

PI reached 90.7 at the top of its gain range (Table 2.7), and it rejected the column's richer feed with an IAE of 0.026 against 0.061. Its advantage often came from the estimate rather than the optimizer. On the spray dryer it shipped 0.444 t off specification against 0.570 t for the loops, and the moisture loop retuned alone shipped 0.413 t (Table 10.8). On the flow synthesis the retuned trim diverted nothing through the rate change, as the MPC did (Table 7.8). Without preview the urea MPC failed its slip peak at 34.6 ppm where storage control passed (Table 12.6). Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Prediction moves inside larger controllers. Dead-time compensators, from the Smith predictor to its robust extensions, were surveyed in 2008 [1]. Bilateral teleoperation has its own historical survey, in which the delay of the channel is the central problem [2]. A 2016 survey traces model-mediated teleoperation from the late 1980s, a method developed to keep a teleoperator both stable and transparent under arbitrary communication delays, and names the challenges of building a reliable one [3]. In the editor's view, the predictor of this volume is becoming the model inside a predictive controller or a mediated interface, and the open work is keeping its delay right when flow, throughput or a network moves it.

Estimation takes late and out-of-order data as normal. In multisensor tracking, measurements from one target arrive out of sequence, and an exact state update for such arrivals was published in 2002, where earlier solutions were approximate [4]. A real-time algorithm for moving horizon state and parameter estimation was published in 2011 [5]. A 2014 review of underwater vehicle navigation notes that radio signals and satellite positioning do not reach the vehicle and that acoustic communication is low in bandwidth and unreliable. It finds that inertial sensors, installed beacons and periodic surfacing are fundamentally limited, and that advances in acoustic communication and in simultaneous localization and mapping open new possibilities [6]. The survey vehicle of Chapter 13 sits at the start of that path, with fixes stored and applied at their instant.

Soft sensors are learned, and material is tracked by its distribution. A 2021 survey reports that new characteristics of industrial processes degrade traditional soft sensor models and reviews deep learning models for data-driven soft sensors, with the problems met in practical applications [7]. In pharmaceutical manufacturing, regulators describe support for continuous production through science- and risk-based approaches, built on quality by design and process analytical technology [8]. Residence time distributions have been proposed as the basis for defining a batch, tracing raw materials and setting sensor frequency on a continuous line [9], the role they play in Chapter 6 and Chapter 7. In the editor's view, a learned soft sensor still needs a laboratory correction aligned with the material it describes, which is the lesson of Section 1.4.8 and not a property of any model class.

> ⧖ HISTORY **1949 to 2021: From Tongs Behind Glass to Semi-Autonomous Operation**
>
> The timeline of Chapter 1 places a remote hand in 1949. The report of that year described tongs worked behind a shielding barrier, such as a lead glass window, through an entirely mechanical master-slave linkage that an operator learned in five to fifteen minutes [10]. A later study of bilateral force reflection for a space robot found that delays as small as 10 ms, from communication or computation, can significantly increase the stiffness and damping the operator must supply [11]. A 2021 review of remote intervention in space reports that force reflection has proved impractical and unstable under small to medium delays, that wave variables restore stability without feeling intuitive, and that many researchers have turned to semi-autonomous teleoperation [12]. In the editor's reading, the passive channel and the mediated model of Chapter 8 stand between the two ends of that arc.

Hidden states gain instruments of their own. On an exhaust catalyst, ammonia and nitrogen oxide sensors give only an inference of the ammonia stored, confounded by the oxide sensor's cross-sensitivity to ammonia. Radio-frequency sensing uses the catalyst housing as a microwave resonant cavity, its readings have correlated highly with the stored inventory, and earlier studies showed its potential as the feedback of closed-loop urea dosing [13]. A thermocouple gives no indication in use that its output has drifted unless it is removed for recalibration, and a review of drift describes in situ calibration against the melting points of built-in reference materials or by Johnson noise thermometry [14]. The urea doser of Chapter 12 estimates its store with an observer, and the electric kiln of Chapter 11 reveals its biased thermocouple only through the product. In the editor's view, a hidden state that gains an instrument turns part of gap 3 of Section 14.5 into a question of validating that instrument.

Navigation relies on its estimator where no fix reaches. A survey of underwater navigation reports that one deep survey vehicle carried out 191 surveys of the seabed at mid-ocean ridges, at an average depth of 2000 m and with a navigation precision of a few meters. The same survey finds surface navigation effectively solved by satellite positioning, while in the middle of the water column acoustic time-of-flight systems and inertial units are the only effective sensors of horizontal position, and it expects model-based state estimators to be central to the advances there [15]. The survey vehicle of Chapter 13 meets that zone over its canyon. In the editor's view, fixes stored and applied at their instant are the part of its filter that practice will keep.

Soft sensors are maintained against the laboratory. Most industrial predictive control products surveyed in 2003 corrected their predictions with the difference between the measured and the predicted output. That feedback assumes a step disturbance at the output that stays constant, and it rejects disturbances entering at the process input sluggishly [16]. A review of industrial soft sensors holds that a soft sensor should be updated as the process characteristics change and that recursive updates cannot cope with abrupt changes. It reports an industrial application in which a bias between measurement and estimate was unavoidable after the operating condition changed [17]. The bias weight of Chapter 10, which trades speed against noise, and the persistence gate of Chapter 4 are

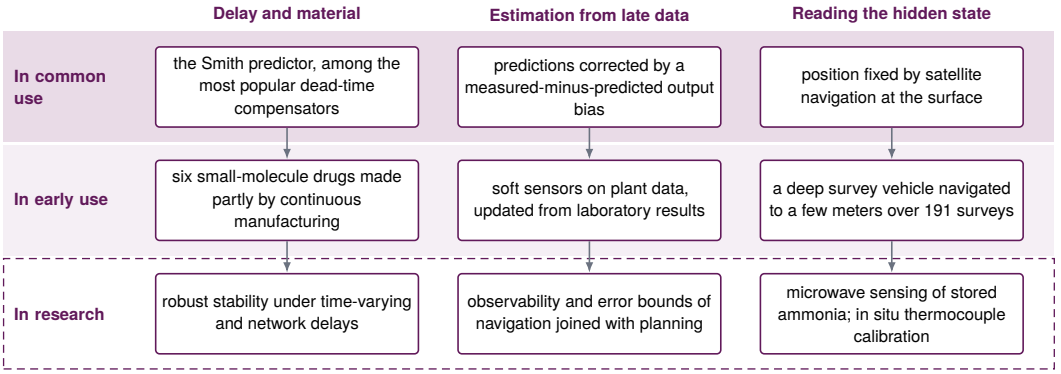


Figure 14.3 Developments in three threads. The top lane holds what the cited source calls popular, solved or used by nearly every product surveyed, the middle lane what it reports in service or on plant data, and the bottom lane what it reports from research, as a proposal or as a need, in the editor's reading. Sources by thread: delay and material [18, 19], estimation from late data [15–17] and reading the hidden state [13–15]. Notice that each thread moves from a fixed correction toward an estimate or a sensor that must itself be checked while the plant runs.

two answers to that tension. In the editor's view, the update rule of a soft sensor is part of its design. Figure 14.3 sets these developments in three threads.

14.4 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 [20]. 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 [21]. 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.

14.4.1 How This Volume Was Made

Each of the twelve 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 seven of the twelve 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. The twelve logs give 163 entries a severity. Of these, 37 are wrong behavior of an app, 90 are documentation, 27 are cosmetic or display, and 9 are of other kinds. The kiln alone holds 10 of the 37, and its log opens with nine defects found while its paper's numbers were being measured. Five cases, chosen because each lies in a predictor, a measurement record or an estimator, show what the agents got wrong.

Estimators that knew the plant. Five logs record an estimator, a soft sensor or a controller's start that read the run's hidden settings or the plant's true state. The flow synthesis's soft sensor ran on the run's own kinetics. With a rate constant halved it predicted the sample point's 95.934 % assay exactly, before any analyzer result, where a soft sensor on the declared model predicts 98.659 %. The grill's estimator started from the plant's true state, and 18 of 33 parameters edited before a decision moved the MPC's first move with every measurement unchanged. The pension fund's projection read the run's market returns, which no valuation had reported. No shipped run showed any of them, because a shipped run leaves every hidden value at its declared one. Audits found them by changing a hidden value before Run and asking whether the estimate moved while the measurements did not.

A feedforward that ignored its own delay. The kiln's first throughput feedforward corrected all ten zones the instant the load cell moved, although the material needs up to 1710 min to reach the last zone, and it scaled the whole design duty of each zone. It raised the peak zone error after a density change from 0.17 to 0.98 °C and the kiln-wide error from 4645 to 11 204 °C min, while the scenario's text claimed the opposite. The paper agent found it by measuring the case the text

described. The corrected feedforward waits for the material, and the chapter reports the honest result that a well-tuned zone loop leaves it little to do.

A reading credited to the wrong glass. The float line's inspection credited each parcel with the gauge's filtered output at the first scan after the parcel reached it. That reading mixed the parcel with its neighbors. With the thickness loop hunting, adjacent parcels differed by up to 0.19 mm, the credited reading by up to 0.086 mm against a guard band of 0.01 mm, and four parcels truly out of tolerance were accepted. The defect was first documented as a limitation of one configuration. A sweep of 20 noise seeds then showed the same mechanism failing the suggested challenge at 2 of them, and each parcel is now credited with a reading of its own glass.

A gate that locked itself out. With the survey vehicle's shipped navigation switched to the acoustic fixes alone, the first fix to arrive on the default seed was an outlier 24.2 m off. The filter's uncertainty was still wide, so the gate accepted it and read most of the error as compass bias, -5.41° against a true $+1^\circ$. Its covariance then collapsed and never reopened, and 132 later fixes were refused. The survey imaged 79.68 %, less than dead reckoning. The behavior was kept as a lesson and the texts now say so. The referee then found that the paper called all 132 refused fixes honest, where three were outliers.

Algebra that the simulation overruled. The agent that wrote Chapter 1 drafted two derivations that were wrong until its script disagreed. The laboratory loop's characteristic equation had one delay step too many, and the Smith predictor's delay tolerance looked larger than it was because an output clamp turned a growing oscillation into a bounded cycle. Bisection by simulation with the limits removed exposed both. The spray dryer's referee found a claim about the laboratory weight that was reasoning, not a measurement, and the case was rerun at three weights before the paper kept it.

The book round. The twelve chapter records list 183 rows of numeric verification, and ten report no drift. The digester's text said a trim held its analyzer near 30 where its own figure shows 32.1 to 34.5, and the wording now follows the figure. The urea chapter said no case's sensor reading crosses 1.2 g/kWh, where the shipped challenge reads 1.689. The paper and the chapter now report 1.689 and 2.876 g/kWh. The tableting paper named its controller actions the opposite way to the series convention, and both its chapter agent and the agent writing Chapter 1 flagged it. At the end of that round the editor's list of defects held nine entries for this volume.

Every error above was caught by a check that could fail: a hidden value moved, a case run as its text described it, a sweep over seeds, a scenario changed by one setting, a derivation simulated with its limits removed. For estimation the record adds one rule to those of the shared section. Before trusting an estimate, ask what it read, and test it on a case where the model is wrong.

14.5 Gaps in Delay Compensation and State Estimation

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow in Section 14.5.1. 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 14.4 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. *Delays that move with what is not measured.* The flow synthesis's plug-flow delay was wrong once an air lock changed the flow, the grill's pulsed auger lengthened its transport as the feed fell, and the teleoperator's channel jittered. A review of dead-time compensation names the conservativeness of robust stability criteria for time-varying delay as a topic that deserves further investigation [18]. A survey of model-mediated teleoperation calls environment modeling its

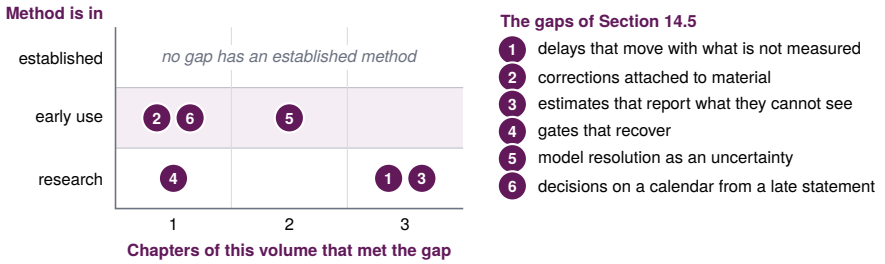


Figure 14.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 7, 2 and 8), one for gap 2 (7), three for gap 3 (11, 10 and 13), one for gap 4 (13), two for gap 5 (12 and 4) and one for gap 6 (9). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the two gaps met by three chapters each, moving delays and unobserved directions, are both still in research.

foremost challenge, because a sudden change of the local model produces a sudden force [3]. The gap is hard because the delay moves with the very flow or network that the upset disturbs. Online delay estimates that state their own uncertainty are not routine.

2. *Corrections attached to material.* The flow synthesis applies one analyzer bias to every batch in the line, and under the shipped supervisor rule its trim released 291 mL after an unassayed drum 15 % strong. Regulators writing on continuous manufacturing note that poor measurement is often due to the sampling system rather than the analyzer, and suggest using the residence time distribution to set the sampling frequency [8]. A study of traceability finds that the back-mixing a robust line needs confounds the boundaries between batches [9]. The gap is hard because a laboratory result describes material that has since mixed with its neighbors. Bias updates that belong to the material sampled, with weights chosen per stream, are still designed plant by plant.

3. *Estimates that report what they cannot see.* The kiln's surface-only correction, the spray dryer's optimistic margin and the survey vehicle's canyon drift all left the estimate consistent with every measurement. The Open challenge box after this list takes the gap further.

4. *Gates that recover.* The survey vehicle's filter locked out after one wild fix, and the chapter leaves a lockout monitor as further work. A survey of industrial predictive control found that most products detect only sensor failures beyond preset limits, while significant bias and drift within normal limits generally go undetected [16]. The gap is hard because, from inside the filter, a refused honest measurement and a wrong estimate look alike. Rules for reopening a covariance after repeated rejections are heuristic.

5. *Model resolution as an uncertainty.* The urea challenge passes on the declared three cells and fails on six or twelve (Table 12.7), and a grid of 40 cells smeared the digester's kappa by 2.64 %. The same survey found that, until performance and robustness are connected, it will not be possible to say when a model is accurate enough, and that all products but one rely on brute-force simulation to evaluate model mismatch [16]. The gap is hard because the resolution that decides a verdict is rarely the one at which the design was checked. A verdict that depends on the number of cells needs that dependence stated with it.

6. *Decisions on a calendar from a late statement.* The pension fund's plan passed only by contributing nearly twice the normal cost. An actuarial analysis found that a one-year delay in the accounts loses information and raises the resulting variances, that an optimal spread period still exists, and that spreads of 20 to 25 years were common practice [22]. A federal agency's projections treat future employer contributions to multiemployer plans as uncertain and a significant risk [23].

The gap is hard because the cap, the deadline and the delay act together and are set by different parties. Design rules for policies that decide once a year on a valuation a year old, under caps, are thin.

OPEN CHALLENGE Estimates That Report What They Cannot See

Gap 3 runs through the kiln, the spray dryer and the survey vehicle, whose estimates stayed consistent with every measurement while an unseen direction drifted. A survey of underwater navigation reports that researchers are only beginning to investigate properties such as observability and lower error bounds for navigation joined with planning [15]. A review of industrial soft sensors notes that a soft sensor can be checked only when analyzed data are available, and that normal changes are hard to tell from faults [17]. In the editor's view the gap is hard because an unobserved direction leaves no residual to alarm on. Progress would be an estimator that reports which directions no measurement has corrected and how far they may have drifted since.

14.5.1 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 [24]. 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 [21]. 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 [25]. 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 [26]. 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 [27]. 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 predictor was only as right as its delay, and a delay kept as a volume of material stayed right when the flow moved.
- A late result served best as the correction of a model aligned with the material it measured, not as the input of the loop.
- No estimator recovered a direction its measurements could not see, and cost already past the actuator was met with equipment or structure, not with gain.

References

[1] J. E. Normey-Rico and E. F. Camacho, "Dead-time compensators: A survey," *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[2] P. F. Hokayem and M. W. Spong, "Bilateral teleoperation: An historical survey," *Automatica*, vol. 42, no. 12, pp. 2035–2057, 2006.

[3] X. Xu, B. Cizmeci, C. Schuwerk, and E. Steinbach, "Model-mediated teleoperation: Toward stable and transparent teleoperation systems," *IEEE Access*, vol. 4, pp. 425–449, 2016.

[4] Y. Bar-Shalom, "Update with out-of-sequence measurements in tracking: Exact solution," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 769–777, 2002.

[5] P. Kühl, M. Diehl, T. Kraus, J. P. Schlöder, and H. G. Bock, "A real-time algorithm for moving horizon state and parameter estimation," *Computers & Chemical Engineering*, vol. 35, no. 1, pp. 71–83, 2011.

[6] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 1, pp. 131–149, 2014.

[7] Q. Sun and Z. Ge, "A survey on deep learning for data-driven soft sensors," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 9, pp. 5853–5866, 2021.

[8] S. L. Lee, T. F. O'Connor, X. Yang, C. N. Cruz, S. Chatterjee, R. D. Madurawe, C. M. V. Moore, L. X. Yu, and J. Woodcock, "Modernizing pharmaceutical manufacturing: From batch to continuous production," *Journal of Pharmaceutical Innovation*, vol. 10, no. 3, pp. 191–199, 2015.

[9] W. Engisch and F. Muzzio, "Using residence time distributions (RTDs) to address the traceability of raw materials in continuous pharmaceutical manufacturing," *Journal of Pharmaceutical Innovation*, vol. 11, no. 1, pp. 64–81, 2016.

[10] R. C. Goertz, "Master-slave manipulator," Argonne National Laboratory, Argonne, IL, Report ANL-4311, Mar. 1949.

[11] J. D. Chapel, "Performance limitations of bilateral force reflection imposed by operator dynamic characteristics," in *Proceedings of the NASA Conference on Space Telerobotics*, vol. 4. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology, 1989, pp. 91–100. [Online]. Available: https://ntrs.nasa.gov/citations/19900020524

[12] P. Kazanzides, B. P. Vagvolgyi, W. Pryor, A. Deguet, S. Leonard, and L. L. Whitcomb, "Teleoperation and visualization interfaces for remote intervention in space," *Frontiers in Robotics and AI*, vol. 8, p. 747917, 2021.

[13] P. Ragaller, J. Mandelbaum, L. Lapenta, A. Sappok, J. Pihl, V. Prikhodko, and J. Parks, II, "Direct measurement of ammonia storage on passive SCR systems for lean gasoline NOx reduction using radio frequency sensing," in *Proceedings of the ASME 2019 Internal Combustion Engine Division Fall Technical Conference*, no. ICEF2019-7224, Chicago, IL, 2019.

[14] R. J. Chesser and N. D. Bull Ezell, "Thermocouple and RTD drift due to neutron irradiation and temperature for space nuclear applications: A review," *IEEE Sensors Journal*, vol. 25, no. 22, pp. 40 965–40 977, 2025.

[15] J. C. Kinsey, R. M. Eustice, and L. L. Whitcomb, "A survey of underwater vehicle navigation: Recent advances and new challenges," https://robots.engin.umich.edu/publications/jkinsey-2006a.pdf, 2006.

[16] 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.

[17] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

[18] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[19] C. Badman, C. L. Cooney, A. Florence, K. Konstantinov, M. Krumme, S. Mascia, M. Nasr, and B. L. Trout, "Why we need continuous pharmaceutical manufacturing and how to make it happen," *Journal of Pharmaceutical Sciences*, vol. 108, no. 11, pp. 3521–3523, 2019.

[20] 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.

[21] 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.

[22] S. Haberman, "Pension funding with time delays and the optimal spread period," *ASTIN Bulletin*, vol. 25, no. 2, pp. 177–187, 1995.

[23] Pension Benefit Guaranty Corporation, "FY 2024 projections report," Report, listed on the agency's Projections Report page, last updated January 27, 2026, 2026, accessed 2026-11-21. [Online]. Available: https://www.pbgc.gov/sites/default/files/documents/fy-2024-projections-report.pdf

[24] 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.

[25] 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.

[26] 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.

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

14.6 Exercises

14.1 Discussion · 8 min

An agent proposes to shorten the closed-loop time constant of a Smith predictor on a dosing line whose throughput varies by a factor of two, because "the predictor removes the dead time". Using Section 1.4.3, Section 1.4.4 and Chapter 7, explain what the proposal risks and what you would ask for before accepting it.

14.2 Discussion · 8 min

A plant team wants to retire the spray dryer's moisture soft sensor and close the moisture loop on the laboratory alone, "because the laboratory is the measurement the customer pays for". Using Chapter 10 and Sections 1.4.5 and 1.4.8, argue for or against, and say how the laboratory result should be used.

14.3 Discussion · 8 min

In a demonstration an agent-built soft sensor on a reactor "tracks the plant exactly" from the first second of a run in which a kinetic constant was changed. Using Section 14.4.1, explain why that result is a warning, describe the test you would run, and say what a correct soft sensor would have shown.

14.4 Discussion · 8 min

The kiln of Chapter 11 and the survey vehicle of Chapter 13 both report confident estimates while the true state has moved away. Compare the two cases, explain why their estimators could not see the error, and propose one measurement or maneuver for each that would let them.

14.5 Discussion · 7 min

After the gate lockout of Chapter 13, a reviewer proposes to turn the chi-square gate off, since "the gate threw away 132 good fixes". Using the chapter and Section 14.4.1, evaluate the proposal and state what you would change instead.

14.6 Discussion · 9 min

A manager reads that the spray dryer's MPC shipped 0.444 t off specification against 0.570 t for the loops and asks for MPC on every dryer. Using Section 14.2, Chapter 10 and Chapter 7, explain where the improvement came from and what comparison you would run before deciding.