INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 15 Control in Mining and Mineral Processing

CHAPTER 12

Summary and Outlook

Pages 259–273

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Mining and Mineral Processing*, John D. Hedengren, Ed., vol. 15 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 259–273.

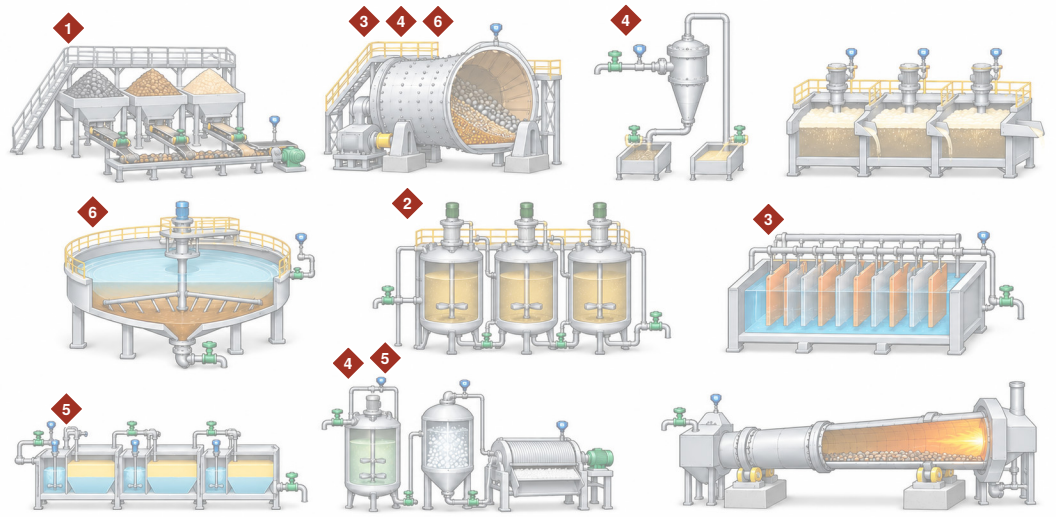
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v15-mining-minerals/

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

CHAPTER 12

Summary and Outlook



Open gaps 1 arriving ore 2 blind estimates 3 inferred limits 4 hidden drift 5 inventories 6 protection

Notice that the mill carries three gaps, and that no gap marks the flotation cells or the kiln. Illustration, not to scale.

IN PLAIN TERMS **What Ten Mineral Plants Taught**

The main idea. Each chapter of this book ran one part of a mineral plant in a simulator, from the ore belt to the kiln. In every one the rock changed in ways that no instrument could see, and the test that judged the product came back late or never.

The challenge. A loop can hold its own reading steady while the plant drifts somewhere else. A late test also limits how hard a loop may correct.

Why it matters. The plants that did best measured what arrived and acted on it early, chose their targets with care, and let protection win when a limit was near.

Ten chapters followed ore from a belt scale to a cathode, a hydroxide product and a clinker. Each met the two problems of the volume's title. The ore changed without being measured for what mattered, and the quality that judged the plant arrived late or not at all. Table 1.1 of Chapter 1 listed, for each chapter, how the ore or feed varies, what is measured late or not at all, and the structure that handles it. This chapter adds what each structure achieved and where it failed, draws a short guide for choosing a method, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

12.1 What the Chapters Showed

Tables 12.1 and 12.2 compare the chapters. The second column condenses Table 1.1, how the ore varies and what is late or never measured. The other columns name a structure that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

The columns make five points, and Figure 12.1 draws one pair of results from each row.

Table 12.1 How Variable Ore and Late Assays Were Handled, Chapters 2 to 6

Chapter	Ore varies, and what is late	What worked	What failed	Measured result
2, ore feed	Calibrations, belt speed, an empty silo. The scale weighs the total and never the blend.	Speeds set through the calibrations. A bump test on a drifted feeder. Holding back ore that would overload the belt.	A master PI that chases transport dips. Every structure on a wrong calibration.	On a speed change 65.4 % on spec for the master PI against 78.5 for ratio feedforward and the MPC. A drifted feeder 37.6 to 40.0 % (Table 2.4). Shift over the rating 172 against 0 s (Table 2.5).
3, SAG mill	Hardness, ball charge, liner wear. The toe is inferred, the size 2 min late.	The cautious impact loop with a slower load loop that replaces lost steel without overshoot.	The recommended impact tuning on a noisy shift. Toe tracking on an inferred toe.	Two points of ball filling took the margin from +0.22 to −3.27° and the product from 2367 to 4136 µm with every loop on setpoint (Table 3.7). Shift toe error 2.42, 3.98 and 5.45° (Table 3.9).
4, hydro-cyclones	Coarseness, grindability, pump wear. Size 2 min late, circulating load never.	A pressure cascade under the size loop against pump wear. Commissioning from the level loop outward.	The same cascade against worn liners or a ninth cyclone. A tighter level loop.	Fineness IAE 154.9 to 3.2 for impeller wear, 25.1 to 73.1 for liners (Table 4.4). A level loop 7.8 times tighter raised the fineness IAE 7.6 % (Table 4.6).
5, flotation cell	Feed grade, throughput, blower. Grades 2.0 and 3.5 min late.	The grade loop alone at an economic setpoint, 0.55 m and 8.4 % copper, over a fast level loop.	An aggressive grade loop. A second quality loop on the collector. Tight regulation at the design depth.	Aggressive loop $3305 /h against 7270 at the same 76.3 % recovery (Table 5.3). Graded run 9890 suggested, 8580 at the as-found setpoints, 7072 with both quality loops (Table 5.4).
6, thickener	Feed solids, settleability. Settleability never, density, torque and bed in minutes.	A high select of density, torque and bed demands with external reset, and a feasible density setpoint.	A low select, which blocks the protection. The loaded setpoint and slow tuning.	Low select torque peak 109.9 % (Table 6.6). Graded shift torque 104.7 against 65.9 %, bed 3.280 against 2.567 m (Table 6.8).

Every loop on setpoint, and the plant somewhere else. The most repeated result of the volume is a plant that moved while its loops held. In Chapter 3 two points of ball filling left the impact angle, the load and the density on setpoint while the media landed inside the toe and the product coarsened by 75 %. In Chapter 10 an electrode reading 0.70 low raised the reagent 24 % while every number on the daily report improved. In Chapter 8 a diffusivity 20 % low left the temperature and copper loops with the integrated errors of a correct run while the cell shorted on nodules. In Chapter 4 the level and pressure loops held while the circulating load rose from 288 to 386 % (Section 4.5). A loop measures its own variable. In each case the variable that decided the result was one that no instrument read.

What the soft sensor cannot read arrives with the assay. A soft sensor removes the assay's dead time only for the disturbances among its inputs. In Chapter 7 the soft-sensor loop moved its cyanide setpoint 0.26 h after a measured feed grade step, and the assay loop 10.26 h after it. A lower leachability, which no input carries, first moved the setpoint at 20.1 h under either loop (Section 7.6.2). The bias update that joins model and assay is then a filter on stale, noisy data. A gain of one raised the actuator travel and the peak, and a gain of zero let the tailings sit above the limit for 172 of 200 hours. The same holds without a soft sensor. With its feedforward on, Chapter 9 passed its shift with

Table 12.2 How Variable Ore and Late Assays Were Handled, Chapters 7 to 11

Chapter	Ore varies, and what is late	What worked	What failed	Measured result
7, gold CIL	Grade, leachability, copper and sulfide parcels. Tailings sampled every 4 h, reported 6 h later.	A fixed cyanide setpoint on the graded week. A bias update gain near a third.	The soft-sensor cascade at its recommended target. No bias update. A bias gain of one.	Hours above 0.40 g/t were 0.0 fixed, 21.6 for the cascade and 172.0 with no update (Table 7.3). A gain of one raised the cyanide travel 63 % over 0.15 (Table 7.2).
8, copper refining cell	Purification upsets, passivation, dirty electrolyte. The limiting current never, copper 2 h late.	Utilization held on the inferred limit with margin, a voltage override, more leveler for a dirty electrolyte.	Fixed current. A diffusivity 20 % low in the estimate. Running for the most copper.	Believed 58.99, true 73.74 %, and the cell shorts (Table 8.3). At 66 % it plates 17 % more copper than at 58.5 and earns 24 % less (Table 8.1).
9, rare earth cascade	Feed grade and concentration. Purity and recovery from a laboratory 1 to 5 h old.	Feedforward from the feed analysis with a purity loop detuned to the plant's day, 0.05 per point and 24 h.	The purity loop at 0.5 per point and 6 h. The crossed pairing. Either part of the suggested structure alone.	Shift passed on 5 of 5 seeds at $2.93 per kilogram, the shipped loop on none, feedforward alone at $3.00 to 3.01 (Table 9.5). MPC 17.78 h below specification (Table 9.6).
10, lithium refinery	Ore grade, fuel header, heating value. Product strength read two units downstream.	A fuel flow cascade under the kiln temperature against header pressure. Each loop tuned for its own clause.	The cascade against a heating value loss. A pH loop that trusts a biased electrode.	Header loss IAE 526.3 against 5.7, heating value 158.4 under both (Table 10.4). An electrode 0.70 low raised the yield from 87.8 to 91.2 % and the reagent 24 % (Section 10.6).
11, cement kiln	Raw meal, coal heating value, false air. Free lime 15.9 to 25.9 min late.	Air ratioed to the coal command with oxygen trim. The free lime loop as a slow trim. A ratio floor and cross-limit.	Oxygen feedback alone. A hot free lime tuning against 20 min of dead time.	Carbon monoxide peak 611 against 20 ppm on one firing move (Table 11.4). Graded run with a reducing flame for 88.7 against 0.0 min, at 1600 against 1721 kJ/kg (Table 11.6).

every purity gain from 0.05 to 0.3 at integral times of 12 and 24 h, and failed with every gain at 6 h. A loop on a late measurement is tuned to the plant's response and the report interval, not to the size of the error.

Measure what arrives, feed it forward, trim from the assay. The structures that worked moved the actuator from something measured before it reached the product. The ore feed set each feeder through its calibration, the kiln ratioed its air to the coal command, and the rare earth cascade keyed its scrub to the feed analysis. A cascade is the same idea one step inside the plant, and its value depended on the path. The lithium kiln's flow loop cut the integral error of a header pressure loss by a factor of 52 to 93, depending on when the loss landed in the scan, and did nothing for a heating value loss. The hydrocyclone's pressure loop cut the fineness error of impeller wear 49 times and made worn liners about three times worse, because restoring the pressure drew dilution water and thinned the feed. The late assay then trims what the fast path cannot see.

The setpoint decided more than the tuning. Several chapters gained more from where a loop was set than from how it was tuned. The flotation cell regulated at its as-found setpoints with the best

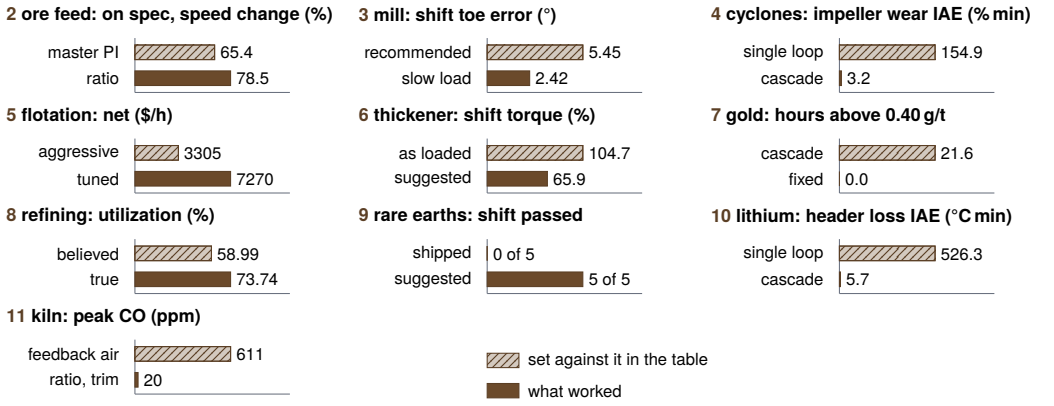


Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 5 and 9. In the panel of Chapter 8 the hatched bar is the utilization the cell believed and the solid bar the true one. Notice that the better result came from a change of structure, measurement or setpoint more often than from tuning alone.

depth tracking of its table and earned \$1309 /h less than at 0.55 m. The refining cell plated the most copper at 66 % of its limiting current and earned the least. The thickener met its graded shift only with a density setpoint the rake could carry. The gold circuit held its tailings limit through the graded week with no outer loop at all, at a fixed 120 mg/L of cyanide, where the recommended cascade spent 21.6 h above it. Each optimum was flat or bounded by a limit that no instrument measured, so choosing it was an engineering decision and not a tuning result.

A faster loop bought nothing the disturbances could use. Tuning for a step response misled on plants whose disturbances are slow and whose measurements are noisy. The mill's recommended impact tuning, 15 times better on a setpoint step, had 3.3 times the integral error of the cautious tuning over the graded shift. The hydrocyclone's level loop, tightened, made the product worse. The kiln's free lime loop, pushed to a gain of 45 °C per percent, cycled with a period of 57 min on 20 min of dead time (Section 11.5.4). The flotation grade loop at $K_c = 12$ cycled on its own analyzer, sent the tailings valve through 2257 % of its range and cost 55 % of the cell's income.

12.2 Choosing a Method

The chapters suggest an order of questions for a circuit that is new to the reader. Each question names the chapter where its answer decided the result, and Figure 12.2 draws the same order as a flowchart.

1. *What in the ore varies, and what reads it, and when?* Fill in a row of Table 1.1 before choosing a structure. A scale that reads the total cannot see the blend (Chapter 2).
2. *Does the disturbance enter before a fast measurement?* Then a ratio, a feedforward or a cascade removes it (Chapter 11, Chapter 10). Check what the correction moves on its way, as the liners of Chapter 4 show.
3. *How old is the quality measurement when it is used?* Count the sampling and the report (Section 1.4.5). Set the integral time to the plant's response and the report interval (Chapter 9), and design the bias update as a filter (Chapter 7).
4. *Is the controlled variable inferred?* Name what the inference cannot see, and back the setpoint off by the error the estimate may carry (Equation (1.9), Chapter 8, Chapter 3).

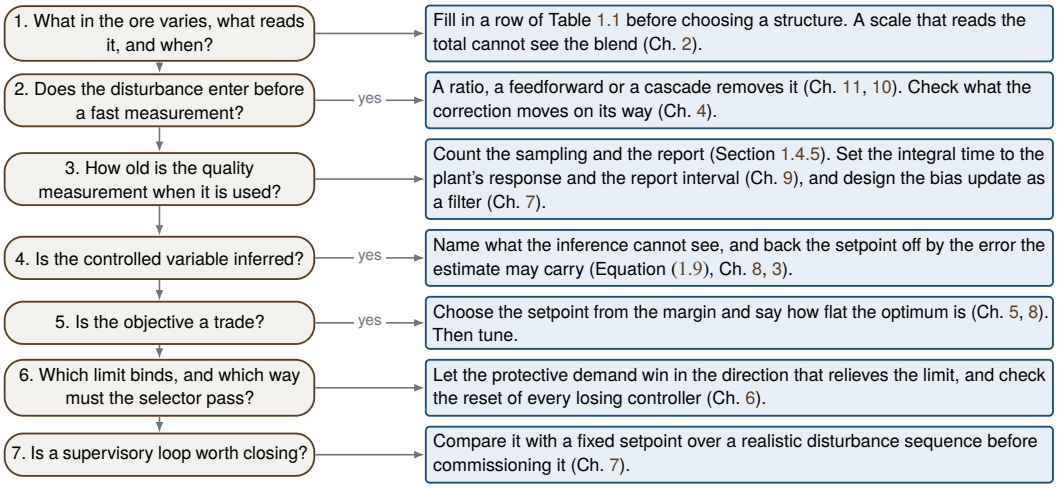


Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions ask what varies, what is measured and when, and are answered before any loop is tuned.

5. *Is the objective a trade?* Choose the setpoint from the margin and say how flat the optimum is (Chapter 5, Chapter 8). Then tune.
6. *Which limit binds, and which way must the selector pass?* Let the protective demand win in the direction that relieves the limit, and check the reset of every losing controller (Chapter 6).
7. *Is a supervisory loop worth closing?* Compare it with a fixed setpoint over a realistic disturbance sequence before commissioning it (Chapter 7).

A predictive controller paid where it handled a constraint on material already in the plant. On the ore belt it was the only structure that kept the loading inside its rating through the infeasible shift under a temporary recipe, against 172 and 148 s over for the master PI and ratio feedforward (Table 2.5). On the kiln it held the flame at an excess air of 1.043 or more through false air where the ratio with trim fell to 0.948 (Table 11.7). Elsewhere most of its gain belonged to an operating point or to information the loops could be given. The flotation MPC's own optimum, entered as fixed stations, came within \$11 /h of it (Table 5.5). The mill's loops given the MPC's breakage estimate as a feedforward came within 0.4 % of its ore (Table 3.11). The kiln's loops given its burning zone target, a ratio floor and a cross-limit reached 19.60 against its 20.93 % min of free lime error. On the gold circuit a setpoint optimizer recovered all but \$1.9 per hour of the plan's gain, and the plan ran over the tank-five discharge limit for 28.3 h (Table 7.4). On the rare earth shift the MPC spent 17.78 h below specification where the compensated PI spent none. Before reaching for prediction, try the feedforward, the setpoint and the selector. Volume 9 treats that choice in full.

12.3 Where Practice Is Heading

Knowing the ore before the plant does. Geometallurgy has become a tool for reducing production risk and improving economic performance, and most of its definitions describe it as a bridge between geology and mineral processing [1]. Sensor-based ore sorting separates coarse ore before fine comminution. It can reject waste, divert material to different process lines or make concentrates at several points of a flowsheet, and its review opens by noting that deposit qualities are constantly decreasing [2]. Both move the measurement of the ore ahead of the circuit. In the editor's view this is

the most direct answer to the first problem of this volume. The scale of Chapter 2 weighs a total and never a blend, and every chapter after it learns of a harder or leaner ore from its effect.

> ⧖ HISTORY **2014 to 2022: The Instruments Arrive, the Assay Stays**
>
> The second history box of Chapter 1 followed instruments for the hidden plant from 1973 to 2004. In 2014 a thesis still called a measurement taken with the mill stopped the common way to find a mill's filling, and named acoustic, inductive proximity and conductive sensors as the new technologies [3]. A 2019 study watched the flow inside a hydrocyclone, hidden by its walls and the opaque slurry, through electrical impedance tomography at several hundred frames a second [4]. In 2022 a study found online elemental analyzers widely available yet under-used in mineral processing, where laboratory samples still guided control [5]. In the editor's reading the instruments have arrived, and in most loops of this volume the assay still decides.

Ore properties inferred from operating data. Where the ore is not measured ahead, it is increasingly inferred from the plant. A long short-term memory network trained on a year of operating data from each of two semi-autogenous mills forecast an operational relative hardness with accuracies above 80 % a short time ahead [6]. Chapter 3 gives the caution that goes with such estimates. Its toe, inferred from the bearing pressure with a declared ball charge, fell while the true toe rose, and a setpoint that tracked it passed the graded shift on none of eight noise seeds. In the editor's view an inferred ore property is a measurement with a model inside it, and it needs the statement of what it cannot see that Section 1.4.5 asks of a soft sensor.

More from the assays a plant already has. Froth image analysis has been used to identify operating regimes, to characterize froths and to build inferential sensors for chemical species in the froth, and it has been joined with multivariate statistical methods to monitor industrial platinum group metal flotation plants [7]. Data reconciliation that adds the stoichiometry of the minerals to the material balance estimates balanced mineral compositions, assays and flow rates, and the estimated mineral composition improves the evaluation of process performance [8]. In the editor's view such estimates will sit increasingly inside the quality loops, and Chapter 7 shows what they must still answer. Between two assays an estimate is only as good as its update rule, and there the update gain, not the model, decided whether the outer loop helped.

Economics at the top of the structure. A top-down plant-wide structure for a grinding circuit defines its economic cost with reference to the final product of the larger plant, and uses the mill's grind curves to set the setpoints of its economic controlled variables [9]. The flotation depth, the refining utilization and the thickener density of this volume were each worth more as a chosen operating point than as a tighter loop. In the editor's view the layer that chooses setpoints from metal prices and reagent costs will matter more as ore grades fall, and it needs the protective selectors below it to keep their place.

Learning methods, used with care. A review of machine learning in mineral processing classifies published work by the nature of its data (synthetic, laboratory or industrial), its level of success and its area. It notes that practitioners may be wary of techniques that overpromise, and proposes work on data collection, technique comparison, industrial participation and cost-benefit analysis [10]. Reinforcement learning has tuned the PID controllers of a multivariable grinding circuit on a nonlinear model verified with industrial data, with the agent's actions constrained to limit the risk of instability. In simulation the tuned controllers tracked setpoints better than manually tuned ones [11]. In the editor's view tuning will increasingly be proposed by software, and Chapter 3 shows the test such a proposal needs. Its recommended tuning won on a setpoint step and lost on the plant's own disturbances.

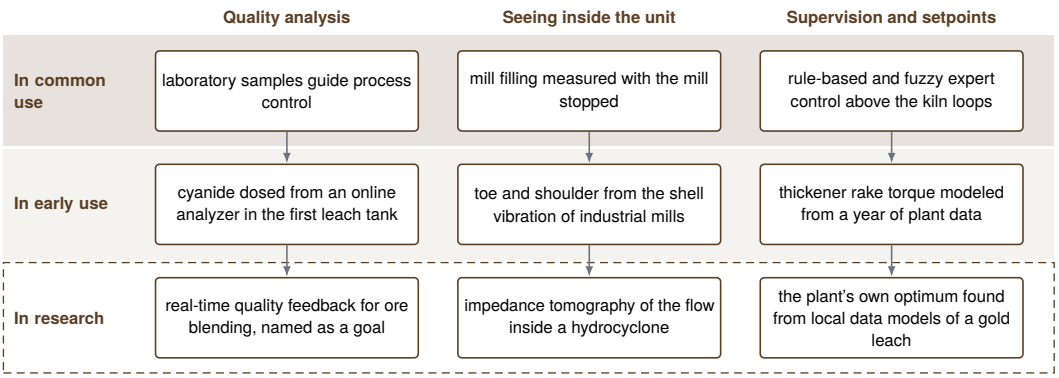


Figure 12.3 Developments in three threads. The top lane holds what the cited source calls standard or common practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, laboratory work or as a need, in the editor's reading. Sources by thread: quality analysis [5, 12, 18], seeing inside the unit [3, 4] and supervision and setpoints [13, 16, 19]. Notice that two of the three threads move the measurement into the stream or the unit.

Analyzers move onto the stream. Cyanide in a gold leach is typically controlled from a measurement in the first leach tank, taken by hand every 2 to 12 hours or by an online instrument every 5 to 30 minutes, and control from the online measurement has been reported to cut cyanide consumption by at least five percent [12]. In cement plants, online analyzers let operators determine the chemical composition of the raw materials at once [13]. In the editor's view such analyzers are the first step toward the measurement that gap 1 of Section 12.5 asks for. Chapter 7 already reads its cyanide this way, and it is the tailings, sampled every four hours and reported six hours later, that still set the pace of its outer loop.

Cells and underflows watched continuously. Open and short circuits commonly occur in copper refining, and they have been found by walking over the cells with a hand-held gaussmeter. A patented system instead senses the magnetic field around the conductor that feeds each cathode and lets a processor find open circuits as well as shorts [14]. Rheometry has seen only limited use in process monitoring, and a 2019 study reported an automated capillary rheometer that takes a side stream of a thickener underflow, answering the need for a unit that measures a stream's flow behavior continuously and unattended [15]. In the editor's view both read a variable that this volume's loops leave to inference: the nodule that shorts the cell of Chapter 8 and the underflow that loads the rake of Chapter 6.

Protection written into the plan. Published work on paste thickeners has added a rake torque constraint to a predictive controller to prevent rake lifting and pumping problems, and has sampled the near future more finely than the distant one [16]. In rare earth extraction, where practice still rests on offline analysis, manual adjustment and experience, an outer predictive controller over inner PID loops has been proposed, with an event trigger that updates the controller only when the state of the process calls for it [17]. In the editor's view Chapter 6 shows what such a controller must still be compared against: a selector whose reset law, fixed twice, moved the comparison each time. Figure 12.3 sets the developments of this section on three lanes.

12.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [20]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

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

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

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

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

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

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [21]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

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

12.4.1 How This Volume Was Made

Each of the ten apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run. Seven of the ten papers have a referee report. Five of these cover only the material added in the predictive-control rounds, and the rare earth report calls itself a self-review. The hydrocyclone, refining cell and lithium papers have none. Defects went into a log for each app, and the ten logs share no field. Eight give some entries a severity, 47 entries in all, and two give none, so this section reports cases rather than totals by kind. For this book one agent per chapter reran the numbers its exercises use and checked the app's helpers, completion steps and figures. 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.

A toe that read the truth. The mill's toe tracking computed its setpoint from the run's true rock hold-up, ball charge and liner wear, which no instrument reports. Two mills with identical readings, one carrying two more points of ball filling and the other the same 40.3 t as rock, received setpoints 3.19° apart in one scan. With that knowledge toe tracking passed the graded shift. A technical review of the added predictive controller asked why toe tracking passed while the MPC, which had to estimate the same toe, did not. The fixed app aims at the toe its filter infers from the bearing pressure with the declared ball charge. No impact tuning then passed with toe tracking, the challenge's suggested settings became the slower load loop, and every toe-tracking number in the chapter moved. The volume's lesson that an inferred toe cannot see a lost ball charge exists because the leak was closed.

A soft sensor built on the plant's own parameters. The gold circuit's soft sensor, and the MPC that shared it, read the leach kinetics, the refractory gold and five more constants live from the run's parameters. Raising the leach rate constant by 20 % moved the estimate by −0.0181 g/t with every measurement fixed, and an edit at 30 h moved the cyanide setpoint at 30.11 h, before any assay. A truth audit found it. The constants now sit at the declared values, so an edit of the plant is model mismatch that the laboratory has to find, which is the problem the bias update exists to solve. No run moved. The review that followed found that the truth tests missed a path by which a new estimator could reach the plant's state, and five mutants were added to catch it.

A thirteen-minute constant that kept coming back. The kiln's texts gave the burning zone's answer to coal a 13 min time constant. That is the burning zone node's own constant. With the air ratioed, a coal step reaches 63 % of its change only after 30.1 min, because the zones upstream warm behind it. The technical review of the added MPC reported the claim, the paper stage timed the step, and the paper and the app's scenario texts were corrected. The figure had already reached Chapter 1, which the orchestrating agent corrected after the kiln chapter's report, and the kiln chapter's agent found one more line of the app's challenge text that still called the kiln a thirteen-minute plant.

A cause tested by removing it. The flotation paper blamed its limit cycle on the interaction between the grade loop and the froth depth, and called the grade tuning defensible for that loop alone. The book's agent held the froth depth exactly at setpoint at every physics step and ran the case again. The grade loop still cycled, from 0.74 to 8.80 % copper. The chapter now says the cycle belongs to the grade loop, and the series made it a rule to correct a cause that a run contradicts when no number has to change. Nine explanations in six chapters of this volume were rewritten under that rule or its predecessor for summaries. The hydrocyclone paper, for one, called a tuning unnecessarily sluggish that the measured plant gain makes 2.7 times more aggressive than intended.

A comparison that moved with the reset. The thickener's external reset was examined three times. A technical review found that the losing controllers settled at twice the offset the texts stated.

A later fix made the reset instant, and the book's agent then found that the instant reset lets two losing controllers take the pump in turn, which drove it from 48 to 100 % in six minutes. The app owner replaced the instant reset with finite back-calculation, and the selector now peaks at 85.45 % against 75.77 % for the MPC, with density IAE 1719 against 1777 (Table 6.7). Part of a predictive controller's measured advantage may belong to the structure it was compared with.

A factor of one hundred. The first mill model divided the liner slip by a hundred once instead of twice. At 65 % wear the effective speed was −277.9 % of critical and the power exactly zero, and the structural test suite passed, because nothing was undefined and no actuator left its bounds. The paper agent found it while writing the paper, when the worn-liner scenario diverged.

The book round. The ten chapter records list 158 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The hydrocyclone record found a printed 2.7 that runs at 2.77 %, the refining record three roundings, and the thickener record numbers that the app owner's repair later regenerated. The others were kept as printed. Four papers named a loop's action opposite to the convention of Section 1.6.2. Four helpers set or promised what the actuators could not deliver. The cyanide helper wrote 633 kg/h to a pump that stops at 500, the kiln helper wrote 20 t/h of coal to a feeder that stops at 18, the lithium helper claimed three setpoints met when one was out of reach, and the mill's feed helper, pressed at the start of an overloaded case, brought the overload alarm at 4.3 instead of 22.7 min. Chapter 1 was written from the converted chapters before their agents had rerun them, and five of its statements were corrected from their reports. It had called the mill's impact angle inferred where a shell microphone measures it, given the lithium cascade a factor of 93 where the chapter measures 52 to 93, estimated about 25 scrub stages where the rare earth chapter uses 20, stated a gold peak tailings result more widely than the chapter shows, and carried the kiln's thirteen minutes.

Every error above was caught by a check that could fail: a hidden value moved, a step timed, a cause removed and the run repeated, a comparison rerun after a fix, a scenario measured for the paper. None was caught by rereading the text that contained it. Four working rules follow. Give a controller and a soft sensor only what the plant measures. Test a cause by removing it. Rerun a comparison after every fix to either side of it. Trace a number to its source before writing it a second time.

12.5 Gaps in Mineral Process Control

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 12.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 12.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the units in the picture that opens this chapter.

1. *Measuring the ore as it arrives.* The scale of Chapter 2 reads the total, and a feeder that lost 20 % of its gain left the blend wrong under every structure until a bump test found it. Hardness, grade and leachability reach the mill, the cells and the leach unmeasured. An instrument that reports them on the time scale of the circuit is not available for most ores. A 2022 study found online elemental analyzers widely available but under-used, and read copper in a suspension with a relative error of 6 % at a processing plant [5]. A 2026 review of ore blending names data from many sources that are hard to join, models too slow for real-time control and poor adaptability across ore types as the barriers to its industrial use [18]. It is hard because hardness and leachability are not elements that an analyzer reads. Progress would be a reading of the arriving ore fast enough for a blend or feedforward loop to close on.

2. *Soft sensors that state what they cannot see.* The gold soft sensor of Chapter 7 removed six hours of dead time for the feed grade and none for the leachability. Data-driven models are used

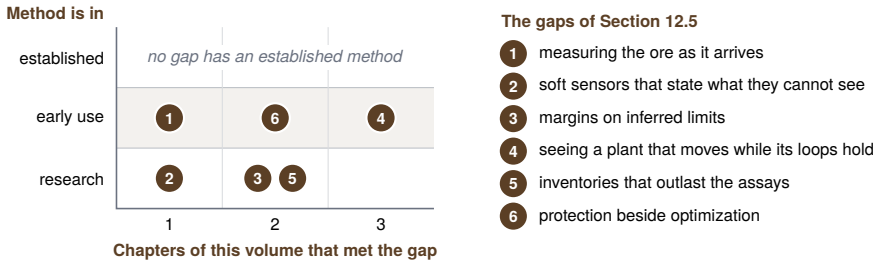


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: one for gap 1 (Chapter 2), one for gap 2 (7), two for gap 3 (3 and 8), three for gap 4 (3, 4 and 10), two for gap 5 (9 and 10) and two for gap 6 (3 and 6). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap most chapters met has a method in early use.

extensively in industrial plants because an accurate mechanistic model of a gold leach is hard to build [19], and in rare earth extraction a soft sensor has estimated the content of a monitoring stage from flows and feed analyses recorded every 30 min, where the laboratory reports daily [22]. It is hard because a model cannot report a disturbance it was never given. An estimate needs, delivered with it, the list of disturbances it is blind to, and a rule that sets its bias update from the report interval and the assay noise.

3. *Margins on inferred limits.* The refining cell of Chapter 8 held a believed 59 % at a true 73.74, and the inferred toe of Chapter 3 moved the wrong way. Experiments on copper nodules suggested a threshold size above which a nodule leads to a short, and found nodules forming without the mold release agent seen in others, for reasons not yet known [23]. It is hard because the limit moves with what the ore and the anodes bring, and an estimate of it is checked only when it fails. A routine method that sets the back-off of Equation (1.9) from data, and revises it as the ore changes, is still a matter of judgment.

4. *Seeing a plant that moves while its loops hold.* The ball charge of Chapter 3, the electrode bias of Chapter 10 and the roping of Chapter 4 each left every loop on setpoint. The evidence was in an actuator's travel, a reagent per tonne or an unmeasured density. Monitoring that reads manipulated variables and consumptions against the ore, across a whole circuit, is not routine. The Open challenge box after this list takes the gap further.

5. *Inventories that outlast the assays.* The lithium train of Chapter 10 took six hours to settle after a grade change, and a leaner feed drained the rare earth cascade's inventory for days before the purity fell (Chapter 9). A rare earth extraction process has dozens or even hundreds of tanks, and models of it seldom consider the delays between them or take them as constant [24]. It is hard because a change shows in the product only after it has passed through the inventory. Design rules for loops whose plant settles over many report intervals, and for handing such changes to production planning, are thin.

6. *Protection beside optimization.* The comparison of Chapter 6 with an MPC moved with each of two fixes of its reset law, and the MPC of Chapter 3 raised a motor overload the loops did not. On one operating thickener the flocculant pump stops when the bed reaches a set level and density, as a safety feature [25], and a study of thickeners at several mines named control instrumentation that was not adequate or not working properly among the causes of their fluctuations [26]. It is hard because a protective layer is tested only when a limit is near. A tested implementation of selectors and external reset that an optimizer must sit above, with a standard check that protection keeps its authority, would let a comparison measure prediction and not the plumbing.

⚠ OPEN CHALLENGE Seeing a Plant That Moves While Its Loops Hold

Gap 4 is the gap most chapters met, and its hidden changes sit where instruments struggle. A monitoring device inside the shell of a running mill is virtually impossible, because the tumbling charge is so severe [27]. The flow inside a hydrocyclone is difficult or impossible to observe directly [4]. The voltage of a glass pH electrode decays with age, and the electrode is calibrated against two buffer solutions [28]. A bias of that kind, 0.70 low, raised the reagent of Chapter 10 by 24 % while every loop held. The change is hard to see because each loop absorbs it. Progress would be monitoring that compares actuator travel and reagent use with the ore across a circuit and warns before the assay does.

12.5.1 Gaps Common to Agent-Era Practice

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

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [29]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [21]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [30]. 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 [31]. 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 [32]. 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 loop measures its own variable. The variable that decided the result was often one that no instrument read.
- A soft sensor removes the assay's dead time only for the disturbances among its inputs, and a loop on a late measurement is tuned to the plant's response and the report interval.
- The structures that worked moved the actuator from something measured before it reached the product, and the late assay trimmed what the fast path could not see.
- Where a loop was set often decided more than how it was tuned.
- Tuning for a step response misled on plants whose disturbances are slow and whose measurements are noisy.

References

[1] V. Lishchuk, P.-H. Koch, Y. Ghorbani, and A. R. Butcher, "Towards integrated geometallurgical approach: Critical review of current practices and future trends," *Minerals Engineering*, vol. 145, p. 106072, 2020.

[2] C. Robben and H. Wotruba, "Sensor-based ore sorting technology in mining—past, present and future," *Minerals*, vol. 9, no. 9, p. 523, 2019.

[3] S. B. Donkor, "On-line sensors for measuring the total ball and charge level in tumbling mills," MSc (Eng) thesis, University of Cape Town, Centre for Minerals Research, 2014, hdl.handle.net/11427/13225.

[4] C. J. Oates, J. Cockayne, R. G. Aykroyd, and M. Girolami, "Bayesian probabilistic numerical methods in time-dependent state estimation for industrial hydrocyclone equipment," *Journal of the American Statistical Association*, vol. 114, no. 528, pp. 1518–1531, 2019.

[5] L. Rybok, T. Kowoll, A. Thomas, A. Nad, and C. Bachmann, "Online elemental analysis for process control in the mineral processing industry," *TOS Forum*, vol. 2022, no. 11, p. 131, 2022.

[6] S. Avalos, W. Kracht, and J. Ortiz, "An LSTM approach for SAG mill operational relative-hardness prediction," *Minerals*, vol. 10, no. 9, p. 734, 2020.

[7] C. Aldrich and X. Liu, "Monitoring of flotation systems by use of multivariate froth image analysis," *Minerals*, vol. 11, no. 7, p. 683, 2021.

[8] M. Sadeghi, D. Hodouin, and C. Bazin, "Mineral processing plant data reconciliation including mineral mass balance constraints," *Minerals Engineering*, vol. 123, pp. 117–127, 2018.

[9] J. D. le Roux, S. Skogestad, and I. K. Craig, "Plant-wide control of grinding mill circuits: Top-down analysis," *IFAC-PapersOnLine*, vol. 49, no. 20, pp. 72–77, 2016.

[10] J. T. McCoy and L. Auret, "Machine learning applications in minerals processing: A review," *Minerals Engineering*, vol. 132, pp. 95–109, 2019.

[11] J. A. van Niekerk, J. D. le Roux, and I. K. Craig, "Reinforcement learning based automatic tuning of PID controllers in multivariable grinding mill circuits," *Control Engineering Practice*, vol. 165, p. 106522, 2025.

[12] W. A. M. van der Merwe, "The development of an online amperometric technique to measure free and WAD cyanide in gold plant leach liquors and effluent streams," M.Eng. dissertation, University of Cape Town, 2012.

[13] E. Worrell, C. Galitsky, and L. Price, "Energy efficiency improvement opportunities for the cement industry," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-72E, 2008.

[14] E. Y. You, A. Berges-Dreyfous, D. B. George, D. Kim, and K. Zhou, "Monitoring electrolytic cell currents," U.S. Patent 7,445,696 B2, application filed March 17, 2005, 2008.

[15] A. G. Chryss, A. Mönch, and K. Constanti-Carey, "Online rheology monitoring of a thickener underflow," in *Paste 2019: Proceedings of the 22nd International Conference on Paste, Thickened and Filtered Tailings*. Perth: Australian Centre for Geomechanics, 2019, pp. 495–504.

[16] J. Ruuska, E. Ruhanen, J. Kauppi, S. Kauvosaari, and M. Kosonen, "Multivariate linear regression model of paste thickener," in *Proceedings of the 61st SIMS Conference on Simulation and Modelling*, ser. Linköping Electronic Conference Proceedings, vol. 176, 2021, pp. 160–164.

[17] Z. Zhang, J. Zhu, C. Xiong, H. Yang, and F. Nie, "Constrained model predictive control based on event triggering for the rare earth extraction process," *ACS Omega*, vol. 8, no. 44, pp. 41 943–41 952, 2023.

[18] Z. Cai, H. Luo, and W. Yang, "The research progress and application status of ore blending technology: A review," *Physicochemical Problems of Mineral Processing*, vol. 62, no. 4, p. 226165, 2026.

[19] Q. Liu, D. He, Z. Wang, and J. Shi, "A data-driven iterative optimization compensation method based on PJIT-PLS for gold cyanidation leaching process," *IEEE Access*, vol. 7, pp. 138 683–138 693, 2019.

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

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

[22] W. Dai, R. Lu, P. Chen, and H. Yang, "Novel dual residual-enhanced deep bidirectional LSTM network for soft sensing of rare earth component content," *Sensors*, vol. 26, no. 10, p. 3152, 2026.

[23] K. Adachi, Y. Nakai, S. Mitsuno, M. Miyamoto, A. Kitada, K. Fukami, and K. Murase, "The mechanism of nodular growth in copper electrorefining," *Journal of MMIJ*, vol. 136, no. 2, pp. 8–13, 2020, in Japanese, with an English abstract.

[24] R. Lu, H. Liu, H. Yang, J. Zhu, and W. Dai, "Multi-delay identification of rare earth extraction process based on improved time-correlation analysis," *Sensors*, vol. 23, no. 3, p. 1102, 2023.

[25] M. Cooks, D. Mostert, L. Coetzee, and D. J. van Schalkwyk, "Thickener experiences at Langer Heinrich mine," in *Paste 2012: Proceedings of the 15th International Seminar on Paste and Thickened Tailings*. Perth: Australian Centre for Geomechanics, 2012, pp. 141–153.

[26] B. Pirouz, S. Javadi, and K. Seddon, "Thickener performance variability: underflow solids concentration and flowrate," in *Paste 2017: Proceedings of the 20th International Seminar on Paste and Thickened Tailings*. Beijing: University of Science and Technology Beijing, 2017, pp. 29–40.

[27] R. Rajamani, J. Delgadillo, and V. Duriseti, "Online SAG mill pulse measurement and optimization," University of Utah, Metallurgical Engineering Department, Semi-annual technical progress report DE-FC26-04NT42088, 2006, u.S. Department of Energy, Office of Scientific and Technical Information.

[28] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[29] 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.

[30] 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.

[31] 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.

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

12.6 Exercises

12.1 Discussion · 8 min

An agent proposes to set the bias update gain of the gold circuit in Chapter 7 to one, "so that the soft sensor uses every assay in full, because the assay is the truth". Using Chapter 7 and Section 1.4.5, explain what the proposal overlooks, and say what evidence you would ask for before accepting any gain.

12.2 Discussion · 8 min

An agent adds a supervisory setpoint to a semi-autogenous mill. The impact angle is aimed at a toe that a model computes, and the agent reports that the graded shift of Chapter 3 now passes. Using Chapter 3 and Section 12.4.1, list what you would check before accepting the result, and say which check found the defect in the first app.

12.3 Discussion · 7 min

An agent recommends adding an inner loop under every quality loop of a circuit, because “a cascade rejects disturbances before they reach the product”. Using Chapter 4 and Chapter 10, say for which disturbances the claim held, for which it failed and why, and what test you would ask for.

12.4 Discussion · 8 min

A steady-state sweep by an agent finds that the graded crop of Chapter 8 is worth most at 62 % utilization and 65 °C, \$35 661, and it recommends that setting. Using Chapter 8 and Equation (1.9), discuss the recommendation.

12.5 Discussion · 8 min

An agent reports that the economic MPC of Chapter 5 earns \$1050 /h more than the shipped loops on scenario 11 and recommends an MPC for every cell. Using Section 12.2 and the chapters it cites, list the comparisons you would ask for, and say what each did in this volume.

12.6 Discussion · 7 min

A colleague argues that Chapter 1 could quote the papers without rerunning anything, because every paper was refereed. Using Section 12.4.1, answer the argument, and propose a routine that would have caught the kiln’s thirteen-minute constant at its first use.