INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 10 Control in Upstream and Midstream Oil and Gas

CHAPTER 16

Summary and Outlook

Pages 377–390

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Upstream and Midstream Oil and Gas*, John D. Hedengren, Ed., vol. 10 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 16, pp. 377–390.

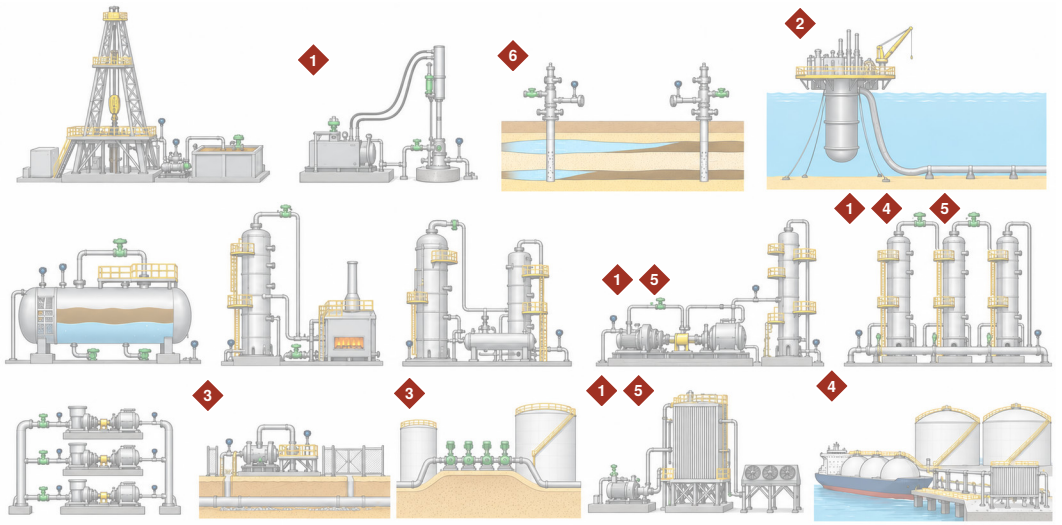
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v10-oil-gas/

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

CHAPTER 16

Summary and Outlook



Open gaps 1 drifting inferences 2 bounded minimum 3 leak or meter 4 chain coordination 5 models without the limit 6 payoff in months

Notice that each number stands beside every plant where a chapter met that gap, so the first gap appears four times and the NGL train carries three. Illustration, not to scale.

IN PLAIN TERMS What Fourteen Oil and Gas Plants Taught

The main idea. Fourteen chapters followed oil and gas from a well being drilled to a terminal that sends gas on to its customers. The same few control ideas came back in plant after plant, such as one valve that answers to whichever limit is closest.

The challenge. The thing that must stay inside its limit often cannot be seen directly. It is worked out from other readings or reported late, so a controller can hold a number that looks right while the real value drifts away.

Why it matters. A reading that looks right but is wrong can let a pipeline plug with ice-like solids, hide a leak or send out gas that is too wet.

Fourteen chapters followed oil and gas from a well being drilled to a terminal sending gas into a grid. Table 1.2 listed, for each plant, the limit that binds, how the controller knows it and the structures the chapter uses. This chapter adds what worked, what failed and the measured result, and reads the comparison for the two questions of the volume, which techniques recur from well to terminal and where measurement, not actuation, sets the limit. 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 where the fluid leaves treating for the liquids recovery plant. The second column condenses the first two columns of Table 1.2. The last column gives the

Table 16.1 What Decided Each Result, Chapters 2 to 8

Chapter	Binding limit, how known	What worked	What failed	Measured result
2, drilling rig	Cuttings, vibration and the pressure window. Cuttings answer a weight move minutes late.	A supervisor that climbs the rate gradient above three PI loops and backs off the most violated limit.	The fixed manual point. Limits set inside the starting state.	28.3 ft against 16.3 ft in 900 s, 97 against 61 % of the best feasible rate (Table 2.4). Cuttings over their limit 42 % of the time while it hunts (Section 2.6).
3, rod-pumped well	Fillage and fluid level, computed through the downhole card from a surface load cell.	An optimizer on the fillage limit, 5.25 SPM at a 2.95 m stroke.	9.0 SPM by hand. A pump-off timer. A card computed with a wrong damping.	55.4 bbl in 6 h with no pound damage, against 56.3 and 392.5 damage units by hand and 51.7 on the timer (Table 3.5). Card error 1.33 to 50.1 % (Section 3.6).
4, smart producer	Drawdown, sand and water from a model and a late meter. The flood fronts are never measured.	Four PI loops in external-reset form at a low selector, under a 180-day planner.	The shipped choke at a 51.4 bar mean drop. A 30-day horizon.	Oil 2.33 times the fixed choke's in 180 days, against 1.07 shipped and 2.00 at 30 days (Table 4.8). Sand 0.30 against 3.24 g/s with the sand loop off (Table 4.7).
5, deepwater flowline	Hydrate margin at the weakest point, 3.23 h downstream, estimated from six transmitters and two meters.	Methanol feedforward on the predicted demand, ratioed to water. A 30 min lead before a restart.	Trim on the estimate alone. A ratio on the liquid rate. The MPC on its reduced model.	Turndown margin –4.43 °C on 47.4 m^3 against 1.03 °C on 28.3 m^3 (Table 5.5). In case M the MPC spends 93.3 m^3 with 8.3 °C h exposed, the feedforward 77.8 and none (Table 5.7).
6, separator	Levels between trips, read directly. Oil and water quality from analyzers behind sample delays.	Averaging level tuning. A second hydrocyclone bank. The supervisor with averaging.	Tight tuning on a slug. The supervisor alone. A longer analyzer delay.	Oil to the second stage 317.5 against 219.1 m^3/h. Off-spec water 1.39, 4.84 and 10.77 m^3 at delays of 30, 120 and 480 s (Table 6.7). In-spec oil 69.8 to 98.7 m^3 (Table 6.9).
7, glycol dehydrator	Water in the gas, read 10 min late. A reboiler ceiling and a pump limit.	A ratio on the water load. Purity by heat, then stripping gas. Economic MPC under the energy budget.	Fixed circulation on a gas step. Ratio and trim under a fuel limit, where the trim makes the gas wetter.	Gas-step peak 6.86 against 7.87 lb/MMscf (Table 7.4). Over the day the worst sample is 6.73 at 232 kW on the MPC and 10.47 on ratio and trim (Table 7.5).
8, amine sweetener	Hydrogen sulfide from a sampled analyzer. Rich loading inferred from an acid-gas meter.	Ratio circulation, a loading floor at a high select, antifoam on pressure drop, analyzer trim.	Fixed circulation. The product analyzer as the only guard of the loading.	True gas over 4 ppm at 13.8 min, analyzer at 24.7 min (Table 8.4). The shift passes at 3.78 ppm, loading 0.458 and 5.25 MW, and Fixed stops at 214.5 min (Table 8.5).

measured result with the table or section of the chapter that reports it.

The columns of the two tables make five points. Figure 16.1 draws one pair of results from each of thirteen chapters.

The selector recurs, and its wiring decides. Six chapters share one actuator among limits through a low selector and three open one through a high selector (Section 1.4.8). On the smart producer

Table 16.2 What Decided Each Result, Chapters 9 to 15

Chapter	Binding limit, how known	What worked	What failed	Measured result
9, turboexpander plant	Methane in the liquid, from a delayed chromatograph. Tray temperature as a proxy. Flooding.	Tray setpoint first, then pressure, with a 20 K return margin.	Either change alone. Both setpoints stepped. An MPC whose model has no quality or flooding.	Worst sample C1/C2 1.51 against 3.65 shipped, 11.36 stepped and 3.69 % on the MPC at 101.7 % vapor load (Table 9.5).
10, NGL train	Ethane in propane, measured one column downstream of the cut that sets it.	The analyzer cascaded onto the deethanizer inferential. A cascade retuned in hindsight.	The inferential alone after it drifts. An MPC holding constant what the surge drum still delivers.	Drifted inferential 0.705 against a 1.80 mol% target, cascade 1.751 (Table 10.4). MPC 153 against 513 min over, peak 3.28 against 2.54 mol% (Table 10.6).
11, compressor station	Surge margin of each machine, computed from flow and head on the commissioned map.	Equal surge margin sharing, a sequencer, the anti-surge valve opened at once on a trip.	Equal speed with one machine fouled. A trip left to the anti-surge PI.	Recycle 15.01 against 0 t in an hour. Surge 0.75 s after a trip (Table 11.4). Over six hours 30.74 against 97.75 t at equal speed, and 31.97 on the MPC (Table 11.5).
12, gas line	Delivery floor. Line pack inferred from three transmitters. A demand forecast.	A 30 min forecast feedforward at a lower setpoint. An afternoon pack for a peak over capacity.	Feedback alone at 48 bar. A 1 % leak threshold against a biased meter.	52 min under the floor on feedback, none with the forecast. A 3 % leak found in 34 min at a 2 % threshold, a 1.5 % meter error alarmed in 34 min at 1 % (Table 12.4).
13, products line	Suction and discharge at four stations, a hill near the terminal, batch cuts at a densitometer.	Low selectors in external-reset form on every pump. Back pressure by MPC on an inferred hill.	Discharge control without selectors. A tight balance across a batch change.	Without selectors a cut at 10 min trips every station by 20.2 min, with them a trip settles at 781 m^3/h (Section 13.5). Over the day 17 false alarms shipped and none suggested (Table 13.5).
14, LNG train	Minimum approach, computed from an analysis refreshed every 12 min. Driver power.	A suction pressure loop under the product temperature. A low selector on the feed.	The computed approach after a relief. The MPC on the feed cut and in peak power.	Estimated minimum 6.5 against 4.2 °C true (Section 14.6). MPC 18.6 against 85.1 min over the driver, peak 2.78 against 1.31 MW, feed cut 18.1 against 6.0 min off spec (Table 14.8).
15, LNG terminal	Tank pressure below the flare. Recondenser capacity from subcooling.	A preview from the published nomination with staging at 95 and 60 %. An MPC on the same schedule.	The loops alone before a low night. A constant low pressure that flashes cargo.	A 30 t/h night fills the recondenser in 91.4 min and flares at 4.69 h (Table 15.4). Over the day the preview uses 30.24 MWh and the MPC 29.75, and a constant pressure boils off 271.5 t (Table 15.5).

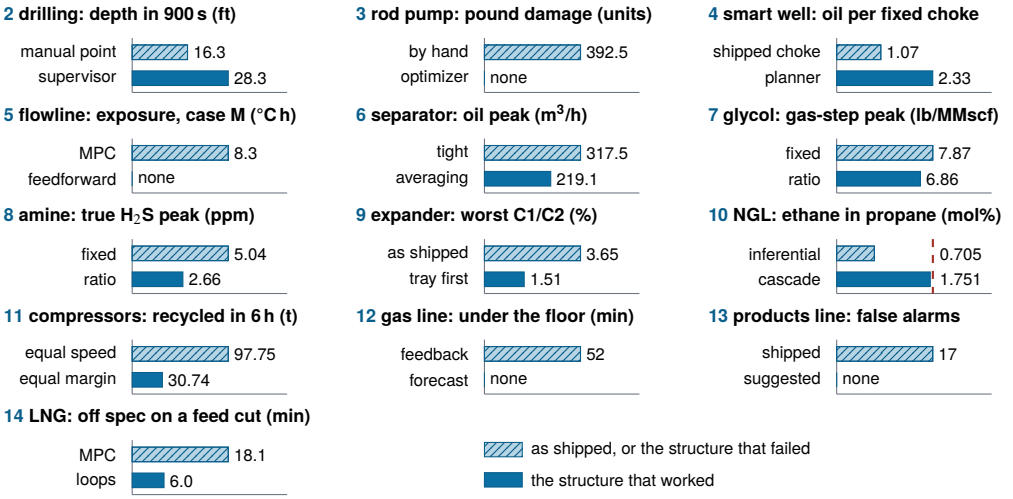


Figure 16.1 One pair of results from each of thirteen chapters, with the numbers as Tables 16.1 and 16.2 print them (the sweetener pair as this section quotes it). Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2 and 4, and the NGL panel marks its 1.80 mol% target with a dashed line. Table 16.2 gives the LNG terminal no like pair. Notice that in four panels the structure that worked removed the harm entirely.

the sand loop held 0.30 g/s where the same rate target without it passed 3.24 g/s. On the products line the selectors turned a station trip into a steady lower flow where discharge control alone tripped every station. On the compressor station the valve of a tripped unit had to open at once, because its anti-surge PI let it surge in 0.75 s. All of this rests on the condition of Figure 1.6. The same two loops with the same tunings held a drawdown limit when both reset terms followed the selected output, and broke it by 1.2 bar for 21.8 min when the waiting loop had wound up.

An inventory decides where a disturbance goes, not whether it arrives. The separator stored a 15 m^3 slug under any tuning, and the tuning decided only whether 317.5 m^3/h or 219.1 m^3/h reached the second stage. The gas line took 12 min to pass a speed step to the city gate, and a peak above its capacity could be met only by packing it beforehand, for about $1,400. The terminal's recondenser filled on a low night whatever the compressors were asked to do. In each case the handle that worked was set before the disturbance: the averaging tuning before the slug, the pack before the peak, a lower tank pressure before the night.

A measured load ahead of a late quality beats a faster loop on the quality. A ratio on the glycol pump held a gas step at 6.86 lb/MMscf where fixed circulation reached 7.87. A ratio on the amine circulation held the true hydrogen sulfide at 2.66 ppm where fixed circulation reached 5.04 ppm. Methanol sent on the predicted demand spent less and kept the margin positive, and the gas line's forecast removed 52 min under the floor. The load must be the one the limit depends on. A methanol ratio on the liquid rate let a water breakthrough thin the inhibitor at the riser base to 9.0 wt% and the margin to −1.73 °C, while a ratio on the water held 1.54 °C (Section 5.6). The glycol ratio keeps the share of water removed, so a warmer inlet passed through it, and the amine ratio reads the last laboratory assay, so a sourer well took the loading past its cap.

Where measurement is the limit, the loop holds the measurement. The three cases of Section 1.4.9 all appeared. A late analyzer turns its delay into product. The amine analyzer reported 10.9 min after the gas, the glycol analyzer 12 min after the water crossed its limit, and on the separator

a longer analyzer delay sent 1.39, 4.84 and 10.77 m^3 of off-spec water overboard. An inferred constraint is held exactly while the truth drifts. The rod pump's card error rose from 1.33 to 50.1 % with a wrong damping, the turboexpander held its tray temperature while the true product reached 3.33 % C1/C2 (Section 9.5), and the subsea estimator with its corrections off read a margin of 3.16 °C where the truth was 0.63 °C, because a colder sea and more heat loss look the same to two thermometers (Section 5.6). The direction of the error decides its cost. The drifted NGL inferential gave away product, while the stale analysis of the LNG train read the approach wide when it was narrow. An inventory computed from a few transmitters cannot see a small leak. The gas line could not tell a 1.5 % leak from a 1.5 % meter error, and on the products line the 60 of 160 detector settings that stayed quiet through a batch change all missed a 0.2 % leak for an hour (Section 13.5).

Prediction paid where it carried a preview or priced several handles. The glycol MPC passed the energy-limited day that ratio and trim failed, and the terminal's MPC passed at 29.75 MWh. On the gas line the MPC spent 10.4 min under the floor against 54.6 min for the cascade with the same forecast, and 301.6 min with its preview off (Table 12.6). Prediction lost where its model left out the limit that binds. The turboexpander's model had no quality or flooding, the NGL plan held constant a disturbance the surge drum was still delivering, the subsea plan smeared a lean front, and the LNG plan cut the minutes over the driver while its peak over it rose from 1.31 to 2.78 MW. On the compressor station it matched equal margin sharing within 0.07 % of fuel and did not beat it. Volume 9 treats that choice in full.

> ⧖ HISTORY **1927 to 2016: The Automatic Driller, Still Unfinished**
>
> The timeline of Chapter 1 opened with a 1927 patent for a dummy driller and closed with automatic control of downhole pressure by a topside choke, applied offshore in 2007 [1]. Mud-pulse telemetry, which sends downhole readings up as pressure pulses at 0.5 to 3.0 bits a second, remained the most common method in 2016. Wired drill pipe, which carries 57,000 bits a second and can send commands down for closed-loop control, was first used commercially in 2006 [2]. A 2016 review found drilling digitized, visual and automated, yet its degree of automation still not enough for real industry [2]. The other chapters stand at the same point: the limit that binds is often what no instrument reads.

16.2 Choosing a Method

The chapters suggest an order of questions for a new plant in the chain. 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, and is it measured, inferred or late?* Fill a row of Table 1.2 before choosing any structure. The rest of the design follows from that column.
2. *Does one actuator serve several limits?* Use a selector, and make every waiting loop's reset follow the applied output (Chapter 4, Chapter 13). Test the wiring by a disturbance that hands the actuator from one loop to the other.
3. *Is there an inventory between the disturbance and the product?* Decide where the disturbance should go, and tune the level or pack for it (Chapter 6). If the capacity behind the inventory is the limit, only an earlier move helps (Chapter 12, Chapter 15).
4. *Is the load measured before the quality?* Add feedforward on the load the limit depends on, and keep a slow trim from the analyzer (Chapter 7, Chapter 5).
5. *Is the binding variable inferred?* Keep a second source that can disagree with it, such as an analyzer cascade or a second transmitter, and alarm on the disagreement (Chapter 10, Chapter 6). Ask which direction of error is unsafe.

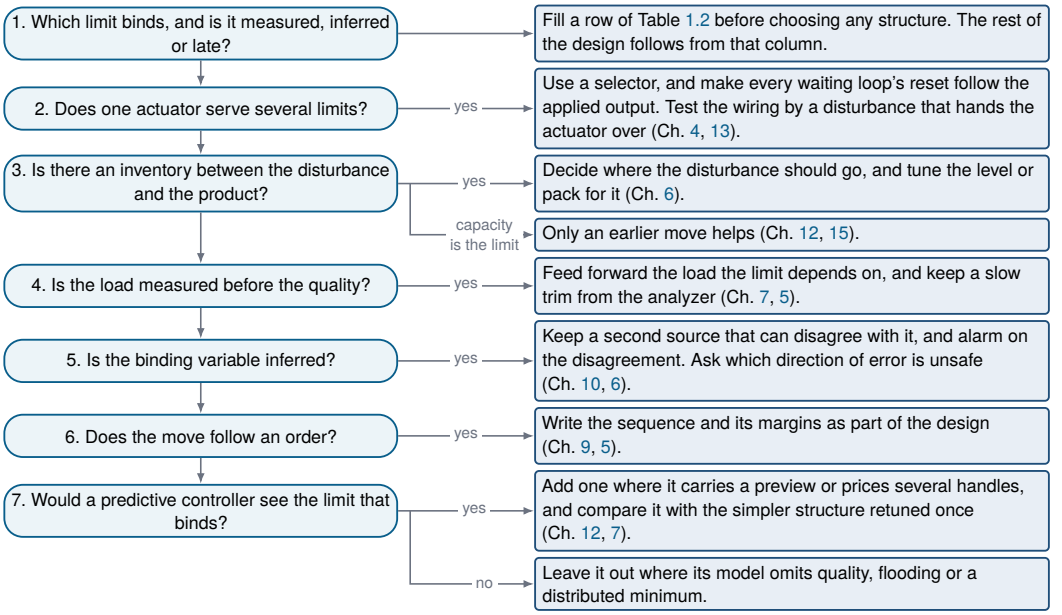


Figure 16.2 The seven questions of this section in order, each with the answer the list gives and the chapters where that answer decided the result. The third and the last question branch. Notice that the first question, which limit binds and how it is known, settles what every later answer acts on.

6. *Does the move follow an order?* Write the sequence and its margins as part of the design. The turboexpander needed both the order and a wider return margin, and the restart needed its methanol lead (Chapter 9, Chapter 5).
7. *Would a predictive controller see the limit that binds?* Add one where it carries a preview or prices several handles, and compare it with the simpler structure retuned once (Chapter 12, Chapter 7). Leave it out where its model omits quality, flooding or a distributed minimum.

16.3 Where Practice Is Heading

Automation reaches the well. Drilling has traditionally been operated by hand. Automatic managed pressure drilling, with downhole pressure held by topside choking, was applied offshore in 2007, and its requirements were stated for normal operations such as connections and for failures such as gas kicks and a blocked choke [1]. A 2013 review found broad interest in drilling-systems automation but no agreement on what it entails, and named well complexity, data overload and the loss of skilled people among its drivers [3]. Real-time production optimization has been surveyed as a chain of data acquisition, model updating and planning that ends in choke, valve, compressor and pump settings [4]. Optimization of gas transmission has its own body of work [5]. In the editor's view the supervisors of Chapter 2, Chapter 3 and Chapter 4 are the form this takes on one well, and the work will be in coordinating many of them.

Measurement moves into models. Virtual flow metering computes multiphase rates from pressures and temperatures instead of a hardware meter. A 2020 review found it able to serve as a standalone meter or a backup, and asked for automatic calibration as fluid properties change, for methods that state the uncertainty of the estimate, and for state estimation that stays accurate in transients [6]. Slug control has been designed with and without pressure sensors upstream of the riser [7]. Digital twins of oil and gas assets were judged in 2020 to be at an early stage, limited to isolated applications [8].

Hydrate practice has been described as moving from avoidance toward the management of risk [9], which is what a computed margin and an exposure integral, as in Chapter 5, make possible. In the editor's view every limit in Table 1.2 that is now inferred will be inferred by a larger model, and the chapters show why each such model needs a second source that can contradict it.

Integrity and emissions are watched more closely. Pipeline leak detection methods have been reviewed for their advantages, limitations and performance, together with the standards that govern them [10]. For fugitive methane, a review of six classes of technology, from handheld instruments to satellites, found minimum detection limits from under 1 g/h to 7.1×10^6 g/h, found the uncertainties poorly constrained, and proposed screening sites with a wide survey before confirming the source up close [11]. In the editor's view the trade measured in Chapter 12 and Chapter 13, a threshold low enough to find a small leak and high enough to ignore a meter, will be met more often as these programs grow.

Fiber turns lines and wells into thermometers. A 2022 review calls Raman distributed fiber sensing of temperature mature, widely applied to finding pipeline leaks and used to check the integrity of oil and gas wells [12]. The same review states its limits: the sensing distance and the spatial resolution cannot both be had, and a cleaner signal costs measurement time. In the editor's view a measured temperature profile is the second source that gaps 2 and 3 of Section 16.5 ask for, on a flowline read through two thermometers and on lines that look for a leak in an inventory computed from a few transmitters.

Electric drives join the power system. New LNG plants are increasingly driven by electricity under a site power management system, which in an electrical contingency typically disconnects loads until the consumption meets the power that remains [13]. A 2022 analysis found that a mixed fluid cascade plant could instead keep running on 30 % less power or more [13]. In the editor's view the driver limit of Chapter 14 and the load sharing of Chapter 11 will more often be set by the power system than by the machine.

Optimizers learn the machines they run. Online feedback optimization steers a plant toward its optimum from measurements and a linear model, without solving a nonlinear optimization problem online. A 2023 study added a model that learns how each compressor departs from its map. On a computed station sharing load among parallel machines, imperfect maps cost 5 % more power and the learning cut the excess to 0.8 %, with a method meant to run at the edge with moderate modeling effort [14]. In the editor's view this answers the drifting maps of Chapter 11, so far only on a model. 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 [21]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

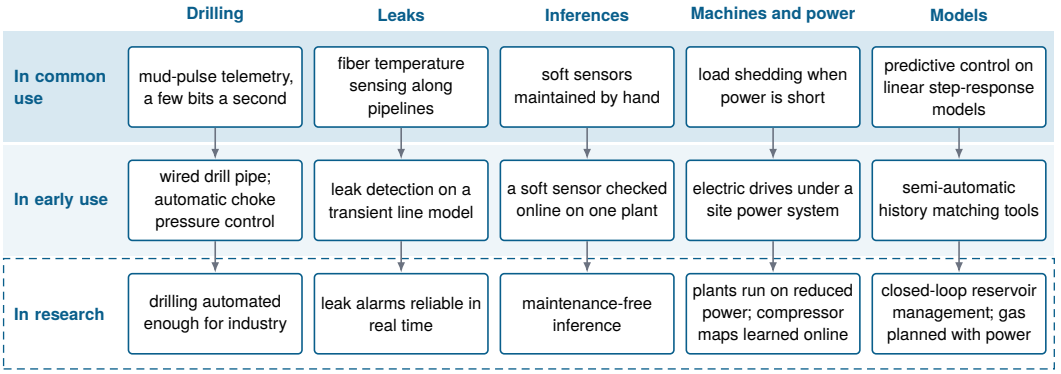


Figure 16.3 The developments of this section in five threads. The top lane holds what the cited source calls the most common, typical or widely applied practice, the middle lane what is in service at some sites or sold as a tool, and the bottom lane what is reported from simulation or as a need. The placing is the editor's reading of the sources, and no box carries a date. Sources by thread: drilling [1, 2], leaks [12, 15], inferences [16, 17], machines and power [13, 14] and models [18–20]. Notice that each bottom box asks that a measurement or a model be trusted to act on in real time, with no person checking it first.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require

an assessment that includes someone independent of the design team [22]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

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 five of the fourteen reports are self-reviews by the agent that wrote the paper, and six papers had a further referee pass on predictive control added later. 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 fourteen logs written beside the papers hold 105 entries with a severity. Of these, 28 are wrong behavior or output of an app, 48 are documentation or content, 20 are cosmetic or display and 9 are other kinds, such as a test harness or a load path. The counts understate the work. Three of the logs hold one entry or none, and eight apps keep a second log of their build in their own folders, which is not counted here. Five cases show the errors that matter for a volume about what a controller can know.

Controllers that knew what no instrument tells. Five apps logged a controller that read values a real plant would not give it. The subsea estimator and its dose conversions read the run's own parameters. At identical measurements an edited heat transfer coefficient moved the estimated minimum margin from 4.575 to 3.660 °C, and handing only the controllers an edited copy moved 49 of 100 six-hour runs. The compressor station's MPC built its model from the learner's parameters, and 16 of 17 edits moved its plan. The products line's MPC read the wall roughness of the plant, so its inferred hill moved by up to 2.63 bar with nothing measured changing. The dehydrator and the sweetener started their controllers from true plant states. No shipped run moved when these were fixed. Each was caught by the same probe, a hidden value changed at identical measurements to see whether the controller moves. In a volume where the hydrate margin and the hill are the limits, a controller that reads them from the plant's parameters looks better than any real one could.

Tests of that boundary that passed bad code. After the first truth tests were written, reviewers wrote deliberately wrong controllers, called mutants, to see whether the tests would fail. On the separator the two original tests caught 1 of 8 mutants. On the NGL train 9 of 12 mutants had passed. On the LNG train a mutant that reused two shipped analyzer statements to hand the MPC the true refrigerant blend on every scan passed the test. Each set of tests was then strengthened. A test that has never failed says nothing about the code it guards.

A limit checked where the instrument was not. The products line's margin tile reported 2.15 bar to the maximum allowable pressure while the terminal valve inlet peaked at 89.7 bar g, 0.26 bar under the limit. The check scanned the segment centers and the station discharges and never the terminal, the lowest point of the last leg. The referee found it by reading the data behind a figure. Once fixed, only the surge case moved, its least margin from 2.1 to 0.3 bar.

Measurement and truth confused in the text. The subsea draft said the loops see only the estimator, yet the flow loop and the pressure override read their transmitters directly. Its screenshot

caption said the flow was held by the pump limit, and a run showed a lower flow under a setpoint computed from a measured water cut of 18.15 % against a true 20 %. Three referee reports, on the separator, the smart producer and the flowline, found numbers inherited from the app's own guide that the engine did not reproduce. One belonged to another instant of the same run. Each was caught by measuring again instead of copying.

The book round. The chapters were written after a repair round that applied the book's earlier defect list to apps and papers, and each chapter agent first confirmed that its paper had not changed since conversion. 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 thirteen report that every rerun number reproduced. The NGL record found three printed values that differ from the run, among them a steam demand of 5235 kmol/h where the run asks about 5406, and left them for the editor. The book round's defect list holds thirteen entries for this volume. The one that touches control was found on the dehydrator. A pump capacity edited in the panel never reached the controls, so at 30 L/min the plant delivered 29.8 while the controller still reported 50 and no saturation, and the trim could wind up against a pump at its limit. The agent writing Chapter 1 first chose an averaging tuning whose outflow peaked at 273 m^3/h against a 240 m^3/h limit, and caught it in its calculation script before writing.

Every error above was caught by a check that could fail: a hidden value moved at fixed measurements, a test run against code wrong on purpose, a limit checked at the plant's worst point, a number measured instead of copied. Four working rules follow. Name the instrument behind every value a controller reads. Test each test with code that should fail it. Check a limit where the plant is worst, not where the transmitter is. Let a person decide what changes.

16.5 Gaps in Oil and Gas Control

Six gaps belong to this volume's topic, 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 plants in the picture that opens this chapter.

1. *Inferences that report their own drift.* The rod pump's card, the turboexpander's tray temperature, the NGL inferential and the LNG train's computed approach each held a loop on a number that was wrong for a time. What is missing is an inference that carries its own uncertainty and says when the data no longer support it, the request the flow-metering review makes for its field. A review of soft sensors found that most have no automated mechanism for their maintenance [16]. The Open challenge box after this list takes the gap further.

2. *A distributed minimum estimated with bounds.* The hydrate margin of Chapter 5 is a minimum over 6.5 km known through two thermometers that cannot tell a colder sea from worse insulation. A dosing law needs a lower bound on that minimum, not a best estimate. A 2021 history of flow assurance finds industry trending toward safe operation inside the hydrate domain [23]. A fiber can measure the profile, but its resolution, of the order of meters, cannot be reconciled with a long reach [12]. The gap is hard because the weakest point moves while the instruments stay at the ends. Progress would be a dose set from a stated lower bound and checked against a measured profile.

3. *Leak detection that tells a leak from a meter.* On both lines of Chapter 12 and Chapter 13 the threshold that ignored a biased meter or a batch change missed a small leak. Transient models, temperature and product properties in the balance are the known routes, and none removes the

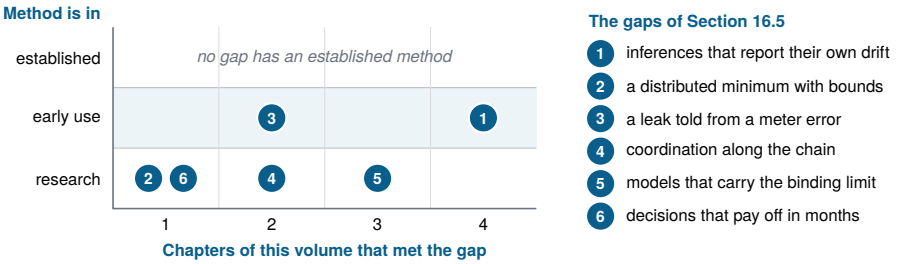


Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: four for gap 1 (Chapters 3, 9, 10 and 14), one for gap 2 (5), two for gap 3 (12 and 13), two for gap 4 (10 and 15), three for gap 5 (9, 10 and 14) and one for gap 6 (4). Up is the maturity of the method available, which is the editor's reading of the sources cited in each item and not a measurement. Notice that only the two gaps of measurement have a method in use anywhere.

trade. A 2019 review finds most methods that work from the fluid's own measurements prone to false alarms, and gaps remaining before reliable real-time detection is achieved [15]. The gap is hard because a leak and a meter error move the same balance the same way. Progress would be a threshold set from a stated meter error that still alarms a leak smaller than that error.

4. *Coordination along the chain.* Each plant in this volume was controlled alone, with its neighbors as disturbances. The NGL header trip showed a decision of which product to sacrifice that no layer made, and the terminal's flare was decided by the send-out the market asked for. Methods that coordinate selectors and supervisors across plants owned by different operators are thin. Gas lines now feed generators that balance wind and solar power, and a 2015 study found the coupling invalidates many traditional approaches to operation [20]. The gap is hard because each plant answers to its own operator and contract. Progress would be a layer that sees which plant's limit binds along the chain and moves the cut there.

5. *Prediction models that carry the binding limit.* The predictive controllers of Chapter 9, Chapter 10 and Chapter 14 lost because their models left out quality, flooding, a surge drum or a peak. A routine way to tell, before commissioning, whether a reduced model contains the limit that will bind is not available. The 2003 survey of industrial predictive control listed a lack of methods to verify whether an identified model is adequate for control [24], and in 2023 new applications were still built mainly on linear models from step responses [18]. The gap is hard because the limit that binds in service may never move in the plant test. Progress would be a check that each limit expected to bind is in the model.

6. *Decisions whose payoff arrives in months.* The waterflood planner of Chapter 4 forecast 13.9 % high from its model's errors, and a fixed policy chosen by hand came within 0.8 % of it. Closed-loop reservoir management needs a way to value a plan against a simple policy when the truth is known only years later. A 2015 review found most of its studies academic and far from large-scale field application, and automatic closed-loop management not yet achieved there [19].

OPEN CHALLENGE **An Inference That Knows When It Is Wrong**

Gap 1 is the one most chapters met. A 2013 review warns that product quality deteriorates if a soft sensor's estimates are blindly believed and used in a control system [17]. Comparing the estimate with each analysis judges it only when an analysis exists. A check between analyses meets a harder problem: a model that must follow changes in the process and the feed sees its new data as faults. One scheme that checked its soft sensors online ran on an ethylene plant and passed only reliable estimates to a predictive controller [17]. Progress would be such a check in routine use, one that tells a drifting inference from a changing plant.

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 [25]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [22]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [26]. 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 [27]. 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 [28]. 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 controller holds the reading it is given exactly, even while the true value drifts away.
- A selector shares one actuator among limits, and its wiring decides whether the limit holds.
- An inventory decides where a disturbance goes, not whether it arrives.
- A measured load ahead of a late quality beats a faster loop on the quality.

- Prediction paid where it carried a preview or priced several handles, and lost where its model left out the limit that binds.

References

[1] J.-M. Godhavn, "Control requirements for automatic managed pressure drilling system," *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[2] T. Ma, P. Chen, and J. Zhao, "Overview on vertical and directional drilling technologies for the exploration and exploitation of deep petroleum resources," *Geomechanics and Geophysics for Geo-Energy and Geo-Resources*, vol. 2, no. 4, pp. 365–395, 2016.

[3] J. D. Macpherson, J. P. de Wardt, F. Florence, C. D. Chapman, M. Zamora, M. L. Laing, and F. P. Iversen, "Drilling-systems automation: Current state, initiatives, and potential impact," *SPE Drilling & Completion*, vol. 28, no. 4, pp. 296–308, 2013.

[4] H. P. Bieker, O. Slupphaug, and T. A. Johansen, "Real-time production optimization of oil and gas production systems: A technology survey," *SPE Production & Operations*, vol. 22, no. 4, pp. 382–391, 2007.

[5] R. Z. Ríos-Mercado and C. Borraz-Sánchez, "Optimization problems in natural gas transportation systems: A state-of-the-art review," *Applied Energy*, vol. 147, pp. 536–555, 2015.

[6] T. Bikmukhametov and J. Jäschke, "First principles and machine learning virtual flow metering: A literature review," *Journal of Petroleum Science and Engineering*, vol. 184, p. 106487, 2020.

[7] F. Di Meglio, N. Petit, V. Alstad, and G.-O. Kaasa, "Stabilization of slugging in oil production facilities with or without upstream pressure sensors," *Journal of Process Control*, vol. 22, no. 4, pp. 809–822, 2012.

[8] T. R. Wanasinghe, L. Wroblewski, B. K. Petersen, R. G. Gosine, L. A. James, O. De Silva, G. K. I. Mann, and P. J. Warrian, "Digital twin for the oil and gas industry: Overview, research trends, opportunities, and challenges," *IEEE Access*, vol. 8, pp. 104 175–104 197, 2020.

[9] E. D. Sloan, "A changing hydrate paradigm—from apprehension to avoidance to risk management," *Fluid Phase Equilibria*, vol. 228–229, pp. 67–74, 2005.

[10] H. Lu, T. Iseley, S. Behbahani, and L. Fu, "Leakage detection techniques for oil and gas pipelines: State-of-the-art," *Tunnelling and Underground Space Technology*, vol. 98, p. 103249, 2020.

[11] T. A. Fox, T. E. Barchyn, D. Risk, A. P. Ravikumar, and C. H. Hugenholtz, "A review of close-range and screening technologies for mitigating fugitive methane emissions in upstream oil and gas," *Environmental Research Letters*, vol. 14, no. 5, p. 053002, 2019.

[12] J. Li and M. Zhang, "Physics and applications of Raman distributed optical fiber sensing," *Light: Science & Applications*, vol. 11, p. 128, 2022.

[13] I. Y. Xuan, C. Skourup, J. B. Jensen, T. Haugen, and N. F. Thornhill, "Flexible operation of a mixed fluid cascade LNG plant for electrical power management," *Energy*, vol. 250, p. 123794, 2022.

[14] M. Zagorowska, M. Degner, L. Ortmann, A. Ahmed, S. Bolognani, E. A. del Rio Chanona, and M. Mercangöz, "Online feedback optimization of compressor stations with model adaptation using Gaussian process regression," *Journal of Process Control*, vol. 121, pp. 119–133, 2023.

[15] M. A. Adegboye, W.-K. Fung, and A. Karnik, "Recent advances in pipeline monitoring and oil leakage detection technologies: Principles and approaches," *Sensors*, vol. 19, no. 11, p. 2548, 2019.

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

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

[19] J. Hou, K. Zhou, X.-S. Zhang, X.-D. Kang, and H. Xie, "A review of closed-loop reservoir management," *Petroleum Science*, vol. 12, no. 1, pp. 114–128, 2015.

[20] A. Zlotnik, M. Chertkov, and S. Backhaus, "Optimal control of transient flow in natural gas networks," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 4563–4570.

[21] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[22] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[23] E. D. Sloan, "Hydrocarbon hydrate flow assurance history as a guide to a conceptual model," *Molecules*, vol. 26, no. 15, p. 4476, 2021.

[24] 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.

[25] 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.

[26] 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.

[27] 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.

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

16.6 Exercises

16.1 Discussion · 7 min

A new well is to share its choke between a rate loop and a drawdown limit. A colleague proposes two PI controllers into a low selector, each with its own anti-windup against its own output limits. Using Figure 1.6, Chapter 13 and Chapter 11, say what to expect and what to check before the scheme is commissioned.

16.2 Discussion · 8 min

Section 1.4.9 names three ways in which measurement limits control: a late analyzer, an inferred constraint and an inventory computed from a few transmitters. For Chapter 8, Chapter 10 and Chapter 12, say which applies, what it cost in the chapter, and what second source could contradict the measurement the loop uses.

16.3 Discussion · 7 min

A ratio controller is usually tuned at the design point, where it protects the product. Using case 3 of Chapter 5 and the warm-up of case 3 in Chapter 7, explain how a ratio that is right at design can protect the wrong quantity, and how the basis of a ratio should be chosen.

16.4 Discussion · 8 min

A vendor proposes an MPC for the mode switch of the turboexpander plant of Chapter 9, and cites the economic MPC of Chapter 7 as evidence that prediction pays in gas processing. Using both chapters and Section 16.2, say what you would ask before accepting it.

16.5 Discussion · 7 min

An agent delivers a new estimator for the hill near the terminal of Chapter 13 and reports that it tracks the true hill within 0.1 bar on every shipped case. Using Section 16.4.1, say what checks you would make before the estimator goes into the back-pressure controller.

16.6 Discussion · 8 min

A field operator plans to replace the physical multiphase meters on its smart producers with a virtual flow meter computed from pressures and temperatures, and to run the water-capacity loop of Chapter 4 on it. Using Section 16.3 and the chapters on inferred limits, discuss what the operator gains, what it risks and what it should keep.