INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 16

Summary and Outlook

Pages 378–391

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 16, pp. 378–391.

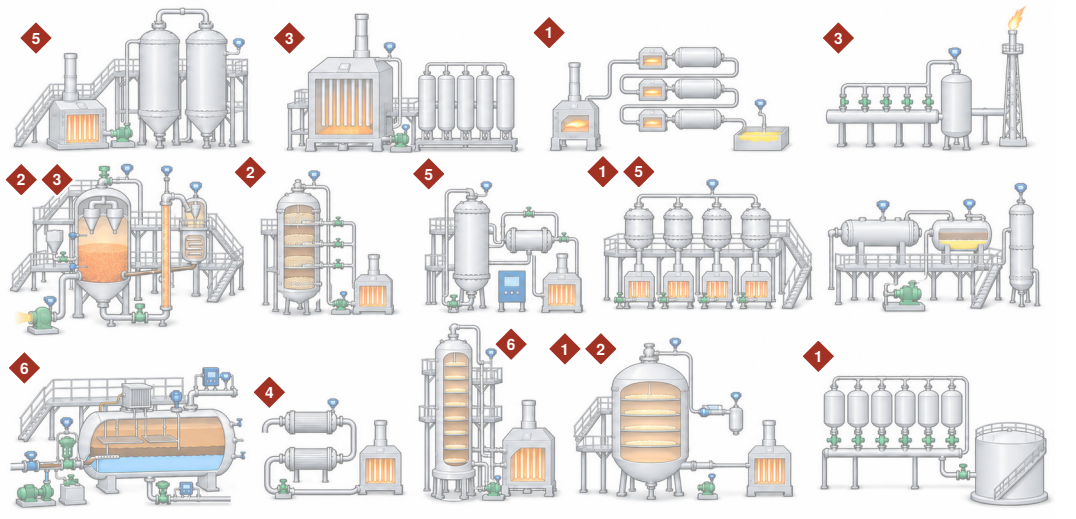
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v11-refining/

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

CHAPTER 16

Summary and Outlook



Open gaps 1 inferred qualities 2 estimated limits 3 shared utilities 4 cleaning schedules 5 quality margins 6 changing feed

Notice that each number stands beside every unit where a chapter met that gap, so the first gap appears four times and the fourth once. Illustration, not to scale.

IN PLAIN TERMS What Fourteen Refinery Units Taught

The main idea. Fourteen chapters followed crude oil from its first wash to the gasoline tank. In every unit the same three things pulled against each other: how much oil goes through, how good the product is, and how hard the equipment is pushed.

The challenge. The limit that stops a unit is often worked out from other readings, not measured, so it can be late or wrong. The instrument that judges the product can drift while it still reads on target.

Why it matters. A unit that trusts a wrong reading makes bad fuel or harms its equipment while every display looks normal.

Fourteen chapters followed crude oil from the desalter to the gasoline tank. Table 1.2 listed, for each unit, the handle that sets throughput, the quality the unit holds and how it is known, and the equipment limit with the structure that guards it. This chapter adds what worked, what failed and the measured result, and reads the comparison for the question of the volume, how each unit balances throughput, product quality and equipment limits. It then gives a short guide for choosing a design, 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.

16.1 What the Chapters Showed

Tables 16.1 and 16.2 compare the chapters, split between the crude and conversion units and the units that finish, supply and blend. The second column condenses Table 1.2 into the three quantities each

Table 16.1 How Each Unit Balanced Throughput, Quality and Limits, Chapters 2 to 8

Chapter	Throughput, quality, limit	What worked	What failed	Measured result
2, desalter	Charge set upstream. Salt sampled, the interface misread by the rag. Grid current and its trip.	A map optimizer on the mixing valve, a rag trim from samples and 15 ppm of demulsifier.	Design settings on a heavier crude. A level loop on a probe that reads half the rag.	Challenge salt 0.819 PTB, a pass, against a grid trip at 493.3 min as shipped (Table 2.5). True water 750 to 632 mm with the reading on setpoint (Table 2.4).
3, preheat train	Charge is the last handle. Column inlet measured. Firing limit as the train fouls.	A valve position override on the charge at a 95 % target. Cleaning outages kept apart.	Feedforward alone on the firing stop. Two outages 8 h apart.	With an exchanger out, the inlet falls to 340.3 °C, or holds 360.0 °C for 555 t of 19 200 (Table 3.5). IAE 559 against 48 °C h with one clean moved (Table 3.6).
4, crude tower	Charge set upstream. Four handles for eight inferred specifications. Overflash and heater duty.	Four loops along the one-way interaction, the gas oil draw on the overflash.	The top loop alone with cautious tuning, as shipped.	Challenge worst 0.28 vol% and 0.65 °C, dry 1.0 min, against 4.01, 1.92 and 174.2 min (Section 4.5.6). One reflux moves two specifications (Table 4.4).
5, vacuum tower	Lift by pressure and heater outlet. Gas oil 95 % point inferred. Ejector break and tube skin.	A lift maximizer that watches the ejector load and cuts the heater outlet as the critical load falls.	Fixed design setpoints. The same maximizer blind to the ejector load.	Challenge 1153.7 t, no break, against 1097.3 t with a break at 352.4 min, and 1161.1 t with a break when the load is ignored (Table 5.6).
6, catalytic cracker	Feed at a low selector. Severity as riser outlet. Blower, compressor, cyclone and valve limits.	A low selector over four constraint controllers with external reset.	The unit left alone through the shift.	$95,391 and 0 min above the cyclone limit, against $84,532 and 296 min (Table 6.9). Feed held at 284.5 t/h of 330 asked (Table 6.8).
7, hydrocracker	Feed. Conversion on a 25 min analyzer. Bed rise, runaway index 0.538 on bed 1.	An override on a rise compensated by one lag on the inlet.	The same override on the raw rise.	On a feed cut from 250 to 220 t/h the raw rise trips at 28.25 min, the compensated rise peaks at 29.9 °C and holds (Table 7.4).
8, diesel hydrotreater	Feed. Sulfur on a 30 min analyzer. Outlet temperature, heater duty and catalyst life.	Mid-ranging the bypass with the heater, a Smith predictor, a 7.5 wppm setpoint.	Mid-ranging off as the exchanger fouls. A 6.5 wppm setpoint.	Challenge worst 9.11 wppm at 1.21 catalyst days a day, against 386.97 shipped and 1.30 at 6.5 wppm (Table 8.5).

unit trades. The last column gives the measured result with the table or section of the chapter that reports it.

The columns make five points about the balance the volume asked about. Figure 16.1 draws one pair of results from each of twelve chapters.

Throughput is the handle of last resort, given up through a selector. On the preheat train the override held the column inlet at 360.0 °C through an exchanger outage at the cost of 555 t of 19 200, while feedforward alone held the full charge and delivered 340.3 °C with the fuel valve on its stop. On the catalytic cracker the low selector held 284.5 t/h against a request of 330 and gave back 23 t/h when the afternoon warmed from 15 to 38 °C (Table 6.8). The vacuum maximizer cut its heater outlet as the cooling water warmed and finished the challenge with more gas oil than the fixed setpoints. In

Table 16.2 How Each Unit Balanced Throughput, Quality and Limits, Chapters 9 to 15

Chapter	Throughput, quality, limit	What worked	What failed	Measured result
9, reformer	Feed. Octane from an analyzer, the laboratory and an inference. Tube metal and regenerator.	A trim on the inferred octane, feedforward from the naphtha, the aged heater distribution.	Any trim that trusts the shifted analyzer. Inlet temperature alone, as shipped.	In band 100 % of the judged time with tube metal at 606 °C, against 77.9 on the analyzer and 13.6 shipped (Table 9.5).
10, alkylation	Olefin feed. Acid strength titrated every 4 h. Refrigeration and the strength floor.	Loops with feedforward, a density soft sensor and moisture feedforward. A supervisor on the capacity rows.	Fixed flows. A strength loop on the titration alone.	Pass at $21,995, supervisor $23,839, fixed flows fail at a low of 93.52 RON (Table 10.7). Strength swing 0.86 points on titration, 0.11 on the soft sensor (Table 10.5).
11, delayed coker	Charge on a drum cycle. Heavy gas oil cut. Tube skin, foam and the drum switch.	Feedforward from the schedule and the shell thermocouples. Antifoam on the foam level.	Loops alone at a switch onto a cold drum. A fixed antifoam ratio.	Pressure dip 0.024 against 0.186 bar, cut within 1.0 against 9.7 K, antifoam 6.8 against 18.9 kg as shipped (Table 11.5).
12, hydrogen plant	Rate on header pressure. Steam to carbon floor, purity from the adsorber. Tube metal.	A lead-lag on demand, a firm header loop, a cycle that follows the feed, tail-gas feedforward.	A plain ratio on a ramping steam supply. The firm loop without the feedforward.	Passes at five seeds, header within 0.395 bar. Without the lead-lag the ratio falls to 2.711, without feedforward the tube reaches 927.1 °C (Table 12.4).
13, sulfur recovery	Acid gas set upstream. Air demand on a 3.2 min analyzer. Stack permit and dew point.	Flow and analysis feedforward under a PI trim.	The trim alone on a flow step. A trim that trusts a drifting analyzer.	Air demand IAE 7.3 with the flow ratio, 1227 without (Table 13.6). A 0.5 mol% drift puts 9.60 t/d up the stack (Section 13.5).
14, fuel gas header	Off-gas set by other units. Pressure and Wobbe index measured. Importing and flaring at once.	A split range with a dead band, Wobbe feedforward, recovery stages set below the seal.	A 4 % overlap. Stages above the seal. Feedforward off.	Challenge flares 6.1 to 6.3 Nm^3 with no time importing and flaring at once, against 7203 and 480 min shipped (Table 14.6).
15, gasoline blender	Blend rate into a tank. Tank-average octane and vapor pressure. Component stocks and give-away.	An optimizer with cumulative targeting and the laboratory correction.	A fixed recipe. Targeting off. The analyzer's bias update alone.	Certified at 87.19 with $14,511 of give-away, against 85.97 off specification shipped. Targeting on and off 87.15 and 86.96 (Table 15.5).

each case the unit gave up rate only while a limit bound and only as much as the limit asked. A unit that keeps its rate and lets the limit decide loses the quality instead.

A guard protects only as well as the signal it reads. The hydrocracker's override tripped the unit 8.25 min after a feed cut when it read the raw bed rise. The inlet answers a quench valve within 0.30 min and the outlet lags by a 3.5 min catalyst time constant, so the raw difference was inverted during the correction (Section 7.10). One lag on the inlet fixed it. The vacuum maximizer blind to the ejector load ran the heater outlet to 425 °C and the ejectors to 149 % of their critical load, against 85 % when the load was watched (Table 5.5). A limit is usually inferred, from a difference of two temperatures or from a load, and the inference has dynamics of its own.

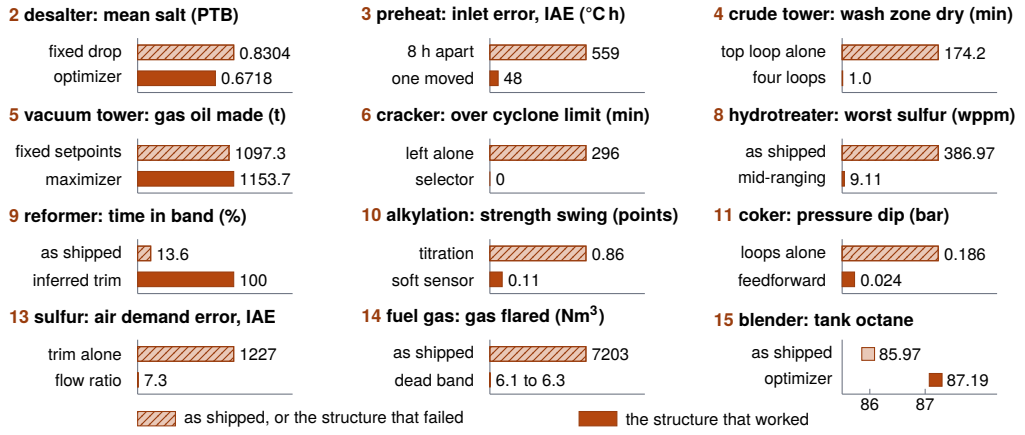


Figure 16.1 One pair of results from each of twelve chapters, with the numbers as Tables 16.1 and 16.2 print them (the desalter pair as this section quotes it from Table 2.6). Each panel has its own scale from zero. The blender alone is drawn with markers, on an axis that starts at 85.5. A larger value is better in the panels of Chapters 5, 9 and 15. The tables give the hydrocracker and the hydrogen plant no like pair. Notice that where a measured load was fed forward or a limit was guarded, the solid bar nearly vanishes.

Quality margin is paid for in catalyst, yield or tube life. The diesel hydrotreater passed its graded day at a 7.5 wppm setpoint with 1.21 catalyst days used per day and failed at 6.5 wppm on 1.30, although its worst sulfur reading improved. On the reformer 10 K of inlet temperature bought 2.01 RON for 1.51 points of liquid yield (Section 9.4). On the coker the same 10 K traded 1.9 points of coke for coils that coke 2.2 times as fast (Table 11.4). The setpoint of a quality loop is therefore an economic choice. A margin wide enough for the analyzer's noise and no wider is the target, and only a sweep over seeds or feeds shows where it lies.

A late quality is held by what is measured early, and lost to an analyzer that is trusted. A flow ratio on the sulfur plant cut the air demand error after a flow step from 1227 to 7.3. The coker's schedule feedforward cut the pressure dip of a cold drum from 0.186 to 0.024 bar. The alkylation soft sensor held the strength within 0.11 points where the four-hourly titration let it swing 0.86. The opposite failure recurs as well. The reformer trim that believed a shifted analyzer held its reading while the true octane fell to 94.24 (Table 9.4). The sulfur trim that believed a drifting analyzer sent 9.60 t/d up the stack with the air demand on target. The blender certified only with its laboratory correction on, at 87.06 against 86.45 (Table 15.5). Feedback holds the reading. A second source is needed to hold the product.

Optimization paid where it set the target. Prediction paid less than it first appeared to. The interior optima were found by an optimizer. On the desalter the map optimizer reached a mean salt of 0.6718 PTB where extremum seeking reached 0.7842 and the fixed pressure drop 0.8304 (Table 2.6). The predictive controllers were mixed. On the crude tower's end point steps the plan cut the freeze point IAE from 123.9 to 18.1 °C min and lost the heavier crude (Table 4.6). On the alkylation unit it held the strength between titrations at an IAE of 1.20 against the supervisor's 6.18 (Table 10.6). On the hydrocracker its higher conversion, 0.5501 against 0.5427, came from the bed 1 target, and the loops reproduced it with that setpoint lowered by hand (Table 7.5). On the preheat train its extra tonnes came from the valve's last five points, and at the override's 95 % ceiling it processed less (Table 3.8). On the cracker it crossed no limit in seven cases at about one percent less feed (Table 6.10). Volume 9 treats that choice in full.

> ⧖ HISTORY **2003: What a Survey of Predictive Control Left Open**
>
> The survey that counted the predictive controllers of Chapter 1 also listed what they could not yet do [1]. Most were built on models from a plant test, which ran day and night and could take up to 5 to 15 days, and no tool told the engineer when the data were adequate. No method verified whether a model had deteriorated after commissioning. Only two adaptive algorithms had reached the market, despite a strong incentive for a controller that tunes itself. The practitioners' questions were about engineering, not algorithms: when a model is good enough, and how to design control for an entire plant. In 2023 a review found new applications still built mainly on linear models from step responses [2]. Gaps 3 and 6 of Section 16.5 are those questions.

16.2 Choosing a Method

The chapters suggest an order of questions for a refinery unit. Each names the chapter where its answer decided the result. Figure 16.2 sets the same questions out as a flowchart.

1. *Which limit binds first, and how is it known?* List the limits and the signal behind each. A limit read from a difference of two temperatures or from a load needs its own dynamics checked, as the compensated rise of Chapter 7 and the ejector load of Chapter 5 show.
2. *Which handle may give up throughput?* Put it at a selector with external reset, with a valve or capacity target below the stop, so that rate falls only while a limit binds (Chapter 3, Chapter 6, Volume 4).
3. *Does the unit feed back its own heat?* Compute the loop gain of Table 1.1 before tuning. A gain near one slows and amplifies every move, and a fouling exchanger can push the inlet out of reach of the heater (Chapter 8).
4. *Is the optimum on a limit or inside the range?* An optimum on a limit is held by a selector. One inside is found by an optimizer, as on the desalter, or turned into a setpoint at zero, as the air demand of Chapter 13 is.
5. *How late is the quality, and what is measured ahead of it?* Feed forward the measured load first, then trim on the analyzer, and keep an inference or a laboratory between the analyzer and the product (Chapter 13, Chapter 9, Chapter 10).
6. *Is the product judged as an average?* Steer the remainder, not the instant, and expect the authority to shrink as the tank fills (Chapter 15).
7. *What does the margin cost?* Choose the quality setpoint from a sweep over noise and feeds against the catalyst, yield or tube budget (Chapter 8).
8. *What does a predictive controller add over these structures?* Compare it with the structure given the same target and the same ceiling. In this volume most of its first apparent gain belonged to one or the other.

16.3 Where Practice Is Heading

Optimization joins planning and scheduling. Enterprise-wide optimization was described in 2005 as a new frontier that integrates planning, scheduling, real-time optimization and inventory control across the functions and decision levels of a company [3]. A review of refinery operations gives that integration of decision levels as the emphasis for the industry, from crude unloading and mixing through the conversion units to blending and distribution, and surveys the methods for scheduling and planning [4]. Model-based real-time optimization of gasoline blending was reported in 2000 [5]. In the editor's view the coordinating layer of Section 1.4.7 will reach further across units, and the blender's cumulative target, set from the tank and not the instant, is the pattern it will follow.

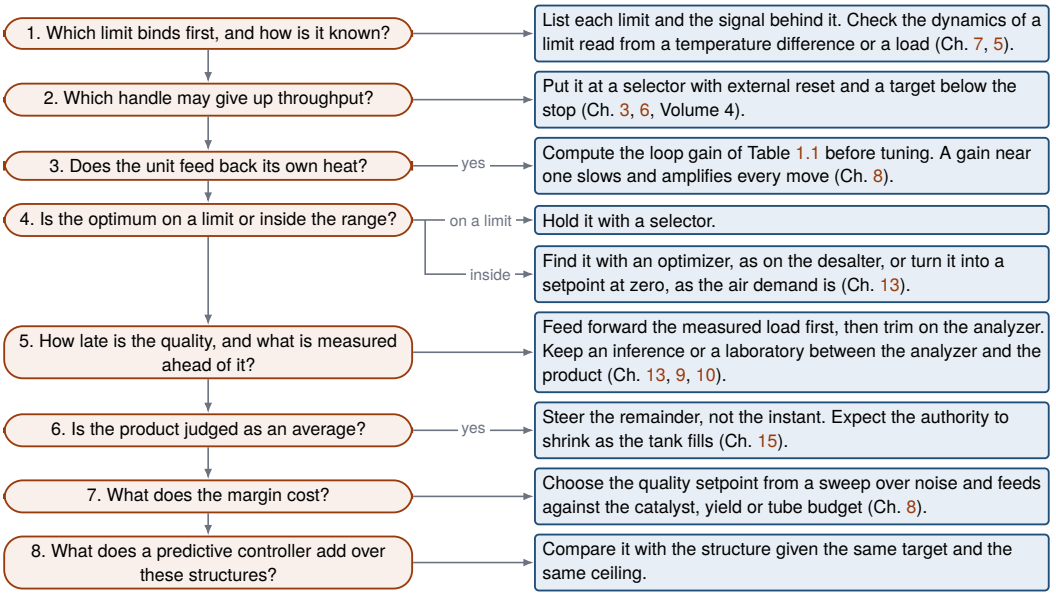


Figure 16.2 The eight questions of this section in order, each with the answer the list gives and the chapters where that answer decided the result. The third and the fourth question branch. Notice that seven of the eight are settled before a predictive controller is considered.

Models are maintained, not only commissioned. A 2015 review of industrial predictive control reports that practitioners now focus on easy commissioning, monitoring and the automated maintenance of applications, with continuous performance monitoring and automated reidentification of models [6]. A review of soft sensors for regression treats their maintenance alongside input selection, training and validation [7]. A perspective on machine learning for process data gives three attributes that process data analytics needs for such methods to apply in the process industries, and stresses a culture of data analytics [8]. The reformer's shifted analyzer and the vacuum tower's inferential that read 558 °C against a true 554 °C at a stripping steam it was not calibrated for (Section 5.5.4) are the cases such maintenance exists for.

Fouling and cleaning are scheduled from models. Optimal cleaning schedules for crude preheat trains have been posed as mixed-integer nonlinear programs, with fouling models reconciled from plant data on a network of 27 exchangers [9]. A review of threshold fouling models finds the most activity in using them to predict fouling, and gives the need to include the chemistry of fouling as one reason the models themselves have advanced little [10]. On the preheat train of Chapter 3, moving one of two colliding outages 32 h later cut the integrated error of three weeks from 559 to 48 °C h, which no tuning in the application could match.

The feed changes. Co-processing of bio-based feedstocks with petroleum fractions, through hydroprocessing and catalytic cracking, has been demonstrated industrially in some cases and reviewed as an option that uses existing refinery units [11]. In the editor's view such feeds will move the heat balances of the cracker and the hydrotreaters and the hydrogen they draw from the plant of Chapter 12, so the limits of these chapters will bind at new places, and models fitted to one crude will need the maintenance described above.

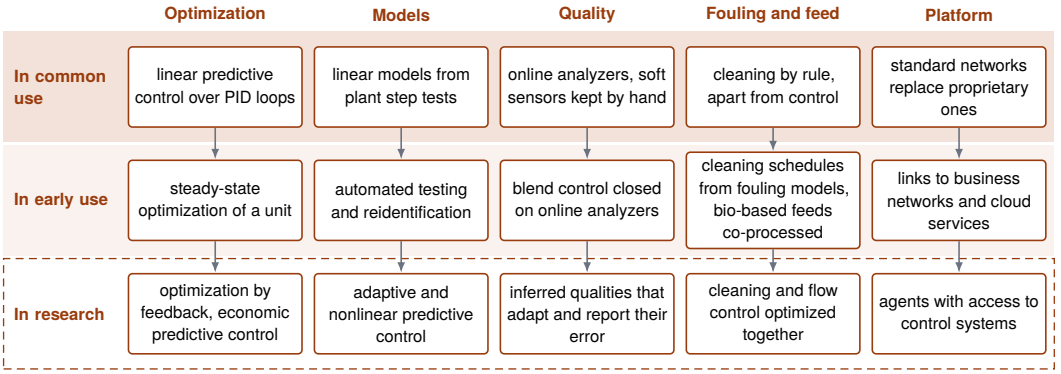


Figure 16.3 The developments of this section in five threads. The top lane holds what the cited source describes as standard practice, the middle lane what is in service at some sites or shown on plant data, and the bottom lane research. The placing is the editor's reading of the sources, and no box carries a date. Sources by thread: optimization [1, 2, 12], models [1, 2, 6], quality [13, 16], fouling and feed [9, 11, 17] and platform [15]. The last platform box is the fourth gap of Section 16.5.1. Notice that every thread ends in something that has to be kept right, a model or a measurement.

Optimization moves into the control layer. Most commercial real-time optimizers repeat two steps, an update of a steady-state model from plant data and a numerical optimization of the setpoints. A 2022 review states that this form is not as widely used as one would expect, and gives as reasons the cost of the model, the wait for steady state before each update and conflicts with the planning layer [12]. The same review surveys two decades of research on the alternative, in which the economic objective is translated into controlled variables and feedback loops bring the unit to its optimum without an optimization solved online. Joining control and optimization in one economic predictive controller is reported as rarely a good solution for practical implementation and tuning [2]. In the editor's view the selector that holds a unit on its binding limit, and the air demand held at zero in Chapter 13, are small cases of that alternative.

Analyzers close more loops. A blend controller that closes the loop on online analyzers downstream of the blender has been in service at one refiner in its present form since 2001, and by 2010 it ran on 17 blenders in six refineries [13]. Near-infrared spectroscopy, the analyzer of Chapter 9 and Chapter 15, takes one minute or less per sample and suits in-line use, and it stays dependent on a good reference method [14]. In the editor's view each loop closed on an analyzer adds a case of the first gap of Section 16.5, because the loop holds the reading whether or not the reading is right.

Control systems join business networks. A federal guide records that control systems, at first isolated on proprietary protocols and specialized hardware, are taking up Ethernet, Internet Protocol and wireless devices and are becoming more integrated into business networks and cloud infrastructures [15]. The integration leaves them much less isolated from the outside world and creates a greater need to secure them. In the editor's view the optimizers and the model maintenance of this section depend on that connection, and the security gap of Section 16.5.1 grows with it. Figure 16.3 sets the developments of this section on three lanes, by how far each has come into use.

16.4 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.

16.4.1 How This Volume Was Made

Each of the fourteen apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run, and a referee report was written on each paper. At least ten of the fourteen reports were written by the agent that wrote the paper, reading its own draft as a referee would. Seven papers had a further referee pass on a controller added later, a predictive

controller on five, extremum seeking on one and a Smith predictor on one. 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, after a repair round on 2026-11-15 had applied the book's earlier defect list to apps and papers. 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 fourteen logs written beside the papers hold 88 entries with a severity. Of these, 30 are wrong behavior or output of an app, 44 are documentation, 9 are cosmetic or display and 5 are other kinds, such as a fragile test or a session blocker. The counts understate the work, because nine apps keep a second log of their build in their own folders, which is not counted here. Four cases show the errors that matter for a volume about throughput, quality and limits.

An optimizer that reported optimal beyond its own limit. The alkylation unit's economic supervisor kept its targets beyond its own reboiler row once a feed increase put them there, and its card read "optimal". In the challenge under the supervisor the targets sat beyond the row for 1,290 of 1,440 minutes, the reboiler was at its steam limit for 1,158 minutes and the alkylate carried up to 21.8 % butanes. The cause was a linearized row that only stopped the step from moving further out. The agent adding the predictive controller found it, reported it without fixing it, and a separate fix agent made the row pull the targets back and report "infeasible" where it could not. The fix moved the supervisor's challenge row to the $23,839 of Table 16.2. A regression test fails on the old code. The scored octane had not shown the fault, because the verdict read the octane correlation and not the butanes.

Protections that read what no instrument measures. The hydrocracker's override cut the inlet setpoint on a runaway index computed from the plant's true kinetics and catalyst activity, so it reacted to an activity change before any thermocouple moved. The fixed override estimates the index from the rise it measures, and no shipped run moved. The catalytic cracker's afterburn and bed fences were computed from the true temperatures and not from the thermocouples the app models, a difference of up to 20.4 °C at the promoter loss of the challenge. That fix moved the challenge margin from $95,402 to $95,391, the number in Table 16.1. The fuel gas header's predictive controller took the valve capacities from the run's settings, so a makeup capacity edited from 8000 to 12 000 Nm^3/h moved its first move from 20.94 to 7.59 % with nothing measured changing. Each was caught by an audit that changed a hidden value at fixed measurements and asked whether a controller moved.

A limit computed from a lagged quality. On the hydrogen plant the predictive controller's ceiling on the adsorber rate scaled the rate by a purity reading that lags it through a drum and an analyzer, and applied it to a rate already cut. With the cycle fixed, the rate fell from 130 to 45 % in 9 min, the true steam to carbon ratio swung from 1.91 to 7.35 and the exercise stopped at 83 min. On the hydrocracker the first price on the bed 1 inlet limit let the plan warm the bed 2.1 °C past that feed's fold for 5.3 min after a trip reset. The hydrocracker case was found by the technical review of the added controller, which ran restart cases the author had not. The hydrogen case was fixed in the same audit round as the reads of the previous paragraph, with a regression test for it.

Gains credited to the wrong layer. Three referee passes on the predictive controllers corrected the credit. The hydrocracker draft credited the higher conversion to the controller, but the rise constraint bound in none of its 480 plans, and the shipped loops with bed 1's setpoint lowered by hand reproduced the result to four figures. The preheat draft said the plan "cuts only the charge it must", but at the override's valve ceiling it processed 92 t less, so the tonnes were the valve's last five points. The crude tower draft led with the case the plan won on every tile and left the heavier crude,

which it lost on every quality, for later. Each correction added a like-for-like row to the chapter's table.

The book round. Each chapter agent first confirmed that its paper had not changed since conversion, and none of the fourteen had. The fourteen chapter records list 216 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all fourteen report that every rerun number reproduced. The book round's defect list holds eleven entries for this volume. None reports a new fault in control behavior. They are text, display and presentation, and one start-up limitation the owner had already chosen to keep. Two papers named a controller's action against the convention of Chapter 1, and the chapters now state it. Two chapter summaries contradicted their own measurements and were aligned. The fuel gas chapter said four times the header volume buys four times the time, against its own 45 and 133.5 s, and the crude tower chapter called a ratio of 0.61 fixed where its own runs give 0.63 over a wider range. The agent writing Chapter 1 first wrote that each optimum lies on a limit, and two chapters contradicted it. The book's tools were checked as well. Three chapter records noted wide drawings, scaled down to the text width, that the build's check report had not listed. The orchestrator traced the cause to a look-ahead error in the width log of the book's style, which had never fired in a real build, so an earlier recheck of four volumes was void. The log was fixed, the earlier volumes were checked again, and three more drawings were redrawn.

Every error above was caught by a check that could fail: a hidden value moved at fixed measurements, a structure given the controller's own target, a case run that the author had not run, a log that should have listed a known wide figure. Four working rules follow. Read every limit a controller enforces from an instrument and say which. Compare a new controller with the old structure at the same target and the same ceiling. Ask an optimizer what it does when it starts outside its constraints. Let a person decide what changes.

16.5 Gaps in Refinery Control

Six gaps belong to the control of refining units, and the gaps common to agent-era practice follow them in Section 16.5.1. In the editor's view none is closed. Each item gives the evidence of the chapters, the state of the literature, why the gap is hard and what progress would look like. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 16.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the units in the picture that opens this chapter.

1. *Inferred qualities that know when they are wrong.* The vacuum tower's inferential misread the gas oil end point by 4 °C at a stripping steam it was not calibrated for, the reformer's trim followed a shifted analyzer, and the sulfur plant's drift was seen only by the stack alarm. The blender certified only with its laboratory correction on. In the editor's view routine detection of a drifting analyzer or inferential against a second, independent source is not standard practice, and a review of soft sensors found that most have no automated mechanism for their maintenance [16]. The Open challenge box after this list takes up why the gap is hard and what progress would look like.

2. *Limits that are estimated, with known lag and bias.* The ejector's critical load, the bed's runaway index and the cracker's fences are computed, not measured. A protection built on such an estimate needs a statement of how late and how wrong it can be, and the hydrocracker's raw rise shows what happens without one. The investigation of a refinery explosion drew the same lesson: operating parameters that give only indirect information on a critical condition can leave a unit unable to tell that it is unsafe [20]. The gap is hard because the estimate has dynamics of its own and its error shows only when the limit is near. Progress would be a stated delay and error band for every computed limit, with the margin of the protection set from them.

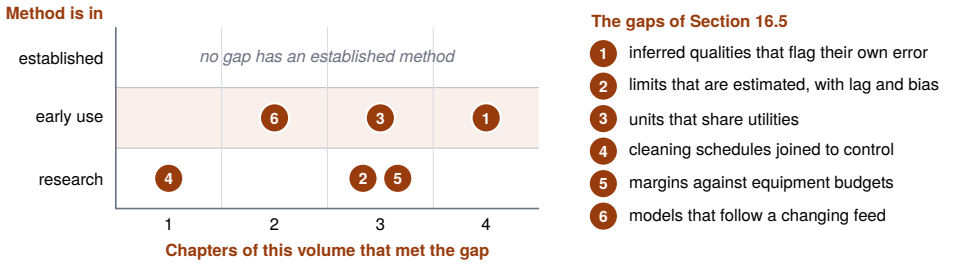


Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: four for gap 1 (Chapters 5, 9, 13 and 15), three for gap 2 (5, 6 and 7), three for gap 3 (6, 12 and 14), one for gap 4 (3), three for gap 5 (8, 9 and 11) and two for gap 6 (2 and 4). Up is the maturity of the method available, which is the editor's reading of the sources cited in each item and not a measurement. Notice that the top band is empty.

3. *Coordination of units that share utilities.* The fuel gas header is an integrator that nobody schedules, the hydrogen plant takes its rate from the pressure of the header it supplies, and the cracker gives back 23 t/h of feed on a warm afternoon. A coordinator that knows which unit's limit binds across the refinery, and moves rate there, remains an engineering problem. A 2022 review of real-time optimization counts the conflict between the optimum of one unit and that of the whole operation among the reasons the method is used less than expected [12]. The gap is hard because a shared header couples the units within seconds or minutes, while the layer that knows the prices acts each hour or each day (Figure 1.10). Progress would be a coordinating layer that reads the binding limit of each unit and moves rate between units within minutes.

4. *Cleaning schedules joined to control.* On the preheat train the phasing of two outages mattered more than any tuning. The cleaning schedule, the override's target and the crude slate are still decided by different people on different time scales. In the literature the cleaning schedule and the flow control of a preheat train have usually been treated independently and solved by heuristics, and a closed-loop scheme that optimizes both has been shown on a simulated industrial train [17]. The gap is hard because the decisions span weeks to years and cleaning crews are contracted with long notice, so operators ask for a schedule that does not change at every update [17]. Progress would be an override target and a cleaning date computed from one fouling model and proven on a plant.

5. *Margins chosen against equipment budgets.* The hydrotreater's 7.5 wppm setpoint came from a sweep over 31 seeds against a catalyst budget. The reformer and the coker paid for the same kind of margin in liquid yield and in the life of the heater coils. Any back-off from an active constraint is an economic loss [2], and how large the back-off must be was named in 2000 as an area in clear need of more research [21]. The gap is hard because the margin depends on disturbances that have not yet happened and on a price for catalyst or tube life that no instrument reports. A method that sets each quality back-off from the analyzer's noise, the feed's variability and the cost of the margin in catalyst, yield or tube life would replace that sweep.

6. *Models that follow a changing feed.* On the desalter no mixing pressure drop met the water specification on a heavier crude at the design demulsifier dose, and the crude tower's predictive controller lost the heavier crude on every quality. A review of catalytic cracking names light, tight oils as a challenge because their low carbon residue makes the coke deposition, and with it the heat balance of the unit, harder to control [22]. The gap is hard because new predictive controllers are still built mainly on linear models from plant tests, which are costly and slow to repeat [2]. A model that tracks the crude slate, and says when it no longer can, is needed before predictive control is trusted on a varying feed.

⚠ OPEN CHALLENGE An Analyzer That Reports Its Own Error

Gap 1 is the one most chapters met. A review of soft sensors found that there is often no objective measure of their quality, so the judgment rests on a person who compares the predictions with the correct values by eye [16]. The difficulty is that a drift of the instrument and a drift of the process are hard to tell apart in the data, can progress very slowly and take expert knowledge to separate [16]. A spectroscopic analyzer adds its dependence on a good reference method [14]. Progress would be an inferred quality delivered with its own error bound and checked against each laboratory result without a person.

16.5.1 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

- A unit gives up oil flow only while a limit binds, and only as much as the limit asks.
- A guard protects only as well as the signal it reads.
- The setpoint of a quality loop is an economic choice.
- Feedback holds the reading. A second source is needed to hold the product.
- A predictive controller is compared with the older structure at the same target and the same ceiling.

References

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

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

[3] I. Grossmann, "Enterprise-wide optimization: A new frontier in process systems engineering," *AIChE Journal*, vol. 51, no. 7, pp. 1846–1857, 2005.

[4] N. K. Shah, Z. Li, and M. G. Ierapetritou, "Petroleum refining operations: Key issues, advances, and opportunities," *Industrial & Engineering Chemistry Research*, vol. 50, no. 3, pp. 1161–1170, 2011.

[5] A. Singh, J. F. Forbes, P. J. Vermeer, and S. S. Woo, "Model-based real-time optimization of automotive gasoline blending operations," *Journal of Process Control*, vol. 10, no. 1, pp. 43–58, 2000.

[6] M. G. Forbes, R. S. Patwardhan, H. Hamadah, and R. B. Gopaluni, "Model predictive control in industry: Challenges and opportunities," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 531–538, 2015.

[7] F. A. A. Souza, R. Araújo, and J. Mendes, "Review of soft sensor methods for regression applications," *Chemometrics and Intelligent Laboratory Systems*, vol. 152, pp. 69–79, 2016.

[8] S. J. Qin and L. H. Chiang, "Advances and opportunities in machine learning for process data analytics," *Computers & Chemical Engineering*, vol. 126, pp. 465–473, 2019.

[9] F. Smaïli, V. S. Vassiliadis, and D. I. Wilson, "Mitigation of fouling in refinery heat exchanger networks by optimal management of cleaning," *Energy & Fuels*, vol. 15, no. 5, pp. 1038–1056, 2001.

[10] D. I. Wilson, E. M. Ishiyama, and G. T. Polley, "Twenty years of Ebert and Panchal—what next?" *Heat Transfer Engineering*, vol. 38, no. 7–8, pp. 669–680, 2017.

[11] S. Bezergianni, A. Dimitriadis, O. Kikhtyanin, and D. Kubička, "Refinery co-processing of renewable feeds," *Progress in Energy and Combustion Science*, vol. 68, pp. 29–64, 2018.

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

[13] M. Chèbre, Y. Creff, and N. Petit, "Feedback control and optimization for the production of commercial fuels by blending," *Journal of Process Control*, vol. 20, no. 4, pp. 441–451, 2010.

[14] C. Pasquini, "Near infrared spectroscopy: fundamentals, practical aspects and analytical applications," *Journal of the Brazilian Chemical Society*, vol. 14, no. 2, pp. 198–219, 2003.

[15] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[16] 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.

[17] F. Lozano Santamaria and S. Macchietto, "Stability of optimal closed-loop cleaning scheduling and control with application to heat exchanger networks under fouling," *Processes*, vol. 8, no. 12, p. 1623, 2020.

[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] U.S. Chemical Safety and Hazard Investigation Board, "ExxonMobil Torrance refinery electrostatic precipitator explosion, Torrance, California, incident date February 18, 2015," U.S. Chemical Safety and Hazard Investigation Board, Washington, DC, Investigation Report 2015-02-I-CA, 2017.

[21] S. Skogestad, "Plantwide control: The search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[22] E. T. C. Vogt and B. M. Weckhuysen, "Fluid catalytic cracking: recent developments on the grand old lady of zeolite catalysis," *Chemical Society Reviews*, vol. 44, no. 20, pp. 7342–7370, 2015.

[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.

16.6 Exercises

16.1 Discussion · 8 min

An agent proposes raising the valve target of the preheat override in Chapter 3 from 95 to 99 % of the fuel valve, to win back crude. Use Table 3.8 for the exchanger outage (scenarios 5 and 6) and the harder outage of scenarios 7 and 13. Say what the proposal gains, what it costs, and what you would ask for before accepting it.

16.2 Discussion · 8 min

The override of Chapter 7 tripped the unit on a feed cut when it read the raw bed rise and held it when the inlet was lagged. A new unit is to protect its beds with an override written by an agent. Say what you would test about the signal the override reads, and relate the hydrocracker's lesson to the vacuum maximizer of Chapter 5 and to the cases of Section 16.4.1.

16.3 Discussion · 7 min

On the graded day of Chapter 8 an agent proposes a sulfur setpoint of 6.5 wppm instead of 7.5, because the worst reading then falls to 7.94 wppm. Discuss the proposal with Table 8.5 and the trade-offs of Chapter 9 and Chapter 11, and say how a quality setpoint should be chosen.

16.4 Discussion · 8 min

An agent reports that a predictive controller on the hydrocracker raised the mean conversion through a feed cut from 0.5427 to 0.5501 and held the bed rise at 27.5 °C. Using Table 7.5, the preheat comparison of Table 3.8 and Section 16.4.1, say what you would ask before crediting the gain to prediction.

16.5 Discussion · 8 min

Three chapters held a quality from an instrument that was wrong without saying so: the tail gas analyzer of Chapter 13, the octane analyzer of Chapter 9 and the blend analyzer of Chapter 15. For each, say what the loop showed, what the product did and which second source caught it. Then propose what a monitoring agent should watch on such a unit.

16.6 Discussion · 8 min

A refinery wants a coordinator, proposed by an agent, that raises the crude rate until the first unit reaches a limit. Using Chapter 3, Chapter 5, Chapter 6 and Chapter 14, discuss what the coordinator must know about each limit, how the limits move during a day, and which decisions should stay with the operators.