INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 16 Control in Iron, Steel, and Metals Production

CHAPTER 12

Summary and Outlook

Pages 259–273

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Iron, Steel, and Metals Production*, John D. Hedengren, Ed., vol. 16 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 259–273.

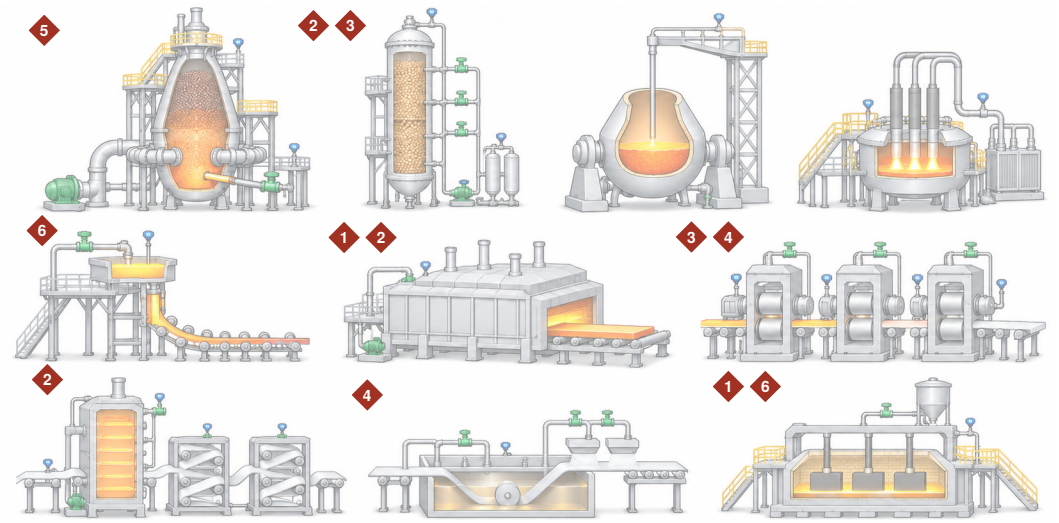
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v16-iron-steel-metals/

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

CHAPTER 12

Summary and Outlook



Open gaps 1 state inside 2 schedules 3 blind side 4 delay as length 5 two clocks 6 diagnosis

Notice that the shaft furnace, the reheat furnace, the strip mill and the smelting cell each carry two gaps, and that no gap marks the oxygen converter or the arc furnace. Illustration, not to scale.

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

The main idea. Each chapter of this book ran one unit of a steel or aluminum plant in a simulator, from the blast furnace to the coating line. Every unit had a slow part, such as a furnace full of hot metal, and a fast part, such as a valve or a motor.

The challenge. A fast loop can hold its own reading steady while the slow part drifts away.

Why it matters. The units that did best planned the slow part with a model, corrected it from the few true readings, and gave the fast part room to move. Many of the open problems lie between the plan and the loops.

Ten chapters followed the metal from a blast furnace to coated strip and added an aluminum smelting cell. Each joined a slow part to a fast one. Table 1.1 of Chapter 1 listed, for each chapter, what is slow, what is fast and what joins them, and Section 1.4.7 named four ways of joining them: a setpoint computed by a model of the slow part, a feedforward from the fast part's measured changes, an inventory that lets the two run at different rates, and a slow loop that returns a fast actuator to its range. This chapter adds what each arrangement 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. The other columns name an arrangement that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

Table 12.1 How the Slow and the Fast Parts Were Joined, Chapters 2 to 6

Chapter	Slow part and fast part	What worked	What failed	Measured result
2, blast furnace	Coke after a 5 h descent. Coal within the hour but bounded. The hearth dipped every 2 h.	A slow mid-ranging loop on coke. A bias update of the soft sensor at each cast. A gentler coal loop.	The coal loop alone past its limit. A fast mid-ranging loop. A loop tuned as if the hearth were read continuously.	Last 12 h at 1477 against 1504 °C with mid-ranging (Table 2.4). At $K_v = 1.0$, 7.0 h cold (Table 2.6). Challenge's coldest hearth 1337 against 1459 °C (Table 2.7).
3, hydrogen shaft	Product over a 6 h residence. Top gas within minutes. Hydrogen from a wind-fed buffer.	A hotter bustle under the sticking push. The supervisor on tons. Feedforward on the measured grade.	The supervisor off. Priority on quality through the lull. The soft sensor trusted while the product turns.	Challenge passed on 5 of 5 seeds against 0 of 5 shipped. Supervisor off, buffer dry and worst sample 54.5 to 61.5 % (Table 3.5). Feedforward 94.16 against 92.85 % (Table 3.4).
4, oxygen converter	Lining heat loss drifting heat to heat. One blow of about 21 min and one sublance reading.	Charge model, sublance, the rate endpoint, a hot aim, learning the heat loss between heats.	A fixed recipe. The off-gas integral from the charge. A fixed margin on a wearing lining.	Hits in ten heats 1, 2 and 9 for recipe, static model and rate endpoint. Integral 3 hits, 9 reblows. Wearing lining 5 hits, 7 with learning (Table 4.5).
5, arc furnace	Bath heat and taps in 50 V steps. Three electrodes every 0.02 s on one floating neutral.	Impedance regulation with the staged tap program and foamy slag.	Current regulation, which reads a neighbor's arc as its own error. The top tap on a fresh pile.	Step overshoot 26 to 40 against 9 to 23 % (Table 5.4). Graded heat passed on 19 of 21 seeds against none shipped (Table 5.5).
6, caster mold	Casting speed sets the bulging frequency. Mold level integrates, sensor 0.3 s late.	Diagnosis first, then a canceller at the roll-passing frequency, a knocker and a lower gain.	Any single remedy. Retuning first, which removes the gain that held the bulging down.	Period 12.5 to 10.0 s at 1.5 m/min, canceller to the 0.14 mm floor (Table 6.4). Challenge 3.34 as loaded, 5.31 retuned, 0.77 mm with all three (Table 6.5).

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

A fast loop holds what it measures. The most repeated result of the volume is a loop that did its job while the slow part went elsewhere. In Chapter 7 the zone loops held their setpoints within 7 °C through a pace change and the slabs left 55 °C cold. A pyrometer trim held the surface at 1212 °C and left the mean at 1174 to 1177 °C at every trim gain tried (Table 7.5). In Chapter 8 the gaugemeter loop held its own reading on setpoint against thermal crown and left the strip 198.6 µm thin, against 41.9 µm with the loop open. In Chapter 2 a drifting soft sensor read 1514 °C while the hearth sat at 1464 °C (Table 2.4). In Chapter 11 the feeder's tally stayed within 1.8 kg of balance while shots 2.8 % short ran the bath down to an anode effect. Each fast loop worked. Each measured a variable that the slow part could leave behind.

The slow measurement corrects the fast one, and slowly. Where a true measurement exists it arrives late, and its correction is a filter. The hot strip mill's X-ray trim, at roughly an eighth the speed of the loop it corrects, held the crown's peak error to 12.5 µm. Faster trims became the delayed loop again and cycled at up to 201 µm (Section 8.7). The blast furnace's bias update at each cast held the hearth at 1492 °C against 1464 uncorrected, and against a drift of 3 °C per hour a weight of 1.0 beat one of 0.2, 5.1 against 9.6 °C rms, because the drift outweighed the dip noise. The reheat furnace's correction from the pyrometer put all 119 slabs on target where the uncorrected model put 30, and

Table 12.2 How the Slow and the Fast Parts Were Joined, Chapters 7 to 11

Chapter	Slow part and fast part	What worked	What failed	Measured result
7, reheat furnace	A slab 2.5 h inside whose mean is never measured. Mill pace and zone loops.	Slab-tracking setpoints every 120 s, the charge temperature as feedforward, a correction from the pyrometer.	Zone loops on fixed setpoints. A pyrometer trim at any gain. The model uncorrected after the exchange factor fell.	On target 21, 21 and 117 of 117 for zones, pyrometer and model. After the fall, 30 and 119 of 119 uncorrected and corrected (Table 7.4). Challenge 22 against 151 of 159 (Table 7.6).
8, hot strip mill	X-ray gauge 0.5 s late, crown and modulus drifting. Gaugemeter at the stand on a 4 ms scan.	The gaugemeter, an X-ray trim at an eighth of its speed and a synchronous eccentricity estimator.	The same gains on the X-ray gauge. The gaugemeter alone against crown.	Entry wave 11.6 against 173.8 µm rms (Table 8.6). Challenge peak 61.2 against 20.3 µm, prime yield 65.4 against 98.8 % (Table 8.7).
9, annealing line	Furnace with a 240 s lag. Line speed and an 80 s residence. An entry that stops for every weld.	A looper allowed to swing, a speed cap, look-ahead, speed scaled by thickness, firing for the speed.	A tight looper loop. The speed held through a thicker coil.	Off specification 71.3 tight against 0.0 % averaging, 77.5 against 2.5 % on 1.3 mm strip (Table 9.4). Challenge 76.0 against 10.1 % (Table 9.8).
10, galvanizing line	Gauge 100 m downstream. Knife pressure fast, line speed in ramps, a heavy grade's gain 5.6 times.	Speed feedforward. A Smith predictor that keeps its delay in meters of strip.	The predictor with its delay in seconds. A PI tuned for the light grade on the heavy one.	Below the minimum after a slowdown 11.6 against 0.0 % with feedforward. Predictor 40.1 in seconds, 0.0 in meters (Table 10.5). Challenge 50.5 against 0.22 % (Table 10.6).
11, smelting cell	Heat and ledge, one dip a day. Resistance read every second, with a minimum it cannot see past.	Demand feeding on the resistance slope, a heat loop on the daily dip, compensation of anode changes.	Feeding on a tally. A low resistance target left from an energy drive.	Fixed feeding's first anode effect at 32.53 h, none in two days on demand (Table 11.5). Challenge five anode effects as shipped, none suggested (Table 11.6).

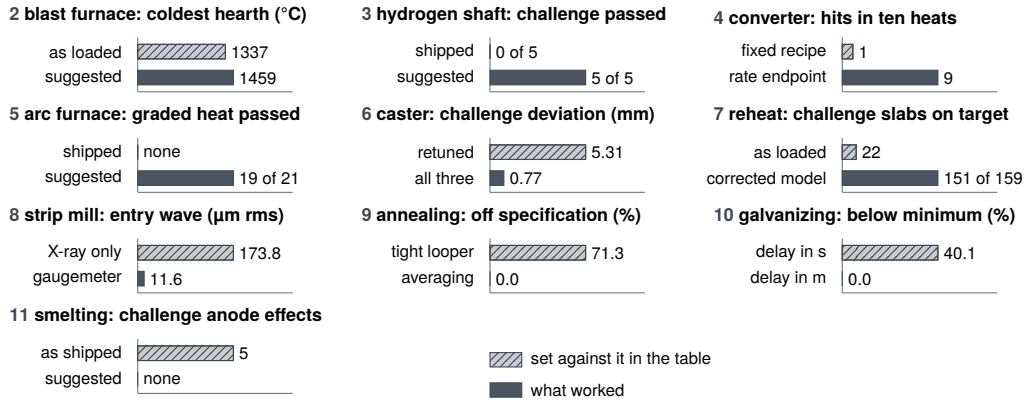


Figure 12.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 3, 4, 5 and 7, and a smaller one in the others. The blast-furnace panel is the coldest hearth of the challenge, and on a scale from zero its two bars look close. Notice that in most rows the better result came from a change of structure or measurement, not of tuning alone.

with noisier readings the largest correction gain lost the most slabs (Table 7.5). The converter learned its lining's heat loss between heats. A weight of 0.5 restored seven hits, and a weight of 1.0 gave seven at 14.5 against 12.2 °C rms, because each update copied more of one heat's random error into the next plan (Table 4.5). The weight of a slow correction is set by the drift it must follow and the noise it must not pass.

Speed is the fast handle, and an inventory lets it move. The exit temperature of a continuous furnace depends on speed and thickness only through vd (Equation (1.12)). The annealing line that held its speed through a change to 1.3 mm strip left 77.5 % of it outside the window, and the line that scaled its speed by thickness left 2.5 %. The looper made the speed change possible, and only when allowed to swing. Held tightly at 75 %, it passed every weld into the speed and put 71.3 % of the strip off specification. Allowed to fall from 96 to 38 %, it held 300 m/min and put none off (Table 9.4). In Chapter 10 a slowdown raised the gain and the delay together by 1.67, and speed feedforward restored the gain, lifting the gain margin from 1.16 to 1.93 against 3.05 at full speed, but not the delay (Table 10.3). In Chapter 6 the speed set the frequency of the disturbance. Bulging moved to a 10.0 s period at 1.5 m/min, exactly roll pitch over speed, and a canceller at that frequency took the level to the noise floor. The hydrogen shaft's buffer stood between the wind and the furnace as the looper stood between the welds and the strip. Through a six-hour lull no setting or schedule among 984 kept both the full tons and every sample at specification (Section 3.6).

A slow, wide actuator returns a fast, limited one to its range. The blast furnace's coal acts within the hour and stops at 200 kg/tHM. Two burden changes asked for about 67 kg/tHM of extra coal against 50 of room, so the coal sat on its limit for about 26 h. A mid-ranging loop moved the coke, which waits five hours, and lifted the last half day from 1477 to 1504 °C for $1.73 against $1.68 million of fuel a day (Table 2.4). It had to be slow. At $K_v = 0.5$ the hearth overshot to 1541 °C, and at 1.0 the loops fought through the coke delay, the coke reached its 420 kg/tHM limit and the hearth spent 7.0 h cold (Table 2.6). Neither arrangement prevented the first dip, to 1448 °C in both, because the burden reached the hearth before either fuel could. The resting point of the fast actuator is an economic choice. A coal target of 190 kg/tHM saved about $33,000 a day against 150, and it left the coal 10 kg/tHM of room where a burden change asked for about 21 (Section 2.5). The arc furnace divides its work the same way, between electrodes regulated every 0.02 s and a transformer tap that moves in 50 V steps.

The structure decided more than the tuning. The clearest measurement is the hot strip mill's challenge. Two configurations on the delayed X-ray gauge differed in every tuning and by almost nothing in result, 61.2 against 60.9 µm of peak deviation, while three that differed in structure spanned 20.3 to 65.9 µm (Table 8.7). The arc furnace showed the same with what a loop measures. Over 21 seeds the impedance regulator with the staged taps passed the graded heat 19 times, and a current regulator passed 16 times only when detuned to $K_I = 10$ (Table 5.5). In Chapter 6 no single remedy met the challenge, and retuning first made the peak worse. On the converter the same off-gas signal gave nine reblows integrated from the charge and nine hits with none read as a rate after the sublance. A tuning study asks how well a structure can do. It cannot rescue a structure that measures the wrong variable.

12.2 Choosing a Method

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

1. *Which clocks does the unit run on?* Estimate the zone time constant, the product's conduction or residence time and the measurement's delay first (Section 1.4.2, Chapter 7, Chapter 3).

2. *Is the product measured, and where?* If not, plan the slow part with a model and correct it from the slowest true measurement (Chapter 7, Chapter 2). For an inferred measurement, write its error identity, as Equation (1.10) does, and give each disturbance to the fast loop or the slow trim (Chapter 8).
3. *Is the delay a time or a length?* Keep a predictor's delay in meters of strip and feed the speed forward to the gain (Chapter 10). Check what a speed change does to the gain, the delay and the disturbance frequency together (Chapter 6).
4. *Can the fast actuator carry the worst load?* Apply the capacity test of Equation (1.3). If it fails, close a slow loop on the wide actuator, slow in gain as well as integral time against that actuator's delay, and choose the resting point from the fuel price and the disturbances (Chapter 2).
5. *Is the product a batch judged at its end?* Plan from the measured charge, aim to the side that the late correction can reach, and learn a drifting constant with a weight that passes little noise (Chapter 4).
6. *Is the line fed in pieces?* Size the inventory against the longest stop with Equation (1.8), let it swing, cap the speed below the average supply and slow ahead of a stop that the inventory cannot carry (Chapter 9).
7. *Has an oscillation more than one cause, or a reading more than one side?* Diagnose before treating. Open the loop, change the speed and read the signature (Chapter 6). Work on the side of an extremum that the reading can tell (Chapter 11).

A predictive controller paid where it held information the loops did not have. On the hydrogen shaft it planned the feed on the wind forecast. At a 30 % buffer setpoint it made 2225 to 2297 t on five seeds with the buffer at 33.8 % or more, where the override let the buffer fall to 10.4 to 14.5 % (Table 3.5). On the blast furnace it scheduled the coke on its five-hour delay, kept the coal off its limit in scenario 10 and cut the IAE from 1016 to 644 °C h for 395 t more coke over three days. A coal loop retuned in hindsight reached 512 by moving the coal three times as far (Table 2.8). Prediction lost where a simpler structure already carried the timing. Through the reheat furnace's mill delay it put 61 of 83 slabs on target against 80 for the slab model's delay rule, and it failed the challenge at 125 of 159 where the slab model passed at 151 (Table 7.7). On the annealing line it filled the looper before an announced long weld, and from the announcement to one hour it left 14.43 against 13.25 % of the strip off specification, more on all ten seeds (Table 9.6). Before reaching for prediction, try the model of the slow part, the feedforward and the inventory. Volume 9 treats that choice in full.

12.3 Where Practice Is Heading

Hydrogen, and the electricity behind it. A design study of fossil-free steelmaking by direct reduction with hydrogen found that it needs 3.48 MWh of electricity per tonne of liquid steel, mainly for the electrolyzers. It named flexibility in production and in electricity demand as a key feature, with grid balancing through storage of hydrogen and of hot-briquetted iron or through the share of scrap [1]. A study that joined a long-term energy system planning tool to a process model of the same route sized the storage it needs. For the United Kingdom's industry the optimal system held 600 GWh of hydrogen storage, and 140 GWh if a fifth of the electricity came from dispatchable generation [2]. The shaft of Chapter 3 meets this problem over one day. Its shipped buffer could not carry a six-hour lull at full tons and full quality. In the editor's view the size of the buffer and the supervisory layer that trades tons against quality will become part of the control design of such plants, not a fixed input to it.

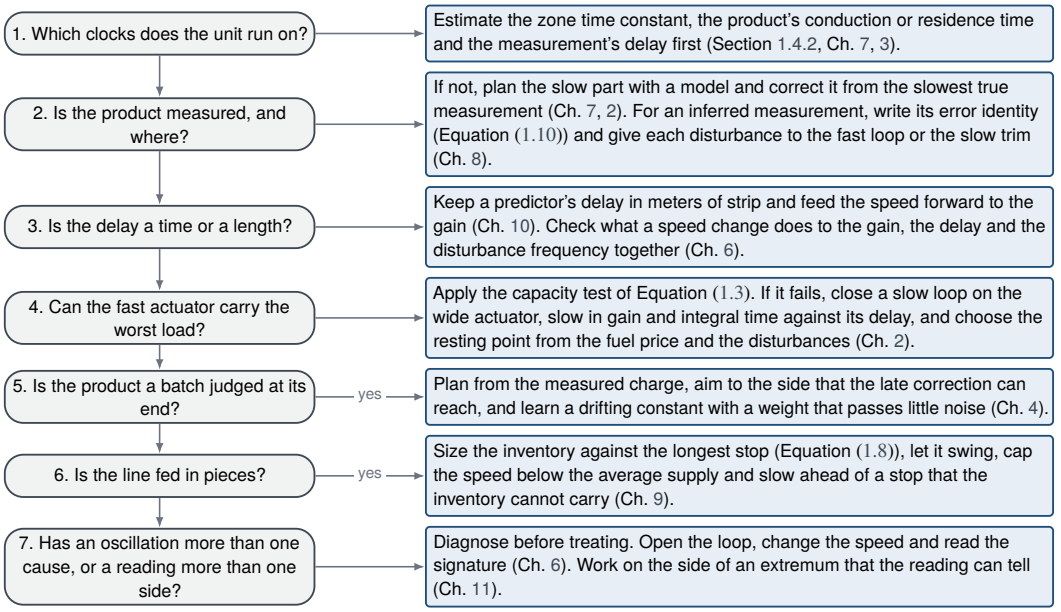


Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the first three questions ask about clocks, measurements and delays, and are answered before any loop is tuned.

> ⧖ HISTORY **2004 to 2026: Learning From the Same Readings**
>
> The second history box of Chapter 1 ended in 1984 with a mold level held within 3 mm. Twenty years later a study of blast furnace hot metal temperature still described two or three measurements by cast, taken in the runner at the exit of the tap hole [3]. In 2025 a model predicted the silicon of the next cast of one furnace from the past shift's operating conditions, with an accuracy of 91 % over 200 predictions, because the hostile environment inside precludes a direct temperature measurement [4]. In 2026 a model forecast the mold level of a slab caster from plant data sampled once a second, 2.2 million points over 13 casts [5]. In the editor's reading the models have moved from physics to data, and they still learn from the dip and the level sensor of the earlier boxes.

Electricity prices in the melt shop and the potroom. Production schedules have been optimized together with the purchase of electricity. A continuous-time formulation that adds energy awareness to the daily schedule was applied to the melt shop of a stainless steel plant, and optimizing the schedule jointly with the net electricity cost mitigated the effect of high day-ahead prices [6]. Aluminum smelters have investigated modulating their power to offer services to the grid, driven by energy prices made volatile by renewable sources [7], and shell heat exchangers have been studied as a way to widen the window over which a pot's power can be modulated [8]. The cell of Chapter 11 shows what such modulation asks of control. A resistance target 0.73 μΩ too low ran at 12.57 kWh/kg while the ledge grew to 37.4 cm (Section 11.9). In the editor's view a modulated cell needs an estimate of its heat between the dips, because each change of power is a deliberate upset of the slowest balance in the volume.

The caster and the mill planned together. Planning the continuous caster and the hot strip mill as one problem has been done through "virtual" slabs, possible forms of steel not yet cast, which pass information between the two plans. The approach was solved by mathematical programming with

heuristics and was implemented at several mills [9]. Such a plan sends slabs to the reheat furnace hot. In Chapter 7 the slab model took the charge temperature as a feedforward and burned 0.923 GJ/t against 0.955 on the operator setpoints with 94 against 41 of 95 slabs on target (Table 7.4). In the editor's view the plan and the furnace model will share more, such as the charge temperature of each slab and the length of each announced delay.

Endpoints predicted from data. Data-driven models have been built to predict the endpoint temperature and the phosphorus, manganese, sulfur and carbon of a converter heat. A model on static features was compared with one that adds time series from the blow, which give a better estimate of the endpoint and allow the process to be tuned, and simple feature extraction improved the prediction of phosphorus, manganese and temperature [10]. The static charge model of Chapter 4 hit only two temperatures in ten at 19.0 °C rms, and all ten at 3.4 °C with every heat nominal (Section 4.5), so its misses came from what it was not told. In the editor's view a better prediction will narrow the hot margin, while the one-way correction and the learning between heats will remain.

Seeing inside the vessel and on the strip. New instruments shrink the part of the plant that a model must infer. A European project has studied muon radiography, which uses the muons of cosmic rays that pass through the furnace, to image the cohesive zone of a blast furnace, where the descending burden begins to soften and melt [11]. On the mill, a shape meter that projects a pattern of LED dots onto the strip measured flatness over more than 98.5 % of every coil, and the meter and the flatness control built on it have been implemented in a hot strip finishing mill [12]. In the editor's view the blast furnace of Chapter 2, whose hearth is read only at the cast, has the most to gain from such an instrument.

A mold with a magnetic actuator. A 2025 study superimposed alternating and direct magnetic fields on the molten steel in a slab mold, in simulation and at an industrial caster, and found that surface defects on the sheets could be roughly halved against a direct field alone [13]. The caster of Chapter 6 lists electromagnetic braking among what a fuller model would add (Section 6.8). In the editor's view a field that shapes the flow becomes one more input to the level loop's plant, and the diagnosis of that chapter will have to tell its effect apart from bulging and a sticking gate.

Hydrogen iron in the arc furnace. Iron reduced by hydrogen carries none of the carbon that natural gas leaves in the pellets, so the carbon must be adjusted in the arc furnace to metallize the iron and to foam the slag. A changing and growing slag has been named one of the challenges of melting it, with more refractory dissolution and a higher power demand. Optical emission spectroscopy, proven in pilot and industrial furnaces, was proposed to watch the arc, the residual hydrogen and the slag during melting [14]. In the editor's view the shaft of Chapter 3 and the arc furnace of Chapter 5 then become one control problem, because the product of the first sets the slag under which the electrode regulators of the second work. Figure 12.3 sets developments of this section and of the volume's chapters 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.

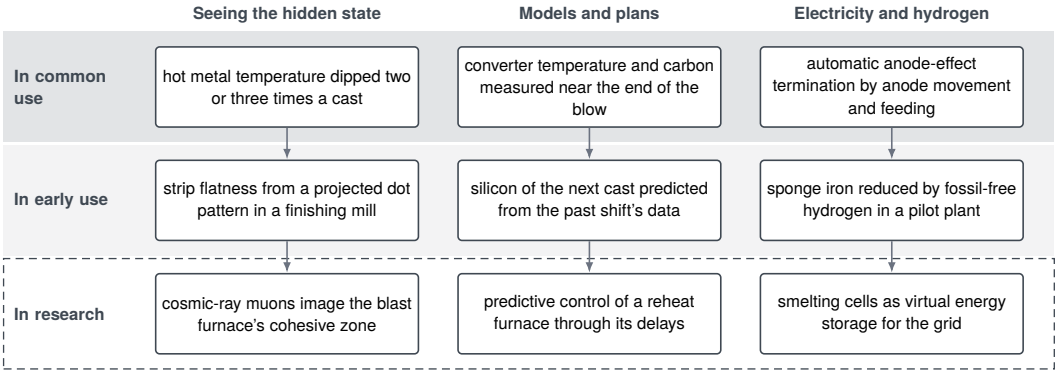


Figure 12.3 Developments in three threads. The top lane holds what the cited source calls standard or common practice, or in use at many plants, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation, an optimization study or as a project, in the editor's reading. Sources by thread: seeing the hidden state [3, 11, 12], models and plans [4, 15, 16] and electricity and hydrogen [17–19]. Notice that the first thread moves from a reading taken at the cast toward an image of the inside of the vessel.

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 each number to a run. All ten referee reports are self-reviews by the agent or stage that wrote the paper, and three carry a second pass on the predictive-control material added later. Eight of the ten papers were measured again in a consistency round and an app-fix round, and the other two were written after them. Defects went into a log for each app, 105 entries with a severity in all. Of these, 18 are wrong behavior of an app or its model, 65 are documentation, 15 cosmetic or display, and 7 something else, such as a failing test. For this book one agent per chapter reran the numbers its exercises use and checked the app's briefings, helpers and completion steps. The editor set the scope of each round and what each agent could change. Five cases show what went wrong and what caught it.

A furnace that skipped its own delay. In the blast furnace app a coke rate, coal rate, blast or burden chosen before Run acted at once. A coke rate of 360 kg/tHM warmed the hearth 5.7 °C in the first hour and 14.7 °C by the fourth, as if the coke had no five-hour descent, and a coal rate showed no inverse dip. The model started each thermal deviation at zero while its delay line held the new inputs for all earlier times, so the furnace began at the base temperature with a history that said the new rates had always been in force. It was found in an app-fix round, when two of the app's own checks failed after a shared engine change made a setting before Run a starting condition. The furnace now starts steady at the settings in force, and no shipped run moved.

A sampling effect that was a tuning effect. Case 3 of the same app was meant to show what the two-hour cast costs a loop tuned for a continuous measurement. Writing the paper, the agent ran the tuning on an ideal continuous measurement, where at the shipped gain of 1.2 the hearth swung from 1401 to 1562 °C too. The case moved to a gain of 0.8, and a consistency round found 0.8 still past the continuous limit: the loop settled at 0.65 but not 0.70 on the continuous measurement and at 0.45 but not 0.50 on the casts. The 48 h minimum-temperature check of the first fix could not show a cycle that grows over a day. The chapter now states both limits and the factor of about 1.4 by which the hold lowers the stable gain.

A span with no length. The hot strip mill's second stand rolled the first stand's present exit, with no transport delay across the span between them. The first stand reads its incoming strip 6.6 m ahead, so a skid band thickened the strip at the second stand 0.66 s before the cold band itself arrived there, and the open-loop skid peak was 119.6 against the 133.5 µm that a steady solve of both stands gives. An agent fixing other defects logged it. The span is now keyed on the delivered coordinate of the steel, so its delay follows every speed change. Every run moved, and the challenge's suggested estimator rate had to fall from 3 to 2/s to pass again. A volume whose lesson is that delay is a distance had a model that left one out.

Controllers that read the plant's constants. In the hydrogen shaft app the predictive controller, the buffer override, the soft sensor's residence filter and the sticking push took the plant's own parameters at each execution. An edit of the bed holdup from 10153 to 12000 kmol moved the

residence filter from 6.03 to 7.12 h, and an edit of the sticking temperature from 950 to 910 °C moved the push's measured margin from 48.8 to 8.8 K, with every measurement unchanged. A truth audit found it. The controllers now read the shipped design values, and a test edits 32 parameters and checks that 15 controller outputs stay fixed at fixed measurements. No number of the paper changed, because no case pins a parameter.

A comparison that rested on its window. The annealing line's predictive controller was compared with the capped looper loop on one scenario. Over the first 2400 s it put 17.46 against 21.60 % of the strip off specification. The referee pass on the new material found that 563 m of the capped loop's off-specification strip came from its speed step at the start of a case that begins at full speed, before any weld was announced. The paper and the chapter now report both windows, and from the announcement to one hour the predictive controller loses on all ten seeds. The same pass left one shipped behavior for the owner. A difficult weld announced during a weld also lengthened the next weld, and correcting it took the suggested challenge configuration from 14.2 to 10.1 % off specification.

The book round. The ten chapter records list 143 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The other three each name one printed value that a run does not reproduce, among them an annealing table cell printed 10.297 where the run gave 12.997. The paper and the chapter now print 12.997, and the other two stand as printed. Seven explanations in five chapters were contradicted or found incomplete by runs and were corrected without a new number. The shaft paper blamed the feedforward's shortfall on the heat a richer lot needs, when the cause is the residence. A fixed holdup passes the richer feed's iron faster, so each pellet spends less time under the gas. The paper, the app's explanation and its hints now name the residence. The mill paper said a slowed mill loses margin because the ultimate gain falls with the delay, and a bisection showed that the proportional limit barely moves while the room for integral action shrinks. The other five chapters tested their causal claims by run, and the claims held. Six statements in four chapters named a loop's action opposite to the convention of Section 1.6.2, and one app's model dialog, the galvanizing line's, still does. Chapter 1 was written from the converted chapters, and five of its statements were corrected from the chapter reports. It had said, for one, that speed feedforward removes both effects of a slowdown, when it removes the gain change and not the longer delay.

Every error above was caught by a check that could fail: a step timed against its delay, a comparison on an ideal measurement, a hidden value moved while the measurements were held, a second window, a cause removed and the run repeated. None was caught by rereading the text that contained it. Four working rules follow. Time every delay the model claims, from a start that is already steady. Give a controller only what the plant measures. Report a comparison in every window that changes its sign. Test a cause by removing it.

12.5 Gaps in Metals 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. *Knowing the product's state inside the furnace.* The reheat model of Chapter 7 read its slabs on target while 30 of 119 were, until a correction from the pyrometer found the exchange factor. The smelting cell's heat loop in Chapter 11 passed its challenge partly because both dips read the superheat low, 3.30 against 8.07 K at the first. Trials with instrumented slabs in a regenerative

reheating furnace found that its total heat exchange factor rises as the fuel consumption or the hot charging temperature falls and as the productivity rises, and offered the result for the dynamic compensation of online control models [22]. In a blast furnace the hostile environment precludes a direct temperature measurement, and the silicon of the hot metal serves as a proxy that arrives with significant delays [4]. The gap is hard because the state that matters lies inside a slab, a hearth or a bath where no instrument lives, and a model's error shows only at the next true reading. An estimate that states its own uncertainty between true readings is not routine.

2. *Schedules that reach the control layer.* The reheat furnace's Level 2 model in Chapter 7 is not told how long a delay will last. The annealing line's predictive controller in Chapter 9, told only what the cap and look-ahead know, put 40.48 against 17.46 % of the strip off specification, and the controller of the shaft in Chapter 3 gained tons because it read the wind forecast. Standard ways to pass delays, weld lengths, coil thicknesses and supply forecasts from planning to control, with a measure of what each is worth, are missing. The Open challenge box after this list takes the gap further.

3. *Inferred measurements that name their blind side.* The gaugemeter's error identity divides every disturbance of the mill in Chapter 8 into one it sees and one it inverts. The shaft's soft sensor in Chapter 3 overstated the product by 2.4 points at hour 8 while the product turned (Table 3.4). In a hot strip finishing mill the gauge cannot be measured at the stand, for lack of space and a hostile environment, so the rolling force is the only measurement of the gauge controller, whose gain is usually set low to stay stable as the mill modulus varies [23]. A review of soft sensors noted that material takes a given run time through a process, so readings taken at the same moment at different places should not be related directly [24]. The gap is hard because an inferred value is checked only against a late or rare true reading, and its blind side shows only when a disturbance of that kind arrives. Most soft sensors in metals plants have no error identity, and a method that derives one from the sensor's model and places the slow trim's crossover from it is not available.

4. *Delays kept in material coordinates.* The galvanizing predictor of Chapter 10 that held its delay in seconds failed once the true delay exceeded its own by about 10 s, and the one that held it in meters carried a 5.6-fold gain as well. The model of Chapter 8 first had no delay across its span. The Smith predictor, one of the most popular dead-time compensators, separates the process model into a fast model free of delay and a delay model [25]. On a galvanizing line the coating gauge stands where 100 to 200 m of strip has left the zinc pot, and the knives are adjusted from its deviation [26]. The gap is hard because most formulations write the delay model as a time, while the plant's delay is a length divided by a speed that the line keeps changing. Controller blocks still take a dead time in seconds, and a block that takes a length and a measured speed would remove a class of error.

5. *Two actuators on two clocks without a fight.* The mid-ranging loop of Chapter 2 at $K_v = 1.0$ fought the coal loop through the coke delay, and the model has no anti-windup coupling between the two loops (Section 2.7). In 1982 one blast furnace held its hot metal temperature by the injection rate of heavy fuel oil [27], and the trend in coal injection rates is toward a long-term average of 200 kg/tHM, 40 % of the total fuel rate [28]. The gap is hard because the fast actuator now carries a large share of the fuel and runs near its limit, while the actuator that relieves it acts hours later. Tuning rules for mid-ranging across a delay of hours, with a guarantee that the slow loop cannot drive the fast one onto a limit, are thin.

6. *Diagnosis before treatment.* The three causes of oscillation in the caster of Chapter 6 were separated only by opening the loop and changing the speed, and each remedy applied to the wrong cause cost gate travel. The smelting cell's reading in Chapter 11 cannot tell on which side of its minimum the alumina lies. A 1997 study of oscillating control loops noted that a full diagnosis typically requires a special test for which the loop is taken out of routine operation [29],

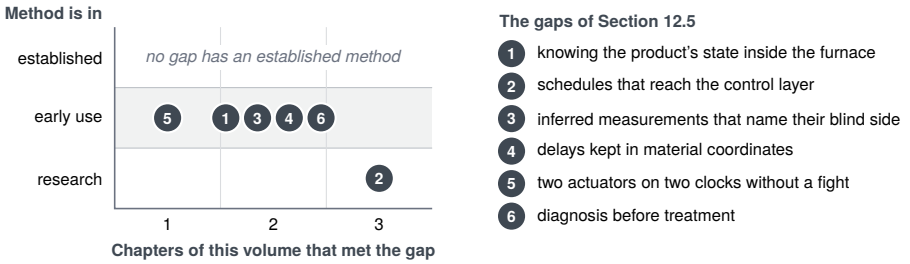


Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 7 and 11), three for gap 2 (3, 7 and 9), two for gap 3 (3 and 8), two for gap 4 (8 and 10), one for gap 5 (2) and two for gap 6 (6 and 11). Up is the maturity of the method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

and a study of thin slab casting called mold level hunching from unsteady bulging a main obstacle to higher casting speed, whose origin and detailed mechanisms were not then known clearly [30]. The gap is hard because the causes share one band of frequencies and a running mold cannot be opened at will. The caster chapter lists an automated stiction test for a fuller model (Section 6.8), and a test that names the cause of an oscillation in a running mold is not standard.

OPEN CHALLENGE Schedules That Reach the Loops

Gap 2 is the gap most chapters met. A study of continuous annealing placed the optimization of process settings between the production scheduling layer and the process control layer [31]. A predictive strategy for a reheat furnace handled three typical delays and raised its control accuracy from 74.79 to 99.17 %, in simulation [16]. Hydrogen made from intermittent renewable power meets a steel plant's steady demand with possible mismatches between the two [18]. The gap is hard because a schedule belongs to planners and a loop to control engineers, and neither owns the message between them. Progress would be a standard message that carries delays, weld lengths, thicknesses and forecasts to the control layer, with a test of what each is worth.

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 [32]. 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 [33]. 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 [34]. 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 [35]. 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 fast loop holds what it measures. In several chapters the slow part that it did not measure went elsewhere.
- A slow, true measurement corrects a fast or inferred one, and its weight is set by the drift it must follow and the noise it must not pass.
- Speed is the fast handle of a continuous line, and an inventory lets it move.
- A slow, wide actuator returns a fast, limited one to its range, and it must be slow against its own delay.
- The structure of the loops decided more than their tuning.

References

[1] V. Vogl, M. Åhman, and L. J. Nilsson, "Assessment of hydrogen direct reduction for fossil-free steelmaking," *Journal of Cleaner Production*, vol. 203, pp. 736–745, 2018.

[2] A. J. Pimm, T. T. Cockerill, and W. F. Gale, "Energy system requirements of fossil-free steelmaking using hydrogen direct reduction," *Journal of Cleaner Production*, vol. 312, p. 127665, 2021.

[3] J. Jiménez, J. Mochón, J. S. d. Ayala, and F. Obeso, "Blast Furnace Hot Metal Temperature Prediction through Neural Networks-Based Models," *ISIJ International*, vol. 44, no. 3, pp. 573–580, 2004.

[4] O. Raza, N. Walla, T. Okosun, K. Leontaras, J. Entwistle, and C. Zhou, "Prediction of silicon content in a blast furnace via machine learning: A comprehensive processing and modeling pipeline," *Materials*, vol. 18, no. 3, p. 632, 2025.

[5] X. Liu, D. Han, Z. Li, H. Wang, W. Li, M. Fu, and X. Wang, "Time series forecasting model of mold level fluctuations integrating multi-scale feature decomposition and trend flexible alignment," *ISIJ International*, vol. 66, no. 7, pp. 867–875, 2026.

[6] H. Hadera, I. Harjunkoski, G. Sand, I. E. Grossmann, and S. Engell, "Optimization of steel production scheduling with complex time-sensitive electricity cost," *Computers & Chemical Engineering*, vol. 76, pp. 117–136, 2015.

[7] N. Depree, R. Düssel, P. Patel, and T. Reek, "The "virtual battery": Operating an aluminium smelter with flexible energy input," in *Light Metals 2016*. Wiley, 2016, pp. 571–576.

[8] P. Lavoie, S. Namboothiri, M. Dorreen, J. J. J. Chen, D. P. Zeigler, and M. P. Taylor, "Increasing the power modulation window of aluminium smelter pots with shell heat exchanger technology," in *Light Metals 2011*. Wiley, 2011, pp. 367–374.

[9] P. Cowling and W. Rezig, "Integration of continuous caster and hot strip mill planning for steel production," *Journal of Scheduling*, vol. 3, no. 4, pp. 185–208, 2000.

[10] D. A. Sala, A. Jalalvand, A. Van Yperen-De Deyne, and E. Mannens, "Multivariate time series for data-driven endpoint prediction in the basic oxygen furnace," in *Proc. 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2018, pp. 1419–1426.

[11] L. Bonechi *et al.*, "BLEMAB european project: Muon imaging technique applied to blast furnaces," arXiv preprint arXiv:2110.10327, 2022.

[12] Y. Isei, T. Kato, M. Osugi, and T. Ohta, "Development of a shape meter employing the LED dot pattern projection method for a hot strip finishing mill," *ISIJ International*, vol. 60, no. 12, pp. 2876–2885, 2020.

[13] Y. Miki, A. Matsui, H. Ohno, and N. Aramaki, "New electromagnetic flow control system for optimization of molten steel flow in continuous casting mold," *ISIJ International*, vol. 65, no. 8, pp. 1162–1171, 2025.

[14] H. Pauna, T. Kokkonen, P. Cavaliere, M. Bayat, A. Mirowska, T. Alatarvas, M. Huttula, and T. Fabritius, "Hydrogen direct reduced iron melting in an electric arc furnace: Benefits of in situ monitoring," *Journal of Sustainable Metallurgy*, vol. 11, no. 4, pp. 4655–4667, 2025.

[15] T. D. Bradarić, Z. M. Slović, G. J. Stepanoski, and S. Kosanović, "Static-dynamic model for endpoint control used in Smederevo's BOF shop," *Metallurgical and Materials Engineering*, vol. 27, no. 4, pp. 449–456, 2021.

[16] Z. Yang, X. Luo, and J. Qiao, "A combined adaptive entropy-TOPSIS and model predictive control strategy for mixed loading and delay operations in the reheating furnace," *Optimal Control Applications and Methods*, vol. 44, no. 5, pp. 2775–2795, 2023.

[17] Intergovernmental Panel on Climate Change, "Chapter 4: Metal industry emissions," in *2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 3: Industrial Processes and Product Use*. Geneva: IPCC, 2019. [Online]. Available: https://www.ipcc-nggip.iges.or.jp/public/2019rf/pdf/3_Volume3/19R_V3_Ch04_Metal_Industry.pdf

[18] F. Superchi, A. Mati, C. Carcasci, and A. Bianchini, "Techno-economic analysis of wind-powered green hydrogen production to facilitate the decarbonization of hard-to-abate sectors: A case study on steelmaking," *Applied Energy*, vol. 342, p. 121198, 2023.

[19] C.-J. Wong, A. A. Larkin, J. Bao, M. Skyllas-Kazacos, B. J. Welch, N. Ahli, M. Faraj, and M. Mahmoud, "Optimisation of power modulation for Hall-Héroult cells: Process operability and constraints as virtual energy storage," arXiv preprint 2511.07893, 2025.

[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. Ji, G. Li, L. Wei, and Z. Yi, "Modeling and determination of total heat exchange factor of regenerative reheating furnace based on instrumented slab trials," *Case Studies in Thermal Engineering*, vol. 24, p. 100838, 2021.

[23] Y. K. Lee, Y. J. Jang, and S. W. Kim, "Adaptive feed-forward automatic gauge control in hot strip finishing mill," *ISIJ International*, vol. 47, no. 10, pp. 1444–1451, 2007.

[24] 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.

[25] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, "Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators," *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[26] H. Wang, K. Zha, K. Yang, W. Zhang, S. Zhu, F. Xu, and Q. Xie, "Data-driven prediction and process parameter analysis of zinc coating thickness in hot-dip galvanizing," *ISIJ International*, vol. 66, no. 7, pp. 857–866, 2026.

[27] M. Hatano, Y. Matoba, K. Otsuka, M. Yoshiki, and T. Miyagi, "Automatic control of hot metal temperature of blast furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 7, pp. 534–542, 1982.

[28] C. Zhou, “Computational fluid dynamics (CFD) modeling for high rate pulverized coal injection (PCI) into the blast furnace,” Purdue University Calumet, final report to the U.S. Department of Energy, Hammond, IN, Tech. Rep., 2008.

[29] N. Thornhill and T. Hägglund, “Detection and diagnosis of oscillation in control loops,” *Control Engineering Practice*, vol. 5, no. 10, pp. 1343–1354, 1997.

[30] U.-S. Yoon, I.-W. Bang, J. H. Rhee, S.-Y. Kim, J.-D. Lee, and K. H. Oh, “Analysis of mold level hunching by unsteady bulging during thin slab casting,” *ISIJ International*, vol. 42, no. 10, pp. 1103–1111, 2002.

[31] Y. Wang, X. Wang, Z. Dong, and Z. Wang, “A Data-Driven Multiobjective Dynamic Robust Modeling and Operation Optimization for Continuous Annealing Production Process,” *ISIJ International*, vol. 60, no. 6, pp. 1225–1236, 2020.

[32] 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.

[33] 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.

[34] 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.

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

12.6 Exercises

12.1 Discussion · 8 min

After the pace change of Chapter 7 an agent proposes to double the pyrometer trim gain “so that the slabs leave hot enough”. Using Table 7.5 and Section 1.4.2, explain why this cannot work, which structure does, and what that structure still needs.

12.2 Discussion · 8 min

On the blast furnace of Chapter 2 an agent proposes to raise the mid-ranging gain from 0.2 to 1.0, at the same 12 h integral time, “to return the coal to its target faster after a burden change”. Using Table 2.6 and Section 1.4.3, say what the proposal will do, why, and what you would try instead to reduce the first cold dip.

12.3 Discussion · 8 min

On scenario 8 of Chapter 8, with thermal crown, an agent reports that the gaugemeter loop “held the gauge on setpoint through the whole coil” and recommends removing the X-ray monitor trim to save its maintenance. Using Equation (1.10) and Table 8.7, explain what the report missed and what evidence would settle it.

12.4 Discussion · 8 min

A vendor’s Smith predictor block for the coating loop of Chapter 10 takes its dead time in seconds. An agent sets it to the 40 s of the 150 m/min design speed. The line runs between 90 and 180 m/min. Using Table 10.5 and Section 1.4.5, say what will happen, and what the block needs.

12.5 Discussion · 7 min

For the twelve-heat challenge of Chapter 4 an agent proposes a 50 °C hot aim with no learning, “because learning adds a model that can be wrong, while a margin cannot”. Using Table 4.6 and Section 1.4.4, weigh the proposal.

12.6 Discussion · 9 min

An agent reports that the looper MPC of Chapter 9 cuts the strip off specification in scenario 11 by 19 percent, from 21.60 to 17.46 %, against the capped loop. Using Table 9.6 and Section 12.4.1, list the questions to ask before accepting it, and the check from the record that would have caught its weakness.