INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 13 Control in Chemical Plants

CHAPTER 12

Summary and Outlook

Pages 282–295

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Chemical Plants*, John D. Hedengren, Ed., vol. 13 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 282–295.

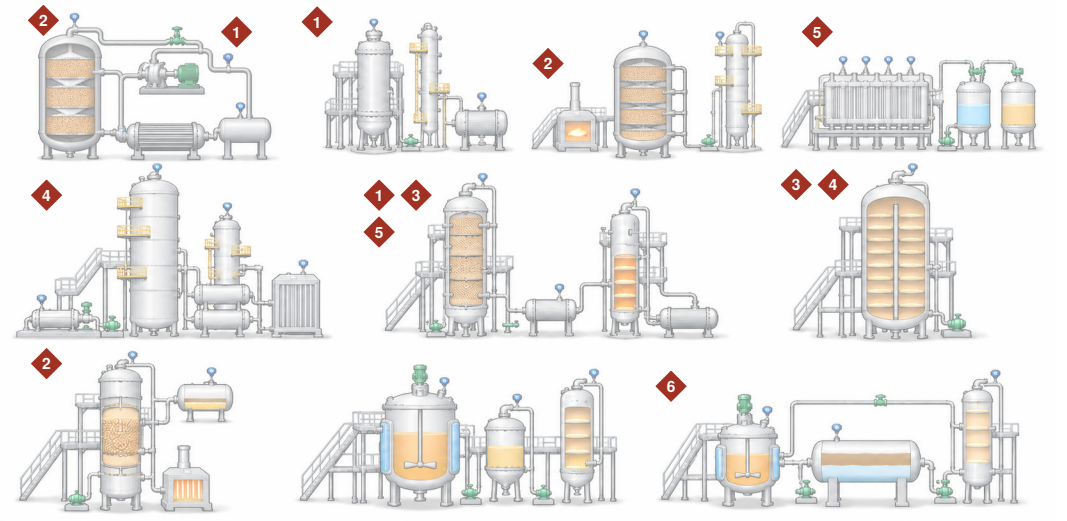
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v13-chemical-plants/

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

CHAPTER 12

Summary and Outlook



Open gaps 1 inventory structure 2 distance to the cliff 3 broken inferences 4 performance unwatched 5 late obligations 6 noisy gradients

Notice that the capture plant carries three gaps, and that gap 2 marks two converters and a column alike. Illustration, not to scale.

IN PLAIN TERMS What Ten Chemical Plants Taught

The main idea. Each chapter of this book ran one chemical plant in a simulator. In most of them, unused material goes round again, a reaction stops partway, and energy is the largest bill.

The challenge. Material that goes round again turns a small error into a large one that grows for hours. The best way to run a plant often sits next to a cliff, such as a reactor that goes out and will not light again by itself.

Why it matters. A plant can hold every reading on target while it wastes energy or spoils product, so what is controlled matters more than how it is tuned.

Ten chapters put ten chemical plants under control. Table 1.2 listed, for each plant, what recycles, what equilibrium limits, where the energy goes and the structure the chapter implements. This chapter adds what worked, what failed and the measured result, and reads the comparison for the question of the volume: how recycle, equilibrium and energy cost shape each plant's control structure. It then gives a short guide for choosing a structure, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

12.1 What the Chapters Showed

Tables 12.1 and 12.2 compare the chapters, five in each. The second column condenses the recycle, equilibrium and energy columns of Table 1.2 into the difficulty each plant posed. The last column gives the measured result with the table, figure or section of the chapter that reports it.

Table 12.1 How Recycle, Equilibrium and Energy Shaped Each Structure, Chapters 2 to 6

Chapter	Difficulty	What worked	What failed	Measured result
2, ammonia loop	Loop gas circulates 5.95 times the makeup. The purge sets the inerts and a 5.1 h time constant. The converter heats its own feed and can go out.	Makeup on pressure, hot spot on quench, and the slow inert and ratio loops in service from the start of the shift.	The fast loops alone: inerts past 16 % at 3.1 h and 21.0 % at the end. A quench opened to 22 % put the converter out.	Challenge passed with a worst pressure error of 0.04 bar and 12.47 % inerts (Section 2.6). Inert IAE 56.6 under prediction against 1843.7 % min supervised (Table 2.7).
3, urea loop	Excess ammonia returns with no exit ($g_R = 1$), so the N/C ratio integrates every feed error. No instrument reads it.	A PI on the N/C inferred from density, with a laboratory bias. A cascade onto the condenser drum through the turndown.	The ratio held on the flows. Wet ammonia drained the loop and the run stopped at 9 h 48 min.	Suggested settings pass 22 of 22 runs, samples 2.965 to 3.033, biuret 0.882 wt% (Table 3.6). The MPC passes with 0.0001 to 0.0054 wt% of margin (Section 3.6).
4, contact acid plant	Four adiabatic beds on an exothermic equilibrium. An ignition floor near 395 °C. No recycle in the gas.	Four bed inlet loops on the bypasses, and the dilution water on ratio with an analyzer trim.	Bypasses held through a turndown to 60 %: bed 1 out at 2 h 17 min, then the others. Either half of the fix alone.	Conversion 99.878 against 95.969 % shipped, stack at most 260 against 48 468 ppmv. The MPC reaches 99.904 % with bed 1 at 409.4 °C (Table 4.6).
5, chlor-alkali cell	No recycle and no equilibrium. The current is the load, power is priced by the hour, and the gas side moves faster than feedback.	Feeds in ratio to the current with trims. Chlorine pressure fed forward to the pressure-difference cascade. A price scheduler.	Feeds in Manual. Feedback alone on the gas side, which trips the rectifier 22 s after the compressor stops.	81.0 t of 86.8 t off specification with the feeds in Manual (Table 5.6). Graded day 72.3 t on specification for $10 908 against $13 853 at constant current (Table 5.7).
6, air separation	One condenser-reboiler couples two columns, 3.11 bar per bar. The compressor runs near its 752 kW rating.	Full cascade, pressure on the guide vanes with a power override. Production given up when the exchanger fouls.	Duty bought back with the pressure setpoint, a 786 kW peak at 5.85 bar. Purity loops blind to a heat leak.	Challenge passed, oxygen within 0.049 %, peak 722 kW (Section 6.6). A heat leak took recovery from 90.8 to 64.9 % while purity rose (Table 6.5).

The columns of the two tables make five points. Figure 12.1 draws one like pair of results from each chapter's row that has one.

A recycle tells the controller which inventory to hold. The three plants with a recycle each needed a loop on the recycled quantity itself. In Chapter 2 a front end 5 % low on ratio settles the loop at 1.80, an amplification of eight, while the pressure never leaves 150.0 bar (Section 2.10). Only the slow inert and ratio loops see it, and they need hours. Put in service at 20 min they finish the challenge at 12.89 % inerts. Waiting two hours gives 13.21 and fails. The urea loop has no exit for its excess ammonia, so a ratio held on two flow meters integrates the meters' error, and only a loop on the inferred N/C held it. In Chapter 7 the return gain is negative, and an inner loop on the lean loading made the outer gain the same at ten minutes and at eight hours. Equation (1.5) predicts all three behaviors from one relation and the value of its return gain.

An equilibrium puts the optimum next to a cliff. Hotter is faster and worse on an exothermic equilibrium. The ammonia converter's own optimum lies near 480 °C and is worth 11 % of production, yet in manual a quench of 15.3 % keeps it lit and 15.4 % puts it out (Section 2.6). The acid plant's best inlet is the coolest one the catalyst can still light, and its MPC held bed 1 within a kelvin of its 410 °C bound. The MTBE column shows the same shape as a fold. A setpoint 1.4 K under tray 12's

Table 12.2 How Recycle, Equilibrium and Energy Shaped Each Structure, Chapters 7 to 11

Chapter	Difficulty	What worked	What failed	Measured result
7, amine capture	Solvent returns in a 70 min lap with $g_R = -3.4$. The reboiler duty is bought at an hourly price.	Capture on circulation over a lean-loading loop on the duty. A supervisor that times regeneration to the price.	A weak inner loop, 3 MW/°C, which leaves the circulation on its stop. A tight level loop that passes the imbalance to the reboiler.	Gain 0.177 at once and 0.040 after 8 h with the duty frozen, 0.177 held with the inner loop (Figure 7.2). Weak loop ends at 93.03 against 94.01 % (Table 7.4).
8, dividing-wall column	No recycle. The spare liquid split moves only the duty, and the minimum moves with the feed.	The four-point structure, a loop on the split, with feedforward from the feed rate.	A fixed split of 0.40. The four-point loop or the feedforward alone. An MPC whose model lacks an input disturbance.	A fixed split costs 17.7 % more steam (Section 8.4). Challenge worst hour 0.3 % above the minimum, 32 of 32 seeds, against 12.4 % shipped (Table 8.6).
9, MTBE column	Etherification on the trays. At 42.5 MW three states convert 98.22, 90.05 and 94.89 %. The duty decides the branch.	The temperature loop on tray 11 at 120 °C, below the low branch's range, so a loop on the low branch cuts the duty.	Tray 12 at 142 °C, 1.4 K under the high branch's peak. On the low branch the same loop raises the duty to 47.76 MW.	Suggested settings keep at least 99.70 mol% MTBE and 97.2 % conversion at seven seeds. Shipped falls to 67.58 mol% (Table 9.7).
10, base oil plant	No recycle and no equilibrium. Yield is traded against a viscosity index the laboratory reports late.	Viscosity on the cut, mid-ranging to the reactor, a bias update and a one-way index cap.	The viscosity loop alone on a blurred cut. Mid-ranging signed the wrong way, which walks the reactor 4.5 °C hotter.	163.5 t downgraded at exactly 6.000 cSt against 5.4 t (Table 10.6). Challenge worst error 0.084 cSt, nothing downgraded, against 63.7 t as loaded (Section 10.5).
11, Williams–Otto RTO	An irreversible reactor whose model lacks a reaction. The optimum lies on the 8 % waste limit.	Modifier adaptation every hour with filter gain 0.2 and a 0.15-point waste back-off.	Two-step, which stops where its own model's slopes vanish. Filter gain 0.3 with no back-off.	Suggested 31 of 31 seeds, 97.47 % of the best. Two-step 86.10 to 92.89 % over 19 settings (Table 11.7).

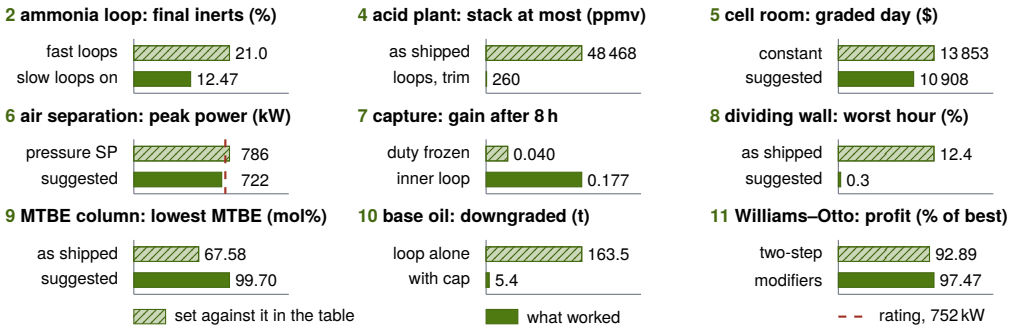


Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" and "What failed" columns of Tables 12.1 and 12.2 print them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 7, 9 and 11. Chapter 3 has no like pair in its row and no panel. Notice that each pair sets one structure or measurement against another.

peak lost the column, and the fix moved the measurement to a tray where a loop on the wrong branch acts in the right direction. In each case the structure carried a margin from the cliff or a measurement chosen so that crossing it could be undone.

Energy is spent by the handle no purity loop watches. A dividing-wall column holding all three purities burned 17.7 % more steam at the wrong split. A fouled air separation exchanger cost eight percent of the oxygen while every controlled variable sat on its setpoint and the compressor power fell. A heat leak raised the oxygen purity while it cut the recovery to 64.9 %. The base oil loop held 6.000 cSt exactly and downgraded 163.5 t. The structures that worked closed a loop on the spare handle (the four-point split loop), or guarded the quality the main loop could not see (the index cap). Purity and viscosity are specifications, and recovery and specific energy are performance.

Timing the energy is worth what the storage allows. The chlor-alkali scheduler made a day's caustic for \$9 008 against \$10 617 at constant current, limited mainly by its tank (Table 5.6). The capture supervisor earned a 0.99 % larger margin over a price day at the same captured tonnage, and on a five-hour peak the lean tank reached its floor after 98 min (Section 7.6). The price decides when, and the tank decides how much.

Prediction paid on slow coupled inventories, and a model needs the right slopes. The predictive controller cut the ammonia inert IAE from 1843.7 to 56.6 % min, kept the oxygen on specification through a day the full cascade spent 724.7 min under it (Table 6.6), and earned \$5 105 more than the capture supervisor on one price day, 70 % of which a supervisor given the same prices recovered (Section 7.10). It failed the dividing-wall challenge with the benzene at 98.29 mol%, because its model lacked an input disturbance, and it passed the urea shift by thousandths. The optimizer of Chapter 11 fitted every measurement and stopped at 92.83 % of the best, because its model's slopes vanished where the plant's did not. Modifier adaptation corrected the slopes and averaged 98.96 % (Section 11.9).

12.2 Choosing a Method

The chapters suggest an order of questions for a chemical plant's control structure. Each names the chapter where its answer decided the result, and Figure 12.2 sets the same questions as a flowchart.

1. *What recycles, and what is its return gain?* A gain near one makes a slow composition that a purge or a slow loop must hold (Chapter 2). A gain of one makes an integrator that needs a loop on the recycled ratio itself, measured or inferred (Chapter 3). A negative gain wants an inner loop on the returning quality (Chapter 7).
2. *Which inventory does the makeup hold?* Pairing the makeup with pressure or with production changed where a catalyst loss appeared: 17 t/d of production under one, 2.3 bar of pressure margin under the other (Section 2.10).
3. *Where is the cliff?* Find the ignition floor, the extinction point or the fold before choosing a setpoint, keep a stated margin from it, and measure where the loop's sign holds on every branch (Chapter 4, Chapter 9).
4. *Is the quality known late?* Infer it from fast measurements, correct the inference with each laboratory result (Section 1.4.6), and protect a second quality with a cap or a back-off (Chapter 3, Chapter 10).
5. *Is there a handle that moves only energy?* Close a loop on it or price it, and trend recovery and specific energy beside purity (Chapter 8, Chapter 6).

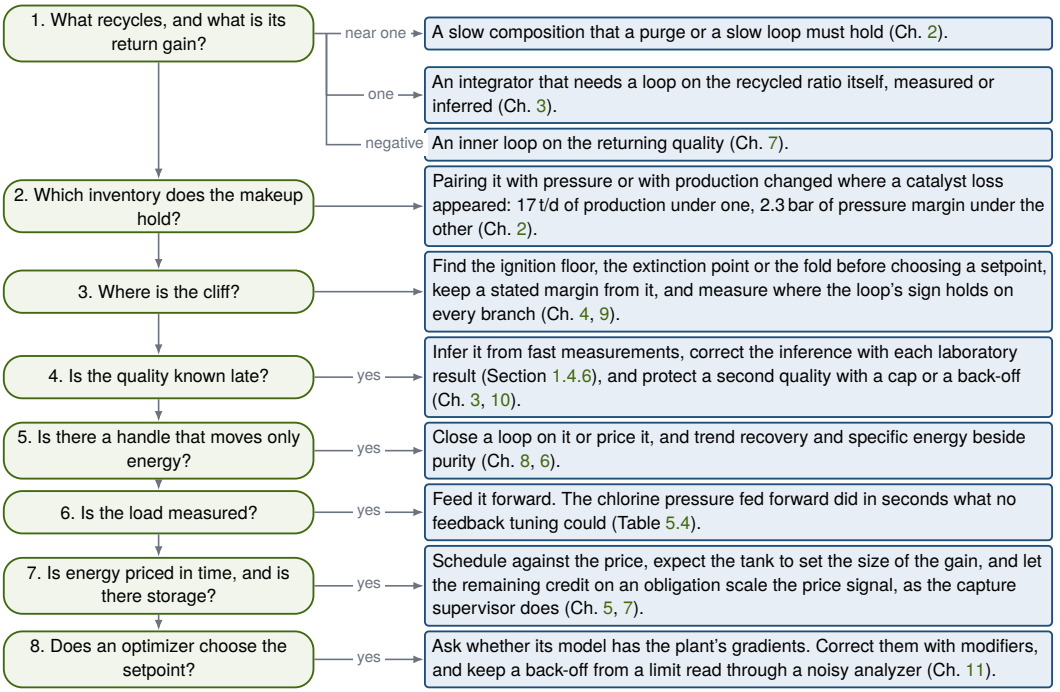


Figure 12.2 The eight questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions are about the plant and are answered before any loop is tuned.

6. *Is the load measured?* Feed it forward. The chlorine pressure fed forward did in seconds what no feedback tuning could (Table 5.4).
7. *Is energy priced in time, and is there storage?* Schedule against the price, expect the tank to set the size of the gain, and let the remaining credit on an obligation scale the price signal, as the capture supervisor does (Chapter 5, Chapter 7).
8. *Does an optimizer choose the setpoint?* Ask whether its model has the plant's gradients. Correct them with modifiers, and keep a back-off from a limit read through a noisy analyzer (Chapter 11).

A predictive controller earned its place where it coordinated two or more slow inventories or handles that the supervisory loops moved one at a time. It did not replace a missing loop or a wrong structure, and Volume 9 treats the choice in full.

12.3 Where Practice Is Heading

Plants follow the power price. A scheduling study of an air separation unit under time-sensitive electricity prices found significant savings from operating in a transient regime, raising production when energy is cheap and cutting it at peak price while using the storage available, and it built the schedule on low-order dynamic models so that the plan stays feasible [1]. A later framework lets power-intensive chemical processes take part in short-term electricity markets, whose prices move in seconds to minutes, and stresses that frequent changes of production target mean transient operation, so the schedule must respect the process dynamics [2]. The chlor-alkali scheduler and the capture supervisor of this volume are small cases of the same idea. In the editor's view the control structure

of a power-intensive plant will increasingly be judged by how far and how fast it lets the rate move without leaving its specifications.

Synthesis loops and capture plants run flexibly. A study of the Haber–Bosch process for power-to-ammonia concluded that the process can be operated over a wide operating envelope while keeping its hydrogen intake and ammonia production flexible [3]. Dynamic models and the optimization of flexible operation of post-combustion capture plants have been reviewed as a field of their own [4]. In the editor's view this moves the hard part of the ammonia loop from the design point to the turndown, where Chapter 2 found a converter that goes out between two quench positions a tenth of a percent apart, and the hard part of a capture plant to the obligation that outlasts its tanks.

> ⧖ HISTORY **2014 to 2020: A Capture Plant That Must Follow**
>
> The timeline of Chapter 1 closed in October 2014, when a capture plant on a 110 MW coal unit opened as the world's first fully commercial plant of its kind [5]. A second commercial plant reported in 2020 that it was designed for 50 to 100 percent of its carbon dioxide load and functioned best at 100. Its low-load point was set at 70 percent, limited by the turndown of the gas turbine whose waste heat raised most of its steam, and a hot start took 2 to 9.5 hours [6]. That year a review found that a capture plant needs more than two hours to settle after a cut in reboiler duty, far slower than the station upstream, and that most control studies still treated it apart from the station [7]. The plant of 2014 was built to capture, and gap 5 of Section 12.5 asks what it takes to follow as well.

Optimizers correct their gradients. An assessment of real-time optimization practice discusses the layer's place in the plant's hierarchy of decisions and its coordination with MPC [8]. Modifier-adaptation schemes reach the plant's optimum on convergence despite structural mismatch between plant and model, and need no parameter estimation. They depend on gradients estimated from noisy measurements, and a 2016 review surveys the methods for that estimate along with the open issues of the field [9]. Chapter 11 shows the cost of that estimate: every third run about three points short, and a back-off that bought robustness to noise with profit.

Intensified units bring control into design. A review of process systems engineering for intensification lists two key open questions: how to derive intensified designs systematically, and how to ensure the operability and optimality of the structures derived while they deliver their function [10]. Dynamic process intensification treats changes to a process's dynamics, operation and control as a source of efficiency in its own right, and names challenges in its design, operation, control and fault detection [11]. The dividing-wall column and the MTBE column are both intensified units, and both chapters found that the hard decision was which variable to control. In the editor's view the control structure of an intensified unit will be chosen with its design, not after it.

Compositions are read in line. Urea plants read the ratio with vibrating density sensors, Raman spectroscopy, chromatographs or the laboratory, and sensors that touch the carbamate corrode and drift. A gamma-ray gauge, which reads the density without touching the stream, has gained attention for that reason [12]. The loading of an amine solvent is still most often found by titration, which can take two hours, and a study of a gas plant's solvent found near-infrared spectroscopy able to replace it [13]. In the editor's view such a reading would have exposed the capture plant's inference that read 0.238.

Cells and catalysts lower the energy. Oxygen-depolarized cathodes make hydroxide from oxygen instead of making hydrogen, and lower the cell voltage by about 1 V at 4 kA/m^2 and the electricity by about 30 %. A reference document of 2014 lists them as an emerging technique, with one plant at industrial scale [14]. Cesium-promoted catalysts for the acid plant stay active at 380 to 400 °C and

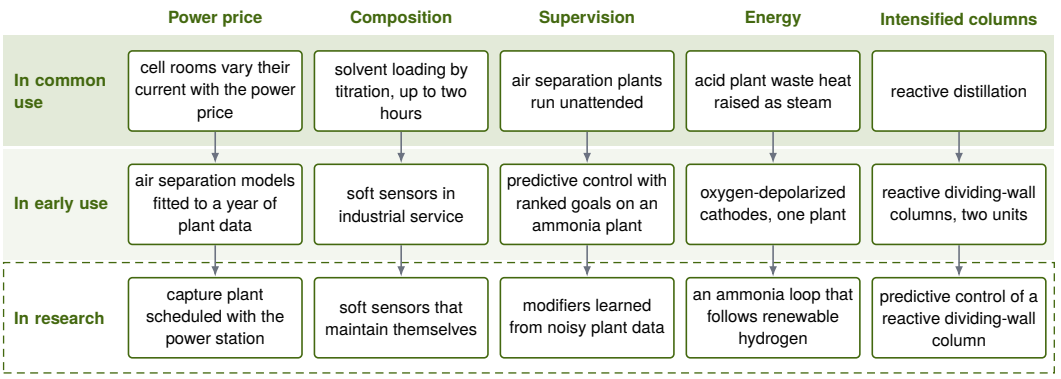


Figure 12.3 Developments in five threads. The top lane holds what the cited source calls standard or common practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, laboratory work or as a need, in the editor's reading. Sources by thread: power price [7, 14, 17], composition [13, 18], supervision [15, 16, 19], energy [3, 14, 15, 20] and intensified columns [21]. Notice that the cell room is the only plant of the volume whose rate follows the price in common use.

can serve in the first bed to lower its inlet temperature [15]. In the editor's view the ignition floor of Chapter 4 would then sit lower and still be a floor.

Plants run with fewer people. Computer-based control lets cryogenic air separation plants run unattended or with minimal staffing. Responses that an operator would give informally at an attended plant must then be designed into the controls, and the instrumented system must shut the plant down safely in an unplanned upset with no manual action [16]. In the editor's view this is the last working rule of Section 12.4.1 applied to a whole plant. Figure 12.3 sets these developments on three lanes.

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 [22]. 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 [23]. 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 every number to a run, and a referee report was written on each paper. Three of the ten reports call themselves self-reviews. The others do not say who wrote them. Defects went into a log for each app with a severity, a cause, a suggested fix and a status. For this book one agent per chapter reran the numbers its exercises use. 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 defect logs. The ten logs hold 101 entries with a severity, from two in the urea and acid logs to 22 in the chlor-alkali and capture logs. Of these, 22 are wrong behavior or output of an app, 62 are documentation, 10 are cosmetic or display and 7 are other kinds, such as a check that tested the wrong thing or a completion rule met on the clock alone. Five cases show the errors that matter for a volume about recycle, equilibrium and energy.

Helpers that found the extinguished converter. The ammonia app offered an operator aid that finds the makeup valve opening for a pressure setpoint by bisection. The bisection assumed that the settled pressure rises with the valve. It does not, because below about 45 % the steady-state solver, started from a fixed guess, lands on the extinguished branch, where nothing is consumed and the pressure runs to its ceiling. At the design line-up the aid set the valve to 20.0 % in place of 59.3 and reported a converter out at 260 bar. A second aid, for the purge, returned a 1.0 % purge for every inert setpoint from 14.5 to 18 %, which also puts the converter out. The paper agent found the first while measuring the app and flagged the same exposure in the second, and both aids now search the lit

branch only. Later, the capture app's predictive controller was found linearizing its absorber on a root with the bottom stage loaded past its maximum, from a fixed starting guess, which gave a capture gain of −7884 percent per unit of loading where the design gives −139. An equilibrium with more than one solution traps a solver the way it traps a plant.

A self-test that measured a pump on its stop. The capture app checked its own headline result, that the lean-loading loop holds the capture gain, with a built-in test. Two statements in the test ran in the wrong order, so the circulation command went to 300 and clamped at its stop of 170. The test divided the capture change by a step of 20 that had in fact been 40, and reported 0.271. The true value is 0.177. The 0.271 had spread to the slider hint, the controller comment, the model dialog, a scenario card and the guide. The paper agent found it by instrumenting the test and seeing the circulation on its stop for eight hours. A check that passes is also a claim.

A price that moved nothing. The capture app charged its carbon price per kilogram on a quantity already converted to tonnes, so the carbon cost was a thousand times too small. At design the stack loses carbon dioxide worth about $3 900 an hour, and the scorecard showed a few dollars after an hour. The price slider, visible from 0 to 300 $/t, changed nothing a student could see. In a volume where energy cost shapes the structure, the probe is simple: move the price and ask whether the margin moves.

A converged helper that left out the recycle. After the ethylene oxide finding of Volume 12, the chapter agents of this volume checked whether each app helper behind a printed number had converged. Nine of the ten records report the check, one of them for a helper that is explicit and needs none, and the tenth chapter uses no helper. Every iterative helper converged. One that converged still answered the wrong question. The capture app's "duty for this lean loading" solves the stripper once through, at the present rich loading. In the recycle the leaner solvent returns a lap later as a leaner rich feed. A request of 0.200 settled at 0.1792, and 0.260 at 0.2765. The chapter now says the duty holds the loading only until the changed solvent comes round, and the app owner has the fix.

A step no learner can take, and a setting tried at one seed. Nine of the ten records report checking that every value an exercise asks for lies on its control's grid, and four report walking the completion steps themselves. The air separation agent found that scenario 4 asks for a level loop gain of 0.4 on a control that moves in steps of 0.5, so the typed 0.4 becomes 0.5 and the step can never be met. The app's own test passed, because it set 0.4 through the program interface and not through the control. In Chapter 11 the first suggested setting passed the challenge at the default seed with 98.19 %. A technical review ran it at 31 seeds and found it passing at 21. The suggestion became a smaller filter gain with a waste back-off, which passes all 31 at 97.47 % (Section 11.6).

The book round. Each chapter agent first confirmed that its paper had not changed since conversion, and none of the ten had. The ten records list 154 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all ten report that every rerun number reproduced. The book round's defect list holds ten entries for this volume. In four apps a loop's action was labeled the opposite of Chapter 1's convention: the acid plant's bed 4 and analyzer loops, the dividing-wall split loop, the MTBE temperature loop in one case text and the base oil mid-ranging control. Each chapter now names the action as Section 1.6.2 does, and notes the app's label where it still differs. The base oil app has since been relabeled, and its Mid-ranging action control now reads Reverse for the setting Chapter 10 calls reverse acting. The MTBE paper's flow sheet labeled the overhead vapor with 7651, which is the boilup. The overhead vapor is about 10 932 kmol/h. The chapter agent left the label as printed and reported it, and the orchestrating agent removed the number from the

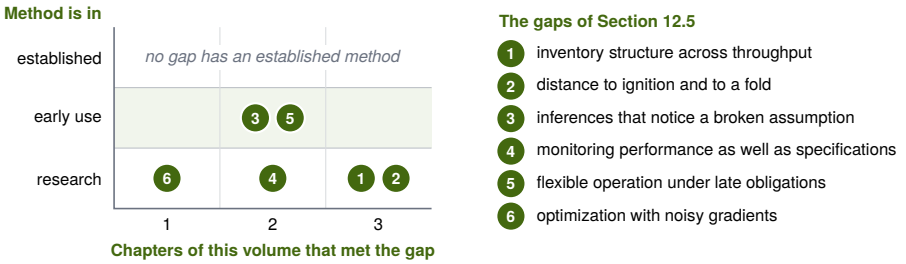


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 2, 3 and 7), three for gap 2 (2, 4 and 9), two for gap 3 (7 and 8), two for gap 4 (6 and 8), two for gap 5 (5 and 7) and one for gap 6 (11). Up is the maturity of the method available, the editor's reading of the sources cited in each item. Notice that the two gaps most chapters met have the least mature methods.

chapter's copy. The agent writing Chapter 1 first took a converter feed of 5000 and fifty passes from a well-mixed shortcut, and a mixer balance gave 4900 and forty.

Every error above was caught by a check that could fail: a helper's answer run in the plant, a test instrumented step by step, a price moved, a value typed into the control a learner uses, a setting run at many seeds, a balance closed around the mixer. None was caught by rereading. Four working rules follow. Start every solver where the plant is, and reject a root the plant cannot hold. Close the balance around the recycle before trusting a gain or a helper. Move every price and see what it moves. Walk every step through the controls a person will use.

12.5 Gaps in Chemical Plant Control

Six gaps belong to the control of chemical plants, and the gaps common to agent-era practice follow them in Section 12.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is 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. *Inventory structure across throughput.* The ammonia loop needed its slow loops in service within the first hours, and its pairing decided whether a loss appeared as production or as pressure margin (Chapter 2). The urea loop and the capture plant needed a loop on the recycled quantity too (Chapter 3, Chapter 7). A method that derives the inventory pairing and the slow composition loops of a recycle plant from its model, with stated time scales for each, and rechecks them when the rate changes, is not routine. A 2023 review finds a lack of understanding of how such control systems should be designed, and calls choosing the structure from a model a very difficult combinatorial problem even for small systems [24]. The slow loops are hard to judge because they act over hours that a steady-state analysis does not see. Progress would be a procedure that returns each pairing with its time scale.

2. *Distance to ignition and to a fold, measured online.* The ammonia converter goes out between two quench positions a tenth of a percent apart, and the MTBE column's branch monitor reads one tray and can name the wrong branch near a fold (Chapter 2, Chapter 9). The acid plant's best inlet is the coolest one its catalyst can still light (Chapter 4). An estimate of the margin to the nearest extinction or fold, from plant measurements, would let a structure keep a known margin rather than a guessed one. The Open challenge box after this list takes the gap further.

3. *Inferred compositions that notice a broken assumption.* The capture plant's inferred lean loading read 0.238 while the true value was 0.274 after a pressure drift, and the dividing-wall loops read

their setpoints while the true purities had moved (Chapter 7, Chapter 8). An inference that detects when a secondary variable has changed its meaning, before the laboratory does, remains open. A 2009 review finds that telling a drift of the process from a drift of the sensor needs expert knowledge, and that most soft sensors have no automated maintenance [18]. It is hard because a variable that has changed its meaning still moves smoothly with the plant. Progress would be an inference that flags its own broken assumption.

4. *Monitoring performance as well as specifications.* Fouling took eight percent of the oxygen with every loop on setpoint, a heat leak cut the recovery to 64.9 %, and a column on specification burned 17.7 % more steam (Chapter 6, Chapter 8). Monitors that track recovery, specific energy and the position of an energy minimum, and say which loop to change, are thinly developed. A dividing-wall column loses its energy saving away from its best vapor split, and the window of minimum energy may be narrow and move with the feed [25]. It is hard because recovery and specific energy have no setpoint and no alarm. Progress would be a monitor that reports the distance from the energy minimum beside each purity.

5. *Flexible operation under obligations that bind late.* Seven hours of high price put the capture plant's lean tank on its floor within 73 min against a four-hour average obligation, and the chlor-alkali scheduler was limited by its tank (Chapter 7, Chapter 5). Schedules that carry the plant's dynamics, storage and rolling obligations together, and degrade gracefully when a forecast is wrong, are still research. Chlor-alkali plants already vary their current with the electricity price [14]. A 2020 review finds that most control studies treat a capture plant apart from the power station it is coupled to, and that steady-state optimization cannot see the transients [7]. It is hard because an obligation averaged over hours answers for decisions taken hours before.

6. *Optimization with a wrong model and noisy gradients.* Modifier adaptation reached the optimum of Chapter 11 only by perturbing the plant, and a waste back-off traded profit for robustness to noise. Real-time optimization is used less than one would expect [26], and gradients inferred from noisy measurements are key to the reliability of modifier adaptation [19]. It is hard because each estimate costs a perturbation that the profit pays for. Gradient estimates that need less perturbation and come with a bound on the limit violation they allow would make such layers easier to trust.

> ⚠ **OPEN CHALLENGE** **How Far Is the Cliff?**
>
> Gap 2 ties with gap 1 as the gap most chapters met. A 2018 study of an autothermal ammonia converter for power-to-ammonia found that operating outside its envelope shuts down the autothermal operation or overheats the catalyst, and that such converters are run only at their optimum [3]. Multiple steady states appear even in ideal distillation, and the result is highly relevant in practice, because industrial columns set their inputs on a mass or volume basis [27]. The margin to a cliff is hard to measure because it is a property of a branch the plant is not on. Progress would be an online estimate of that margin, with its uncertainty, that a setpoint and an alarm can use.

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 [28]. 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 [23]. 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 [29]. 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 [30]. 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 [31]. 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 plant that sends material round again needs a loop on what goes round.
- An equilibrium puts the optimum next to a cliff, so the structure carries a margin from it or a measurement chosen so that crossing it can be undone.
- Energy is spent by the handle no purity loop watches, so recovery and specific energy are trended beside purity.
- Timing the energy is worth what the storage allows. Prediction paid on slow coupled inventories, and an optimizer needs the right slopes.

References

[1] R. C. Pattison, C. R. Touretzky, T. Johansson, I. Harjunkoski, and M. Baldea, "Optimal process operations in fast-changing electricity markets: Framework for scheduling with low-order dynamic models and an air separation application," *Industrial & Engineering Chemistry Research*, vol. 55, no. 16, pp. 4562–4584, 2016.

[2] J. I. Otashu and M. Baldea, "Grid-level "battery" operation of chemical processes and demand-side participation in short-term electricity markets," *Applied Energy*, vol. 220, pp. 562–575, 2018.

[3] I. I. Cheema and U. Krewer, "Operating envelope of Haber–Bosch process design for power-to-ammonia," *RSC Advances*, vol. 8, no. 61, pp. 34 926–34 936, 2018.

[4] M. Bui, I. Gunawan, V. Verheyen, P. Feron, E. Meuleman, and S. Adeloju, "Dynamic modelling and optimisation of flexible operation in post-combustion CO_2 capture plants—A review," *Computers & Chemical Engineering*, vol. 61, pp. 245–265, 2014.

[5] M. Bui, C. S. Adjiman, A. Bardow, E. J. Anthony, A. Boston, S. Brown, P. S. Fennell, S. Fuss, A. Galindo, L. A. Hackett, J. P. Hallett, H. J. Herzog, G. Jackson, J. Kemper, S. Krevor, G. C. Maitland, M. Matuszewski, I. S. Metcalfe, C. Petit, G. Puxty, J. Reimer, D. M. Reiner, E. S. Rubin, S. A. Scott, N. Shah, B. Smit, J. P. M. Trusler, P. Webley, J. Wilcox, and N. Mac Dowell, "Carbon capture and storage (CCS): the way forward," *Energy & Environmental Science*, vol. 11, no. 5, pp. 1062–1176, 2018.

[6] Petra Nova Power Holdings LLC and G. Kennedy, "W.A. Parish post-combustion CO_2 capture and sequestration demonstration project (final technical report)," U.S. Department of Energy, Office of Scientific and Technical Information, Tech. Rep. DOE-PNPH-03311, Mar. 2020.

[7] X. Wu, M. Wang, P. Liao, J. Shen, and Y. Li, "Solvent-based post-combustion CO_2 capture for power plants: A critical review and perspective on dynamic modelling, system identification, process control and flexible operation," *Applied Energy*, vol. 257, p. 113941, 2020.

[8] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[9] A. Marchetti, G. François, T. Faulwasser, and D. Bonvin, "Modifier adaptation for real-time optimization—methods and applications," *Processes*, vol. 4, no. 4, p. 55, 2016.

[10] Y. Tian, S. E. Demirel, M. M. F. Hasan, and E. N. Pistikopoulos, "An overview of process systems engineering approaches for process intensification: State of the art," *Chemical Engineering and Processing: Process Intensification*, vol. 133, pp. 160–210, 2018.

[11] M. Baldea and T. F. Edgar, "Dynamic process intensification," *Current Opinion in Chemical Engineering*, vol. 22, pp. 48–53, 2018.

[12] S. Z. Islami Rad and R. Gholipour Peyvandi, "A novel calibration method for online determination of nitrogen-to-carbon ratio in urea synthesis," *Scientific Reports*, vol. 16, no. 1, p. 36571, 2026.

[13] M. Yusop Nurida, D. Norfadilah, M. R. Siti Aishah, C. Zhe Phak, and S. M. Saleh, "Monitoring of CO_2 absorption solvent in natural gas process using Fourier transform near-infrared spectrometry," *International Journal of Analytical Chemistry*, vol. 2020, p. 9830685, 2020.

[14] T. Brinkmann, G. Giner Santonja, F. Schorcht, S. Roudier, and L. Delgado Sancho, "Best available techniques (BAT) reference document for the production of chlor-alkali," European Commission, Joint Research Centre, Luxembourg, JRC Science and Policy Report EUR 26844 EN, 2014.

[15] European Commission, "Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers," European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007.

[16] European Industrial Gases Association, "Guideline for safe practices for cryogenic air separation plants," European Industrial Gases Association, Brussels, EIGA Doc 147/22, 2022.

[17] C. Tsay, Y. Cao, Y. Wang, J. Flores-Cerrillo, and M. Baldea, "Identification and online updating of dynamic models for demand response of an industrial air separation unit," *IFAC-PapersOnLine*, vol. 54, no. 3, pp. 140–145, 2021.

[18] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[19] E. A. del Rio Chanona, P. Petsagkourakis, E. Bradford, J. E. Alves Graciano, and B. Chachuat, "Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation," *Computers & Chemical Engineering*, vol. 147, p. 107249, 2021.

[20] C. Smith, A. K. Hill, and L. Torrente-Murciano, "Current and future role of Haber–Bosch ammonia in a carbon-free energy landscape," *Energy & Environmental Science*, vol. 13, no. 2, pp. 331–344, 2020.

[21] A. A. Kiss and M. Jobson, "Taking reactive distillation to the next level of process intensification," *Chemical Engineering Transactions*, vol. 69, pp. 553–558, 2018.

[22] 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.

[23] 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.

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

[25] D. Dwivedi, J. P. Strandberg, I. J. Halvorsen, H. A. Preisig, and S. Skogestad, "Active vapor split control for dividing-wall columns," *Industrial & Engineering Chemistry Research*, vol. 51, no. 46, pp. 15 176–15 183, 2012.

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

[27] E. W. Jacobsen and S. Skogestad, "Multiple steady states in ideal two-product distillation," *AIChE Journal*, vol. 37, no. 4, pp. 499–511, 1991.

[28] 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.

[29] 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.

[30] 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.

[31] 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 to hold the N/C ratio of a new urea loop with a ratio station on the ammonia and carbon dioxide flow meters, retuned twice as fast as before, and reports that its simulation holds the ratio of the two flows within 0.1 percent. Using Chapter 3 and Chapter 2, explain why that report does not show the loop is held, and say what structure and what evidence you would ask for.

12.2 Discussion · 8 min

An optimization study shows that the ammonia converter of Chapter 2 would make 11 % more product with its hot spot near 480 °C, and an agent recommends moving the setpoint there. List what you would check before agreeing, and compare the margin the acid plant of Chapter 4 keeps from its ignition floor.

12.3 Discussion · 7 min

A monitoring agent reports that a dividing-wall column and an air separation unit are healthy because every loop has been on its setpoint for a week. Using Chapter 8 and Chapter 6, explain why that report could be true while both plants lose money, and name the measurements you would add to it.

12.4 Discussion · 8 min

Section 12.4.1 describes two operator aids in the ammonia app that returned an extinguished converter, and a capture helper that converged but left out the recycle. You are asked to write the acceptance tests an agent-written steady-state helper must pass before its answer is shown to an operator. List them, and say which defect each would have caught.

12.5 Discussion · 8 min

A plant manager asks whether a new amine capture unit, a chlor-alkali cellroom and an air separation unit on one site should each get a predictive controller. Using the comparisons of Chapter 7, Chapter 5, Chapter 6 and Chapter 8, argue for an order of investment and say what each predictive controller would have to beat.

12.6 Discussion · 9 min

An agent installs a two-step real-time optimizer on a reactor and reports that the fitted model matches the plant's five measured fractions within half a point and that the optimizer has settled. Using Chapter 11 and Section 12.4.1, explain why neither fact shows the plant is at its optimum, and propose the evidence you would require before the layer is left running.