INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 6 Multivariable and Plantwide Control

CHAPTER 12

Summary and Outlook

Pages 278–291

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Multivariable and Plantwide Control*, John D. Hedengren, Ed., vol. 6 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 278–291.

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/v06-multivariable-plantwide/

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

CHAPTER 12

Summary and Outlook



Open gaps 1 one-way coupling 2 decoupler cost 3 switching structures 4 downstream variables 5 weak directions 6 perishable chains

Notice that the gap of switching structures joins the reverse osmosis skid, the recycle plant and the Tennessee Eastman plant, the three units that held a limit only by changing what they controlled. Illustration, not to scale.

IN PLAIN TERMS **What Ten Plants Taught About Structure**

The main idea. Each chapter of this book ran a plant with several handles and several results to hold, from a seawater plant to a dairy chain. Before any tuning, someone had to decide which handle answers to which result, and where the plant's output is set.

The challenge. A wrong choice failed however it was tuned. A common test of the pairs missed a pull that runs one way only, and a fix for one result could cost another far down the line.

Why it matters. Testing each choice against a real upset, judged where the next unit feels it, told the designs that worked from those that failed.

Ten chapters put the structure of a control system on trial before any loop was tuned. Each plant had more than one actuator and more than one variable worth holding, and in each a wrong arrangement failed in a way no tuning could repair. A pairing drove a train to a valve stop, a fixed fresh feed flooded a column, and a full inverse lost the orbit to one reversed corrector. This chapter sets the results side by side with the measure that predicted each, draws a short guide for choosing a structure, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

12.1 What the Chapters Showed

Table 1.2 of Chapter 1 listed the structural question of each chapter and the measure it reports, without numbers. Tables 12.1 and 12.2 add the results. The first holds the unit-level chapters, where several actuators act on one unit, one line or one ring. The second holds the plant-level chapters, where the question is where production is set and where the upsets go. Each row names the difficulty, an arrangement that worked, one that failed, and the measure with the measured result, cited to the table, equation or section of the chapter that reports it. Figure 12.1 sets one like pair of numbers from each row side by side.

The columns of the two tables make five points.

The relative gain rejected the pairings that fail outright. Each pairing whose relative gain was negative failed whatever the tuning. The swapped pairing of Chapter 3, at $\Lambda_{12} = -0.517$, made 15.5 times the integrated error of the natural pairing and 13.6 t of off-specification paper against 2.1 t. The reversed pairing of Chapter 2 drove the valve to its stop. The array did not report the interaction that remained. The reverse osmosis gains are triangular, so $\lambda_{11} = 1.00005$ while the pump moves the recovery by 0.703 %/% of speed, and the cabin of Chapter 10 couples its two gas loops through

Table 12.1 Pairings, Decouplers and Inverses, Chapters 2 to 6

Chapter	Difficulty	What worked	What failed	Measure and result
2, reverse osmosis train	The valve's gain on permeate flow is –0.001 in $m^3/h/\%$ at design and changes sign either side. The pump moves recovery strongly.	Pump on production, valve on recovery, one static feedforward from the pump to the valve at the gain ratio 0.914.	Production on the valve. The train ends at 82.5 m^3/h against a 500 m^3/h setpoint (Section 2.6).	$\lambda_{11} = 1.00005$ (Equation (2.14)). Peak recovery error 0.937 points, 0.213 with the feedforward, 24.35 swapped (Table 2.3).
3, paper machine	Weight and moisture both respond to stock and steam. The scanning gauge adds a scan to the dead time, 21.9 s for moisture.	Stock on weight, steam on moisture. The stock-to-steam half of a static decoupler. A 10 s scan.	The swapped pairing, still walking 22 minutes after a dilution, at 72.8 g/m^2 and 13.12 % moisture.	$\Lambda_{11} = 1.517$ (Equation (3.25)). Moisture excursion 1.80 to 0.23 points with the decoupler, off-specification paper 2.38 to 4.54 t (Table 3.5).
4, milk evaporator	Solids are read at the end of the vapor chain, 380 s after a steam step (Section 4.4).	Steam on solids, a ratio feedforward compensated for feed solids, averaging level loops.	Feed on solids, which cuts throughput. Tight level loops, which pass boiling noise to the dryer.	One point of solids costs 0.12 t/h of feed or 0.030 t/h of steam. Shift: 0.64 t off and 155.5 t fed as shipped, none off and 160.0 t suggested (Table 4.7).
5, coating line	Line speed is the throughput handle. A speed change disturbs both spans and the coat weight, which a gauge reads 19 m downstream.	Pump feedforward from the speed and inertia compensation. For a late batch change, a lead on the pump and a 25 s trim.	A speed ramp with the pump flow held. A trim that waits for the gauge.	Scrap after a speed change 208.90, 22.30 with the feedforward and 12.80 m with both (Table 5.6). Case 9: 63.00 loops, 21.05 MPC, 20.55 m with the lead (Table 5.8).
6, storage ring	Singular values fall from 379 to 0.94, a condition number of 405.2, with a cliff at mode 48 (Section 6.4).	The inverse truncated at the cliff or regularized. A fast set, a monitor residual test and a feedforward.	The full inverse. It moves the magnets most and fails with one corrector wired backward.	Jitter 0.198 against 0.0164 µrad, orbit 0.09 against 0.21 µm. Reversed corrector: 132.6 to 136.0 µm full, 0.38 at 48 modes (Table 6.4).

one sum with a relative gain of one. A triangular plant needs its off-diagonal gains read one by one, because the array multiplies them together.

A decoupler is priced by its other loop. The static decouplers of Chapter 2 and Chapter 3 cut the disturbance they targeted by a factor of four or more, from 0.937 to 0.213 points of recovery and from 1.80 to 0.23 points of moisture. On the paper machine the stock-to-steam half carried almost all of the benefit, and the full decoupler raised the off-specification paper from 2.38 to 4.54 t, because the weight loop now carried a steam-driven bias while the production loop slowed the machine. On the reverse osmosis train the decoupler gain had a minimum, and a gain of 1.6 left a peak of 0.864 points, close to the 0.937 of no decoupling at all (Table 2.4). Section 1.4.6 found the same trade on the running example.

Where production is set decides where the upset goes. With the fresh feed fixed, 10 % more feed raised the recycle of Chapter 7 by 62.3 % in 12 h and left a vapor margin of 3.1 %. Setting production at the reactor holdup gave about the same production with the recycle 4.2 % lower, at the price of 14 % more holdup for each 10 % of production. The evaporator could hold its solids on steam or on feed at nearly the same speed, but the feed pairing spent throughput. On the coating line the speed was the throughput handle, and the scrap of a speed change fell from 208.90 to 12.80 m once the

Table 12.2 Production Handles, Inventories and Chains, Chapters 7 to 11

Chapter	Difficulty	What worked	What failed	Measure and result
7, reactor and column with recycle	The recycle magnifies a change of fresh feed. The column floods 10.7 % above design production (Section 7.4).	CS2, production set by the reactor holdup. In the challenge, CS2 with the level at 75 %.	CS1 with the fresh feed on flow. CS3 with a fixed recycle, which floods at 4.47 h.	Snowball elasticity 4.61 (Table 7.3). Recycle up 62.3 % on CS1, down 4.2 % on CS2 (Table 7.6). Challenge fails on CS1, passes on CS2 (Table 7.7).
8, Tennessee Eastman plant	Unstable open loop, a trip at 3.59 h. Gas recycle with an inert and a purge.	Inventory loops first, composition loops above them, pressure at 2800 kPa g. An override and a production cut on the A feed loss.	Inventory loops alone on a feed change. The decentralized structure as shipped on the A feed loss.	Trips at 5.84 and 8.40 h, none with composition loops or with the override and cut (Table 8.5). Challenge cost 215.7 at 2705 and 187.3 $/h at 2800 kPa g (Table 8.7).
9, peptide purification	The chromatography cut is judged at the filter. Late variant in the pool shrinks the crystals and slows filtration.	A downstream-aware cut that stops on the predicted late variant, or a fixed window at the times it realizes.	A cut that maximizes recovery at the pool's purity limit. The UV threshold rule.	Late variant 0.430 against 0.698 %, filter 46.8 against 62.0 min, 59.9 g acceptable against none. Purity-maximizing cut 43.4 g (Table 9.6).
10, sealed cabin	Oxygen and nitrogen loops share one total pressure, and the machines can only add gas. Eclipses limit power.	A power-aware schedule that keeps electrolysis out of the dark. A nitrogen balance that flags a leak.	The week without the balance, or on the adaptive schedule.	Relative gain one. 1 kPa more oxygen shuts the nitrogen valve 2.52 h later (Section 10.4). Week passed at 11 of 11 seeds, at 0 of 11 by either failure (Table 10.5).
11, dairy supply chain	Each stage orders from the next one up. Orders lag, milk spoils and the plant must take its farms' milk.	Counting the supply line ($\beta = 1$) and sharing the supermarket's sales. The MPC with shared sales.	The shipped human rule, $\beta = 0.34$.	Peak amplitude ratio 1.89 (Section 11.5). Plant variance ratio 12.70, 7.61 at $\beta = 1$ (Table 11.4). Challenge 898.2 k$ and 44 h of outage, 791.3 and 6 suggested (Table 11.5).

pump and the winders were driven from it. In Chapter 8 the inventory loops held the plant until a feed changed, and only the composition loops above them carried it through.

Weak directions carry noise and model error. The full inverse of Chapter 6 made the slow correctors jitter 12 times as much as a 48-mode truncation, for an orbit of 0.09 instead of 0.21 µm. The disturbance decided what the weak modes were worth. The girder drift put 86.0 % of its power in modes 45 to 48 and 7.3 % in the four weakest. With one corrector wired backward the full inverse lost the orbit at every bandwidth tried, while the truncation at 48 modes held 0.38 µm. A large condition number warns about the noise and about the model at once.

A chain is judged at its far end. Three chapters judged a decision by its effect downstream. The shipped ordering rule of Chapter 11 has a peak amplitude ratio of 1.89, at a wave of 15.9 days, so a swing near that period grows at each stage, and counting the supply line cut the plant's variance ratio from 12.70 to 7.61. Tight level loops in the evaporator passed the boiling noise on as a product flow with a standard deviation of 0.60 t/h, against 0.023 t/h for averaging loops (Section 4.6). The cut of Chapter 9 that maximized recovery at the pool's purity limit lost the batch at the filter. In each case the measure that predicted the result was computed at the end of the chain, not at the decision.

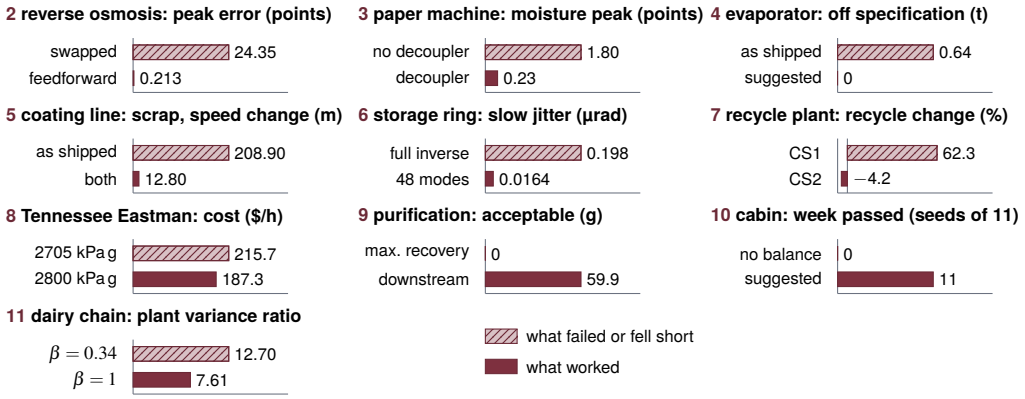


Figure 12.1 One like pair of numbers from each row, as the "Measure and result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. Lower is better on every panel except the purification and the cabin. Notice that no better number came from a faster loop alone: each came with a different pairing, feedforward, decoupler, production handle, inverse, cut rule, balance or ordering rule, or, on the Tennessee Eastman plant, a higher pressure setpoint.

12.2 Choosing a Structure

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

1. *How many degrees of freedom are there, and where is production set?* Count the valves that affect a steady state (Section 1.4.2), place the throughput handle, and point every level loop away from it (Section 1.4.3). Close the inventory loops first. The plant of Chapter 8 trips within hours without them.

2. *Does a recycle magnify a change of production?* Compute the snowball elasticity and the capacity margin (Section 1.4.4). Where the margin is small, set production at a flow inside the loop or at a holdup, as CS2 does in Chapter 7.

3. *What does the next unit respond to?* Hold the variable that the downstream unit feels, as the cut of Chapter 9 holds the predicted late variant rather than the pool purity.

4. *What is the relative gain at the operating point, and does any gain change sign across the range?* Use a central difference at the design point. Reject a negative relative gain. Where the array reads one, check whether the plant is triangular. The reverse osmosis valve gain runs from +0.938 at 46 % to −1.045 at 54 % opening (Table 2.2).

5. *Which half of a decoupler carries the benefit, and what does the other loop pay?* Run each half alone and sweep the gain, as Chapter 3 and Chapter 2 do.

6. *Is the gain matrix ill-conditioned?* Read the singular values and project the disturbance on the modes before choosing how many to correct (Chapter 6).

7. *Does a chain of inventories amplify?* Compute the amplitude ratio of each level loop or ordering rule (Section 1.4.8). Tune levels as buffers where a downstream unit feels the flow, and count what is in transit (Chapter 11).

A predictive controller paid for itself where it carried a delay or a constraint that the loops did not see, and even there simple elements often matched it. On the coating line it scrapped 21.05 m against

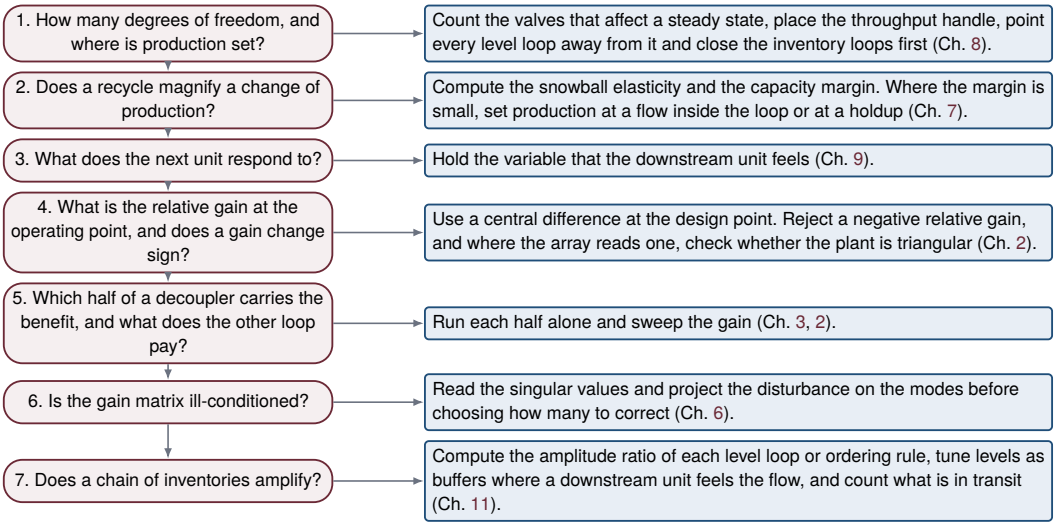


Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the production handle and the recycle come before any pairing is chosen: a fixed fresh feed flooded the recycle plant's column, and the Tennessee Eastman plant tripped without its inventory loops.

63.00 m for the loops when a thinner batch and a faster order arrived together, and loops with a lead on the pump and a 25 s trim scrapped 20.55 m. On the reverse osmosis train it held the feed pressure at 69.81 bar where the loops spent 108.6 min above the 70 bar rating, and a pressure override on the recovery target held 69.87 bar (Table 2.8). On the paper machine's grade change it made 2.68 t of off-specification paper against 4.32 t for the shipped decoupler and lost to the decoupler retuned by the same rule, at 1.51 t (Table 3.7), chiefly because its model holds the speed that the production loop moves. In the evaporator it cut the integrated error of a throughput step from 54.92 to 20.28 % min and put 0.466 t off specification on a shift the loops passed (Table 4.8). Volume 9 treats that choice in full. Figure 12.2 draws the seven questions in order.

12.3 Where Practice Is Heading

Structure before algorithm. Plantwide control asks which variables to control, which to measure, which inputs to manipulate and which links to make between them, questions that most control theory assumes are answered at the outset [1]. A recent review argues that decomposing a control system into simple elements such as cascade, ratio, feedforward, decoupling and selectors is a powerful way to design control for complex processes with limited information. It challenges the view that model predictive control can solve every control problem and calls for more teaching of these elements [2]. The chapters of this volume point the same way on their own cases. The retuned decoupler of the paper machine beat the predictive controller, and the pressure override of the reverse osmosis train matched it.

Decomposition at plant scale. Large control and optimization problems can be decomposed by analyzing their network for its latent block structure, which then determines the decomposition used for distributed control [3]. Distributed predictive controllers built on recurrent neural network models have been shown, on a simulated chemical process network, to compute more efficiently than a centralized controller while performing very close to it [4]. In the editor's view, partitioning a

plant for such controllers is the structure question of this volume asked again, and the measures of Chapter 1 remain a way to check a proposed partition before it is trusted.

> ⧖ HISTORY **2004 to 2023: From a Design Procedure Back to Simple Elements**
>
> The history boxes of Chapter 1 end with a review of 2000. A procedure published in 2004 split the design of a plant's control system into a top-down analysis and a bottom-up design of the control layers, and its author wrote that more work was needed to understand how to decompose and coordinate those layers [5]. In 2011 an industrial and an academic author found that much of the research had still not gained enough acceptance to influence industry profoundly, and stressed the need for a formalized yet simple approach [6]. In 2023 a review proposed that far more academic research go into the theory of the simple control elements that industry already uses [2]. In the editor's reading, the ten chapters of this volume stand at that point: each structure that worked was built from such elements and checked by a run.

Benchmarks that repeat. The Tennessee Eastman plant has served as the benchmark in comparison studies of data-driven fault diagnosis and process monitoring [7]. A revision of its code found that changes in simulation software since the 1990s broke the repeatability of simulation results, and presented a modified code [8]. Chapter 8 met the same sensitivity in its own port, where the open-loop trip came at 3.47, 3.59 or 4.06 h depending on how the constants were rounded (Section 8.4). In the editor's view, a plantwide benchmark is useful only with its code, its numerical precision and its scoring published together.

Information along supply chains. When each company in a chain orders from the one upstream, the variance of orders can exceed that of sales and grows moving upstream, and actions to mitigate the distortion have been discussed since 1997 [9]. The dairy chain of Chapter 11 measured the same growth in milk, and sharing the supermarket's sales took the plant's variance ratio under order-up-to from 11.36 to 2.73 at a seven-day window (Section 11.5).

Feedback by mode on large arrays. Orbit feedback at light sources corrects a large array through the modes of its response matrix. A study of one feedback system reports that Tikhonov regularization, which lowers the gain of the weak modes, is used in the feedback systems of many light source facilities [10]. A controller with its own dynamics for each mode was implemented on a booster ring run with stored beam as a test bed, because the two extra matrix products per update are usually too heavy for a storage ring with more than a hundred actuators [11]. For beams of the order of 4 µm, a design study began to look beyond the classic PID regulator to model-based methods and stated that a reliable model is required [12]. The truncated and regularized inverses of Chapter 6 belong to the first of these steps.

Measurements that see what the next unit feels. Process analytical technology has become a pillar of process control for numerous small-molecule drugs, while the close chemical similarity of protein contaminants makes suitable methods hard to find. In-line soft sensors that combine spectroscopy with chemometrics have reached the needed detection limits for some contaminants at laboratory scale, and the review that reports them names closed-loop control as a next step [13]. On paper machines, a laser ultrasonic sensor of sheet stiffness was tested at three industrial sites and placed at the demonstration stage [14]. The soft sensor of Chapter 9 predicts the late variant at the column, while the filter decides the batch. In the editor's view, a sensor of what the next unit feels lets a structure hold that variable directly.

Treatment trains as control networks. A roadmap for water treatment observes that most systems need operators to adapt to changing water quality and component degradation, and that the data of their many sensors are used mainly to detect failures. It calls for sensing and control networks with

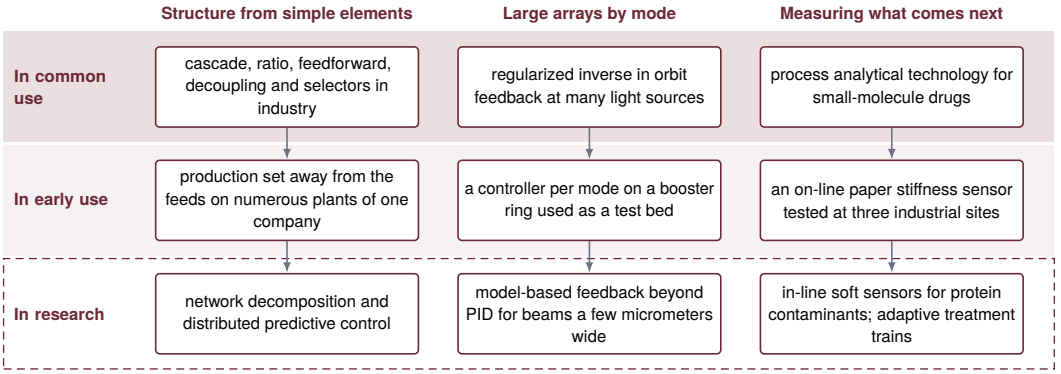


Figure 12.3 Developments in three threads. The top lane holds what the cited source calls commonly used or used at many facilities, the middle lane what it reports in service at some plants or on a test bed, and the bottom lane what it reports from simulation, at laboratory scale or as a need, in the editor's reading. Sources by thread: structure from simple elements [2–4, 6], large arrays by mode [10–12] and measuring what comes next [13–15]. Notice that each thread ends in a model that must itself be checked on the plant before a structure can rely on it.

models of the interdependencies between unit processes, and notes that the savings of predictive control are hard to quantify without actual facilities [15]. The reverse osmosis train of Chapter 2 models one such dependency, from the pump to the recovery, and leaves the pretreatment out. In the editor's view, the pairing and override questions of this volume return at the scale of a whole treatment train. Figure 12.3 sets these developments in three threads.

Agents that propose structures. A 2026 survey of language models in process systems engineering found genuine promise for tasks built on natural language, such as querying documentation, and found that applications needing real-time execution, constraint satisfaction or formal safety guarantees remain challenging [16]. An agent built to turn industrial specifications and process diagrams into optimization models for process control was evaluated on 50 scenarios, and its authors state that the evaluation assessed the formulation, not closed-loop industrial performance [17]. In the editor's view, agents will soon draft pairings, decouplers and inventory structures in quantity. The checks of this volume, a relative gain at the operating point, a snowball gain, a run of each decoupler half and a run on a disturbance, are how an engineer will accept or reject them.

12.4 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 [18]. 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 [19]. 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.

12.4.1 How This Volume Was Made

Each of the ten apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace numbers to runs, and a referee report was written on the paper. For at least three papers the report is a self-review by the agent that wrote it, and for the reverse osmosis paper it covers only material added in a later round. Defects found along the way went into a log for each app. For this book one agent per chapter reran every number that an exercise uses. The editor set the scope of each round and the rules on what each agent could change, and statements that runs contradicted were left for the editor to decide. The record holds the following errors and the checks that caught them.

The defect logs. The ten logs hold 70 entries with a severity. Of these, 18 are marked as wrong behavior of an app (a wrong output, a blocker, a functional fault or a wrong credit), 47 as documentation, four as cosmetic and one as minor. The reverse osmosis log adds 15 numbered entries without a severity. Seven cases fit the topic of this volume.

A relative gain from the wrong derivative. The reverse osmosis app's briefing computed $\lambda_{11} = 0.99$ from a valve gain of -0.13, a one-sided difference taken away from the maximum. With the briefing's own gains the arithmetic gives 1.007, and the central difference at the design point gives a valve gain of -0.001 and $\lambda_{11} = 1.00005$. The same log found the pump-to-recovery gain quoted as 0.73 in five places where the model gives 0.7029, which moves the decoupler ratio from 0.95 to 0.914. Both were caught by recomputing each quoted number from the shipped model.

A pairing switch that was not bumpless. Choosing the swapped pairing before Run made the first scan command the concentrate valve from 50 to 88 %, the pump speed its new controller had been initialized with, and the permeate fell from 450 to 27.8 m^3/h. The controllers kept their state when the pairing changed. The defect showed only when the pairing was changed in automatic, which the shipped scenarios avoided by starting in manual, and it was reported with the scripts that reproduce it. Each loop is now seeded on the actuator it drives, no shipped scenario run moved, and the swapped row of the paper was measured again.

Two tunings compared as two handles. The first draft of the evaporator paper repeated the app's explanation that the feed-paired loop is slower because the feed must pass the sumps. The self-review identified both handles open loop, found loop gains of 1.32 and 1.0 for the shipped tunings, and reran the feed pairing at the steam loop's loop gain. It then reached the setpoint band only 0.2 minutes behind, 0.3 on the current engine. The speed difference was the tuning, and the lasting difference is the cost in milk.

A production loop that read the plant's truth. The paper machine's production loop computed its tonnage from the true dry fiber flow times a hidden retention, which no instrument reads. Two identical machines, one given a lower retention, differed by 1.36 t/h in measured tonnage before the gauge had posted a pass of the changed sheet. The read acted as a hidden feedforward through the speed. Once the loop measured the scanner's weight times the speed, every automatic result with the production loop in service moved. The decoupled grade move made 1.07 t of off-specification paper before the fix and 4.54 t after, and decoupling became the trade of Section 12.1. A second log, the coating line's, records a trim started from the true viscosity, where no shipped run moved. Both logs document the same probe, which changes a hidden value and asks whether the controller moves while the measurements do not. This is the check the shared section calls reading code for what it reads.

A rule described as its labels intend. The purification app labels a UV threshold "Stop below", and its texts describe a cut that collects until the detector falls below it. The referee pass found that the opening threshold is tested at every scan, so the stop threshold never acts. Settings of 0.10, 0.35 and 0.59 gave identical runs. A final validation over 31 seeds at two step sizes then found that the suggested challenge settings pass at 18 of 31 seeds at the default step. The paper states both, and the fix, which moves every run that uses the rule, awaits the editor.

A claim about every structure. The first draft of the recycle paper said that PI loops hold three levels and two purities under every structure. On CS3 the reflux holds the drum and the recycle purity floats. The self-review read the claim against the paper's own case table and narrowed it.

An offset that looked like a plant property. The agent writing Chapter 1 first clamped the controller outputs of the decoupled example rather than the applied actuators, and the loops held a steady offset that looked like a property of the plant. The catalyst-loss run exposed it, and now only the actuators are clamped.

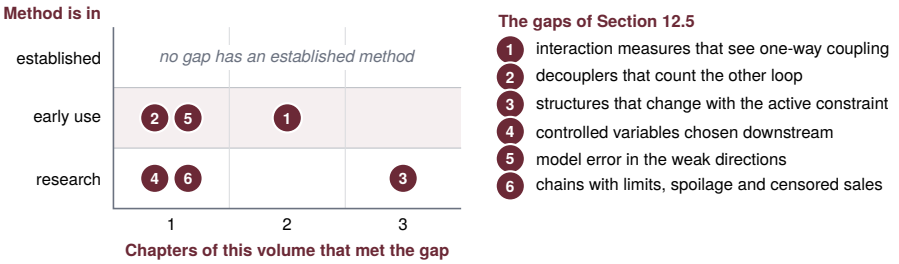


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 2 and 10), one for gap 2 (3), three for gap 3 (7, 8 and 2), one for gap 4 (9), one for gap 5 (6) and one for gap 6 (11). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap met by the most chapters, structures that switch with the active constraint, is still in research.

The book round. The ten chapter records list 162 rows of numeric verification, each a stated value beside its reproduction. Every chapter number rerun reproduced except one already recorded in the paper's notes, a worst posted moisture of 1.25 points on the paper machine's challenge that now runs at 1.244. Two statements contradicted by runs were reported, and both papers and chapters have since been corrected: the recycle product first crosses 97.5 mol % at 7.62 h while the column still floods, and the cabin's case 1 settles at 101.37 %. A printed saturation pressure of 1.002 kPa, which reads 1.0007 kPa in the model, was corrected with them, and the milk app now passes 54 checks where its chapter says 42.

Every error above was caught by a check that could fail: a derivative recomputed at the operating point, a structure switched in automatic, two handles compared at equal loop gain, a hidden value changed while the measurements stayed fixed, a threshold varied, a claim read against a table, a disturbance run as well as a setpoint step. None was caught by rereading the sentence that contained it. Four working rules follow from this record. Compute every structural measure from the model at the operating point and say how. Compare structures at equal loop gain and on a disturbance. Read what each controller reads. Let a person decide what changes.

12.5 Gaps in Multivariable and Plantwide Control

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow them in Section 12.5.1. 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 12.4 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. *Interaction measures that see one-way coupling.* The reverse osmosis train and the cabin both returned a relative gain of one while one loop disturbed the other strongly. A review of plantwide control lists the closed-loop disturbance gain and the performance relative gain array beside the relative gain as tools to analyze interactions in decentralized control [1], and a later review notes that single loops may perform poorly on some interactive processes, where decoupling or multivariable control should be considered [2]. The gap is hard because a measure that multiplies the two off-diagonal gains cannot tell one strong path from none. A routine measure, reported beside the array, is needed that gives the size and direction of one-way coupling and flags a gain that changes sign inside the operating window, as the concentrate valve's gain does.

2. *Decouplers that count the other loop.* The static decoupler of Chapter 3 put the right steam on at the wrong time, because its two paths take 44 s and 84 s, and it doubled the off-specification paper of the loop it did not help. A review of simple control elements suggests one-way decoupling, with

one element set to zero, to avoid instability and sensitivity to model error, and partial decoupling of only the important outputs [2]. The gap is hard because the cost falls on the loop the decoupler was not built for. Design rules that choose which half to build, add dynamics behind a late gauge and price the cost to the other loop are thin.

3. *Structures that change with the active constraint.* The recycle plant floods once production passes its vapor wall, and the Tennessee Eastman plant on its A feed loss needed an override and a production cut to survive. At maximum production sixteen settings all tripped within a day. The reverse osmosis train of Chapter 2 held its pressure rating through an override that gave up recovery at the limit. The Open challenge box after this list takes the gap further.

4. *Controlled variables chosen by their downstream effect.* The soft sensor of Chapter 9 predicts about 0.78 % late variant whatever the crude (Table 9.7), so its downstream-aware rule does no better than a fixed window. A review of plantwide control warns that outputs chosen by a controllability analysis may be those easy to control rather than those important to control [1], and a guideline on drug substance manufacture asks that the linkages between process steps be evaluated so that the cumulative generation and removal of impurities is controlled [20]. The gap is hard because the variable that matters is measured only after the next unit has acted. Choosing such variables needs measurements that see the upstream composition and a validated model of the downstream unit.

5. *Model error in the weak directions.* The full inverse of Chapter 6 broke on one reversed corrector, and a response measured through the plant held 0.05 µm. A study of one fast orbit feedback system found that the weak modes, though dominated by measurement errors, still carry information, and that measuring the closed-loop gain needs an excitation larger than the beam's noise, which itself disturbed the gain [10]. The gap is hard because the directions that most need checking respond least to a test the plant can tolerate. A way to measure and validate the weak directions of a large response during operation, without disturbing it, is still needed.

6. *Chains with limits, spoilage and censored sales.* No ordering rule of Chapter 11 is string stable, and in a perishable chain with capacity limits the variance floor is a reference, not a bound. A review of fresh dairy quality finds that most models estimate shelf life from the product's initial condition and seldom state their uncertainty [21], and a study of grocery stores reports that an error in the stock record cannot be observed when it occurs and is found only by a deliberate count [22]. The gap is hard because the sales a rule sees are clipped by the stock it already lacks. Design rules for chains that pour out milk and clip orders are thinner than for linear chains.

> **OPEN CHALLENGE** **Structures That Switch With the Active Constraint**
>
> Gap 3 runs through the recycle plant, the Tennessee Eastman plant and the reverse osmosis train, which held a limit only by changing what they controlled. A perspective on plantwide control notes that there is usually no effort during design to build a control system that can run at maximum throughput [6]. A review of simple control elements finds that switching between active constraints is usually done with selectors and split range control, and that the analysis of such switched systems is mathematically very difficult [2]. In the editor's view the gap is hard because each switch changes the structure whose stability was checked. Progress would be a routine test that proves a structure through every switch it may make.

12.5.1 Gaps Common to Agent-Era Practice

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

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability

and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [23]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [19]. 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 [24]. 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 [25]. 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 [26]. 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

- The arrangement of a control system, chosen before any tuning, set what tuning could reach: a wrong pairing failed whatever the gains.
- A relative gain of one did not prove the loops independent. A triangular plant needed its off-diagonal gains read one by one.
- A decoupler was priced by its other loop, and where production was set decided where the upset went.
- A chain was judged at its far end, and every error in the record was caught by a check that could fail, never by rereading.

References

[1] T. Larsson and S. Skogestad, "Plantwide control: A review and a new design procedure," *Modeling, Identification and Control*, vol. 21, no. 4, pp. 209–240, 2000.

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

[3] W. Tang, A. Allman, I. Mitrai, and P. Daoutidis, "Resolving large-scale control and optimization through network structure analysis and decomposition: A tutorial review," in *Proc. American Control Conference (ACC)*, 2023, pp. 3113–3129.

[4] S. Chen, Z. Wu, D. Rincon, and P. D. Christofides, "Machine learning-based distributed model predictive control of nonlinear processes," *AIChE Journal*, vol. 66, no. 11, p. e17013, 2020.

[5] S. Skogestad, "Control structure design for complete chemical plants," *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[6] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[7] S. Yin, S. X. Ding, A. Haghani, H. Hao, and P. Zhang, "A comparison study of basic data-driven fault diagnosis and process monitoring methods on the benchmark Tennessee Eastman process," *Journal of Process Control*, vol. 22, no. 9, pp. 1567–1581, 2012.

[8] A. Bathelt, N. L. Ricker, and M. Jelali, "Revision of the Tennessee Eastman process model," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 309–314, 2015.

[9] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[10] S. Kongtawong, Y. Tian, L. Yu, X. Yang, G. Wang, K. Ha, and T. Shaftan, "Numerical simulation of NSLS-II fast orbit feedback system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 997, p. 165175, 2021.

[11] S. Gayadeen, S. R. Duncan, and M. T. Heron, "Fast orbit feedback control in mode space," in *Proceedings of ICALEPCS2013*, San Francisco, CA, USA, 2013, pp. 1082–1085, paper THMIB07.

[12] P. Kallakuri, D. Paskvan, J. Carwardine, N. Sereno, and N. Arnold, "Modeling the fast orbit feedback control system for APS upgrade," in *Proceedings of IBIC2017*. Grand Rapids, MI, USA: JACoW, 2018, pp. 196–198.

[13] L. Rolinger, M. Rüdt, and J. Hubbuch, "A critical review of recent trends, and a future perspective of optical spectroscopy as PAT in biopharmaceutical downstream processing," *Analytical and Bioanalytical Chemistry*, vol. 412, no. 9, pp. 2047–2064, 2020.

[14] L. Kong, A. Hasanbeigi, and L. Price, "Emerging energy-efficiency and greenhouse gas mitigation technologies for the pulp and paper industry," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep., Nov. 2012.

[15] D. Sedlak, M. Mauter, J. Macknick, J. Stokes-Draut, P. Fiske, D. Agarwal, T. Borch, R. Breckenridge, T. Cath, S. Chellam, A. Childress, D. Dionysiou, D. Giammar, E. Hoek, S. Jiang, L. Katz, J. Kim, R. Kostecki, J. McCutcheon, Y. Polsky, Z. Stoll, and P. Xu, "National alliance for water innovation (NAWI) master technology roadmap," U.S. Department of Energy, Report DOE/GO-102021-5617, 2021.

[16] B. Gopaluni, V. Kotamraju, and S. Bhushan, "Large language models in process systems engineering: Opportunities, architectures, and industrial deployment challenges," arXiv preprint arXiv:2606.11589, 2026.

[17] Y. Ye, X. He, D. Boshoff, W. Kuo, and L. Li, "little m: An AI agent for industrial process optimization," arXiv preprint arXiv:2609.16680, 2026.

[18] 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.

[19] 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.

[20] International Conference on Harmonisation, "Development and manufacture of drug substances (chemical entities and biotechnological/biological entities) Q11," ICH harmonised tripartite guideline, Step 4 version, https://database.ich.org, May 2012.

[21] L. Lima, A. I. Pereira, M. Pintado, and M. Carocho, "Fresh dairy quality in the era of smart food systems: Deterioration pathways, analytical technologies, and predictive modeling," *Foods*, vol. 15, no. 19, p. 3439, 2026.

[22] Y. Rekik, R. Oliva, A. A. Syntetos, and C. H. Glock, "Inventory record inaccuracy in grocery retailing: Impact of promotions and product perishability, and targeted effect of audits," *Journal of Business Logistics*, vol. 47, no. 3, p. e70079, 2026.

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

[24] 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.

[25] 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.

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

12.6 Exercises

12.1 Discussion · 8 min

An agent proposes a pairing for a new two-by-two plant and reports a relative gain of 1.01, "so the loops will not interact and need no decoupler". Using Chapter 2 and Chapter 10, explain why that conclusion does not follow, and list what you would check before accepting the pairing.

12.2 Discussion · 8 min

On the paper machine of Chapter 3 an operator reports that the static decoupler "made the eight-gram grade move worse", and the control engineer reports that it "cut the moisture excursion by a factor of eight". Using Table 3.5, show that both are right. Then recommend a configuration and say what you would measure to settle the argument.

12.3 Discussion · 7 min

The owner of the reactor and column plant of Chapter 7 wants 15 % more product and proposes to raise the fresh feed on the shipped structure, CS1. Using the chapter's results, explain what will happen. Then propose where production should be set and the check to make before any change.

12.4 Discussion · 8 min

An agent retuning the orbit feedback of Chapter 6 proposes the full inverse because "it gives the smallest orbit, 0.09 against 0.21 µm". Evaluate the proposal with the chapter's results, and say what you would ask the agent to run before you accept any inverse.

12.5 Discussion · 10 min

Section 12.4.1 lists seven errors made by the agents that built this volume's apps, papers and Chapter 1. For each, name the check that caught it. Then write a procedure, in five to eight steps, for an agent that will propose a change of control structure on a running plant (a new pairing, a decoupler or a moved throughput handle), stating what the agent may read, what it may write and what a person must approve.

12.6 Discussion · 7 min

A site standard asks for every level loop to be tuned tight "so that inventories stay on target". Using Chapter 4, Chapter 11 and Section 1.4.8, argue when the standard is wrong, and what a supply chain adds to the argument.