INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 17 Control of Energy Conversion and Storage

CHAPTER 14

Summary and Outlook

Pages 352–366

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control of Energy Conversion and Storage*, John D. Hedengren, Ed., vol. 17 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 14, pp. 352–366.

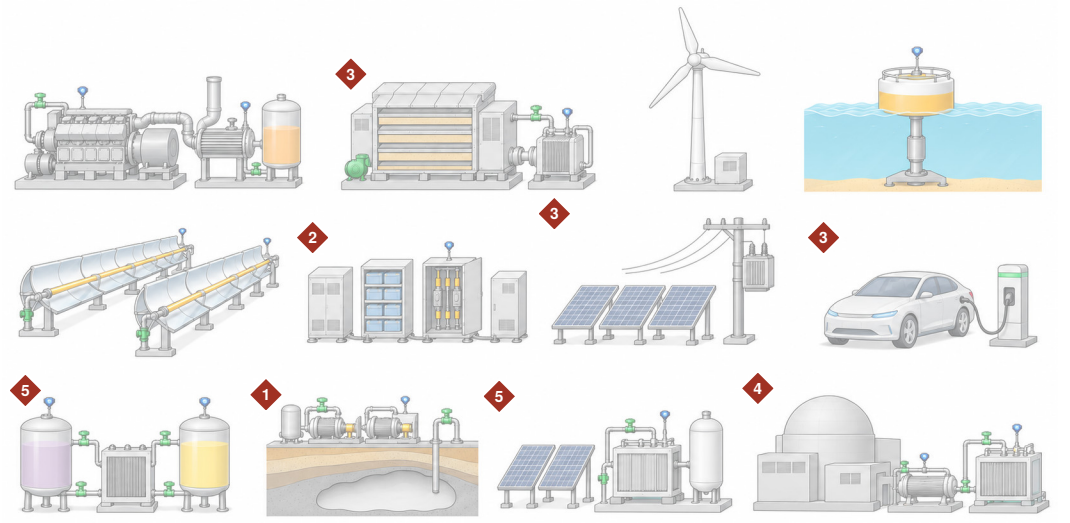
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v17-energy/

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

CHAPTER 14

Summary and Outlook



Open gaps 1 slow clock 2 inertia and meter 3 unmeasured limit 4 forecast error 5 silent degradation

Notice that the fuel cell, the feeder and the charger share one gap, and that no gap marks the gas engine, the wind turbine, the wave buoy or the trough field. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Energy Plants Taught**

The main idea. Each chapter of this book ran one plant that makes or stores energy in a simulator, from a gas engine to a cave full of air. Every plant had fast controls that act within seconds and a slow plan that looks hours or days ahead.

The challenge. A fast control can hold its reading steady while the plant wastes money or wears out. The limit that ends a run is often one that no instrument can see.

Why it matters. The plants that did best kept a store between the two layers, judged every plan by what was left in the store, and let safety rules act first.

Twelve chapters controlled twelve plants that convert or store energy, from a gas engine to a salt cavern. Each joined a fast regulating layer to a slow deciding one. Table 1.2 of Chapter 1 listed the two layers of every chapter, and Section 1.4.7 named the four ways they are joined: a store that lets them run at different rates, a setpoint written by the slow layer and held by the fast one, a limit or interlock that lets protection take the actuator from either, and a forecast that lets the slow layer act before the fast one sees anything. 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.

14.1 What the Chapters Showed

Tables 14.1 and 14.2 compare the chapters. The second column condenses Table 1.2. The other columns name an arrangement that worked, one that failed, and the measured result with the table or

Table 14.1 Fast Regulation and Slow Dispatch, Chapters 2 to 7

Chapter	Fast and slow layers	What worked	What failed	Measured result
2, gas engine CHP	Power in 2.3 s, recovered heat in 68 s. A hot-water buffer and a UPS between them.	Metered-load feedforward. A short tracking time. Fast jacket and supply loops in the challenge.	Feedback alone. Any tuning once demand exceeds two engines. A slow cooling loop that the buffer hides.	Import IAE 18.82 against 4.14 MW s (Table 2.7). Challenge UPS low 51.97 shipped against 79.19 % suggested (Table 2.8).
3, fuel cell CHP	Current on the inverter's 50 ms. Anode in 1.38 s, fuel arriving in 17.3 s.	A governor on the arriving fuel with the feedforward. The UPS paying for the wait. An estimate of the gas.	Load-to-current control, which leaves the inverter's foldback to stop the current. Feedforward alone.	Anode margin 0.157 naive against 0.623, power IAE 2.638 against 0.797 kWh (Table 3.6). Challenge passed on 31 of 31 seeds (Section 3.8).
4, wind turbine	Torque and pitch on one speed. Region logic and a gain schedule. A wind preview seconds ahead.	Torque below rated, pitch above, the pitch gain scheduled. A tight objective with a preview.	The schedule off, which holds speed tightly and rides the pitch rate limit. The challenge's weak gains.	IAE 3.058 with the schedule against 0.865 rad without it, pitch travel 108.4 against 255.2° (Section 4.6.2). MPC energy cost 2.6 to 5.2 % over ten seeds (Table 4.6).
5, wave buoy	Force law every 0.1 s. Sea state estimated over 120 s.	MPC on a forecast built from its own force estimate. A survival lock on the estimated sea.	A damper, which resists only motion already there. The MPC without survival in the storm.	Five minutes: damper 0.767, MPC 2.522 on its forecast and 2.713 kWh on the truth (Table 5.4). Challenge 10.889 kWh with no contact (Table 5.5).
6, trough field	Outlet loop on field flow. Setpoint optimizer and day schedule. A two-tank store.	A moving outlet setpoint on a broken day. A schedule into the evening peak. Tuned gains.	A fixed setpoint on a broken day. Gains alone or the schedule alone in the challenge.	Broken day 25.4 against 33.5 MWh, clear day 120.8 against 121.1 (Table 6.6). Challenge 3.2 shipped against 72.8 MWh suggested (Section 6.7).
7, campus microgrid	Virtual inertia and droops within milliseconds. Secondary, SOC manager and genset start over minutes.	A battery that forms the grid with 2 s of virtual inertia. Shedding. The genset started by charge.	A battery that follows the grid. Too little inertia. The one-hour dispatch on adjusted fuel.	Following: bus dead at 60 ms, 11.75 kWh critical lost. Forming: 2.48 Hz/s, nadir 59.516 Hz (Table 7.6). Challenge MPC adjusted 631.64 against 585.31 L (Section 7.7).

section of the chapter that reports it.

The columns make five points, and Figure 14.1 draws one pair of results from each row.

The fast layer holds what it measures, and the slow layer decides the value. In Chapter 10 the power loop reaches 63 % of a step in 18 s, and the chapter finds nothing in it to teach beyond how easy it is. On the day with one scarce hour the same plant booked −$2,242 under feedback and $14,599 under the dispatch layer (Table 10.8). In Chapter 6 the outlet loop held 393 °C through a broken day while the receiver went on losing heat, and a setpoint chosen by the optimizer raised net generation from 25.4 to 33.5 MWh (Table 6.6). On a clear day the same optimizer was worth 0.3 %. Through a single cloud the temperature trend said the controller succeeded and the energy trend said the plant did not. A fast loop that works is necessary. It says little about whether the plant is worth running.

Table 14.2 Fast Regulation and Slow Dispatch, Chapters 8 to 13

Chapter	Fast and slow layers	What worked	What failed	Measured result
8, solar feeder	Inverter curves in seconds. A tap changer on a 45 s timer and a central program.	Curves at the default slope, a compensator matched to the line, a slower band. MPC with a clear-sky preview.	A steep slope against a reactance-weighted compensator. The tap changer alone.	Tap operations 2 against 18 (Table 8.4). Challenge 0 of 6 seeds shipped, 6 of 6 suggested (Table 8.5). Case 9 out of band 21.2 against 0 min (Table 8.6).
9, fast charger	Four current limits in one selector. A charge to target over half an hour. The anode has no instrument.	The anode map in the selector with a margin and a preload. The coolant loop on temperature.	Constant current, which plates unseen. No preload. A margin smaller than the table's error.	To 80 %: 16.84 min and 3.381 Ah plated against 21.34 min and none (Table 9.5). Margin covers a table about 20 % optimistic (Table 9.7).
10, flow battery	Power loop in 18 s. A 24-hour dispatch program with a bounded net-load trim.	Feedback on a clean day. Dispatch on a day with one scarce hour. Feedforward flow.	Feedback through a price it cannot measure. Dispatch judged on regulation.	Clean day $35,083 against $35,067, worst ramp 2.61 against 5.55 MW/min. Scarcity day −$2,242 against $14,599 (Table 10.8).
11, air cavern	Grid trim and heat loops hourly. Inventory policy and seasonal reserve over months.	The heater armed. A rolling plan with a reserve. The operator curve, which holds the most gas.	The threshold rule and the plan with the heater off. A floor in megawatts computed at the present scan.	Heater off: default 0.56 for the rule, 2 MWh for the curve, 0.47 GWh for the plan (Table 11.8). Challenge 1.11 GWh default shipped, none suggested (Table 11.6).
12, solar hydrogen	Rectifier loops in 1.5 s. A state machine and an evening schedule. Gas stores behind interlocks.	Interlocks choosing the mode and economics the setpoint. Stored oxygen. Structure and tuning together.	Tuning alone or structure alone in the challenge. The economic schedule judged on net margin.	Three cloudy days: firm $16,508, economic $15,656, following $15,468 (Table 12.8). Challenge −$44,664 shipped, $5,467 tuned only, $6,959 suggested (Section 12.7).
13, nuclear heat and H_2	Six PI loops on the header, a 36 s scan. An economic plan every 15 min over a day.	Price thresholds on an ordinary day. The economic plan on a day with a spike.	A fixed split. Re-planning on a forecast blind to the spike, which prices the forecast only.	Ordinary day $470,311 fixed, $516,300 thresholds, $529,303 MPC. Spike day $709,818, $924,749 and $1,010,196 (Table 13.5).

A store lets the layers run apart, and its state enters every comparison. Every plant in the volume has a store between its layers, and the store changes what a comparison means. In Chapter 2 the hot-water buffer hid a slow cooling loop for eight minutes. In Chapter 3 the UPS paid for the fuel's 17.3 s wait. In Chapter 7 the one-hour dispatch burned 397.47 against 476.30 L of diesel, but it ended with the battery at 62.18 against 77.98 %, and priced at the replenishing fuel it used 465.02 against 476.30 L. A plan with only one slot recovered almost all of that (Table 7.8). In Chapter 13 the thresholds left the tank at 59.3 against 35.1 %, worth $5,693 if sold overnight and $21,000 at the next evening's peak, enough to reverse the ranking (Section 13.6). A dispatch comparison without the end state of its store is not a comparison.

A forecast pays where feedback cannot see the cause. On a clean day the flow battery's forecast was worth nothing, $16 less than feedback, because the highest net load and the highest price fall in the same hour. One scarce hour at three in the afternoon separated them by $16,841. The nuclear plant's thresholds captured 78 % of what its optimizer added on an ordinary day, and the optimizer

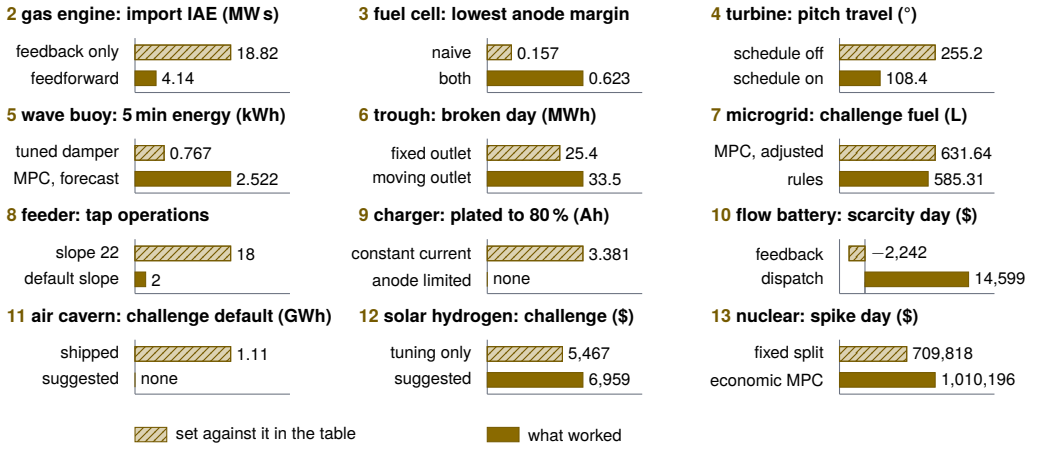


Figure 14.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 3, 5, 6, 10, 12 and 13, and a smaller one in the others. The solar hydrogen panel sets the tuning alone, which falls short of the challenge's $6,000 target, against the suggested answer. Notice that in most rows the better result came from a change of structure, measurement or information, not of tuning alone.

booked $85,447 more on a day with an evening spike (Table 13.5). In Chapter 11 a pair of price thresholds booked $0.77 M more than the rolling plan, because the price is affine in the net load that the plan forecasts. At a shorter scale the turbine's preview supplied 2.1 of the 5.3 points of its gust advantage, and the rest came from a tighter objective that a shipped option of the loops matched on that gust (Section 4.6.5).

Protection and the unmeasured limit act below every layer. The limits that ended a run were rarely the ones on a gauge. On the fast charger the charger ceiling and the anode map held the current for all but 0.32 min of a charge to 95 %, so two of the four limits were decorative, and the one that mattered has no instrument. A constant-current charge looked identical on every instrument while it plated 3.381 Ah, and 13.867 Ah from a cold pack (Table 9.5). In Chapter 8 the band, the timer, the tap clamp and Volt-Watt act below either central controller. Through the worst cases the application can pose no run curtailed solar or stepped the tap faster than its timer, and in a transmission sag both central structures lost to the local devices because the band's timer set the loss (Table 8.7). The buoy passed its storm only with a survival lock on the estimated sea, and in Chapter 12 the interlocks choose the mode before any economics may choose a setpoint.

Tuning repairs the fast layer, and structure repairs the slow one. The challenges separated the two kinds of change. In the trough field's challenge the tuned gains alone generated 68.7 MWh and booked $2,663, while the schedule with the weak gains booked $4,751 on 41.1 MWh. Together they generated 72.8 MWh and booked $6,780 (Section 6.7). The gains recovered the energy and the schedule earned the money. In Chapter 12 neither tuning alone nor structure alone passed. On the flow battery the structural decisions were worth $35,033 and 56.2 MWh of pumping, and no gain change in the default structure came close. In the gas engine's challenge each loop guarded a different part of the target, so either one tuned alone failed the other part (Table 2.8). A fast loop is tuned against its disturbance. A slow layer is chosen against the question the plant is paid to answer.

14.2 Choosing a Method

The chapters suggest an order of questions for an energy plant that is new to the reader. Each question names the chapter where its answer decided the result, and Figure 14.2 draws the list as a chart.

1. *Which clocks does the plant run on?* Place its loops and layers on one axis, as Figure 1.3 does, and check that each layer is several times slower than the one below it.
2. *What makes the voltage or the frequency when the grid is gone?* A source that forms the grid, with enough inertia for the meter that protection uses, comes before any dispatch (Chapter 7).
3. *Is the binding limit measured?* If not, infer it and size the margin by the model error it must cover, as the charger's 15 mV covers a table 20 % optimistic (Chapter 9). Find which limit carries the plant before tuning the others.
4. *What does the fast layer wait for?* A fuel that arrives 17.3 s late or heat that answers in a minute needs a store to pay for the wait and a governor that admits only what can be supplied (Chapter 3, Chapter 2).
5. *What is the store's state worth at the end of the horizon?* Give the plan a terminal value, such as the merit-order water value of Equation (1.9), and judge every comparison with the end state priced (Chapter 10, Chapter 7).
6. *Does the right action depend on something the fast layer cannot measure?* If it does, as a price away from the load peak does, plan on a forecast. If it does not, feedback is as good and regulates better (Chapter 10, Chapter 11).
7. *Which layer must lose?* Write the interlocks, overrides, shedding and survival modes first, and let the economics act only inside the mode they leave (Chapter 12, Chapter 5).

A predictive controller paid where it held information or a model the loops did not have. On the buoy it captured 2.522 kWh in five minutes on its own forecast against 0.767 kWh for the tuned damper (Table 5.4). On the feeder's cloudy day at 150 % solar it held every node in band where the program spent 21.2 min out, and its preview of the clear-sky rise took it from 15.9 to 0 min (Table 8.6). Its advantage often belonged to the information more than to the horizon. The fuel cell's governor given the same estimate of the gas reached a peak utilization of 0.805 against the MPC's 0.807 (Table 3.8), and a trough loop with a feedforward from the same pyrheliometer did as well through passing clouds, 45.7 against 46.8 °C min (Table 6.7). Prediction cost something elsewhere. The turbine paid 2.6 to 5.2 % of its energy in turbulence, the charger took 3.22 min longer in scenario 12 to keep its core cool (Table 9.8), and the microgrid's dispatch lost on adjusted fuel. Before reaching for prediction, give the loops the measurement or the estimate the predictive controller would use. Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Inverters that form the grid. Power systems in Hawaii, South Australia, Tasmania, Texas and Ireland commonly run with inverter-based sources supplying more than 50 to 60 % of demand at an instant [1]. A survey of grid-forming inverters describes the concept as drawing increased attention from academia and industry, takes recently commissioned projects in Australia, the United Kingdom and the United States as examples of the trend, and names the transition between grid-connected and standalone operation among the open challenges [2]. On the campus of Chapter 7 a battery that followed the grid left the bus dead 60 ms after the fault, and the same battery forming it held the island. In the editor's view the virtual inertia of an inverter will be set, documented and audited like a governor's droop, together with the meter that protection reads it through.

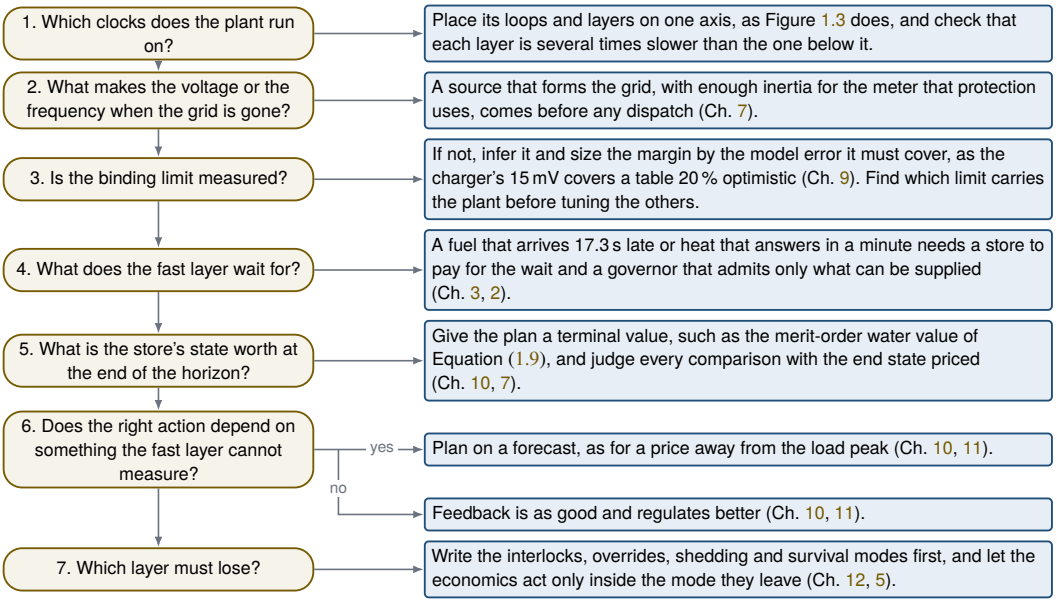


Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the grid's own voltage and the unmeasured limit come before any question about the plan.

> ⧖ HISTORY **2014 to 2020: Inverters Asked to Hold the Grid**
>
> The second history box of Chapter 1 ended with an amendment of 2014 that let inverters support the voltage of the grid they feed. In 2018 a European guidance document asked generating units to stay connected while the frequency changed by up to 2 Hz/s over a 500 ms window, 1.5 Hz/s over one second and 1.25 Hz/s over two [3]. In 2020 a research roadmap saw early grid-forming inverters emerging in islanded microgrids with high shares of renewable generation, and judged that replacing synchronous machines with them at the transmission level would take about 10 to 30 years and a research base of protection, controls and interoperability [4]. In the editor's reading the droop of the early 1990s has returned as the inverter that forms the grid, and the campus of Chapter 7 shows why the window of its meter matters.

Preview in the fast layer. Lidar lets a turbine controller act in anticipation of the incoming wind rather than after it has struck the rotor. A review of lidar-enhanced control reports its use for load reduction through pitch and for energy through yaw, and ongoing work that combines pitch and torque in a feedforward nonlinear predictive controller to reduce loads and raise energy at once [5]. Wave energy has not reached commercial viability, and a review of its control holds that control technology can play a major part in getting there [6]. The buoy of Chapter 5 earned 88.5 % of what a 4 s perfect forecast adds within its first 2 s (Section 5.6), and the turbine's preview supplied about two fifths of its gust advantage. In the editor's view a claim for a previewing controller should report the preview's share apart from the share of its objective.

One store on two clocks. A battery used for peak shaving and frequency regulation together, in a joint optimization with degradation and uncertain load and regulation signals, cut a user's bill by up to 12 % on real data. The joint saving was often larger than the sum of the two services alone, and a simple threshold rule in real time achieved it [7]. The flow battery of Chapter 10 is wired the same way, a day-ahead program with a net-load trim beneath it. In the editor's view the stacking of fast and

slow services on one store will make its end state, not its profit on one day, the quantity a dispatch is judged by.

Charging on an estimated limit. Charging an electric vehicle in fifteen minutes is a goal that present lithium-ion cells with graphite anodes cannot reach without harming their performance and safety [8]. Optimal control on a coupled electro-thermal-aging model has been used to trade charge time against degradation under electrical and thermal constraints, with experiments comparing a balanced protocol with a traditional one [9]. On the charger of Chapter 9 a pack whose resistances had risen 60 % plated 3.412 Ah on a strategy that plated nothing when the pack was new. In the editor's view a charging protocol will be delivered with the model error its margin covers, and with an estimate of aging that keeps that error inside the budget.

Electrolysis as a flexible load. Water electrolysis is expected to couple electricity with mobility, heating and chemicals, and a review compares alkaline, proton exchange membrane and solid oxide electrolyzers on the properties that decide flexible operation: load range, load gradients, start-up time and stand-by losses [10]. The two hydrogen chapters meet those properties as control constraints. The solar plant's suggested challenge run paid $4,400 in starts, and the nuclear plant's optimizer held its electrolyzer at its 18 MWe turndown through the evening peak. In the editor's view the dispatch of electrolyzers will be limited as much by starts, turndown and gas purity as by the price of electricity.

Data centers on their own generation. United States data centers used an estimated 192 TWh of electricity in 2024, 4.7 % of the national total, and a reference case puts them at 649 TWh and 11.8 % in 2030, a third of all load growth over those six years. The same report asks that the variability of their demand be characterized, to learn how the grid or generation on site can serve it, and that demand flexibility be explored [11]. Fuel cells that power data centers are slow to follow a sudden rise in demand because of fuel delivery. Capping the servers limits their performance, and storage alone must be greatly oversized for the largest surge [12]. The engines of Chapter 2 and the fuel cell of Chapter 3 both carry such a load and pay for the wait with a UPS. In the editor's view the ramp a data center may ask of its own generation will become a stated term of its supply, set with the fuel's delay in view.

Hydrogen in the pipeline gas. Natural gas utilities have begun to blend hydrogen into their existing pipelines. Hydrogen carries less energy per volume, so at a blend of 20 % the volumetric fuel flow must rise by 15 % to deliver the same energy. On a stoichiometric engine whose controller had no warning of the change, sudden steps in hydrogen swung the air-fuel ratio and released bursts of nitrogen oxides, and better-tuned controllers responded faster and emitted less [13]. The fuel cell plant of Chapter 3 meets the same disturbance as a fall in methane content, and an analyzer read from the start of the run cut the energy drawn from its UPS from 48.37 to 36.50 kWh (Table 3.6). In the editor's view a reading of gas quality, fed forward, will become part of the fuel system of every engine and fuel cell that burns blended gas.

Estimates in place of meters. A study of best practices at concentrating solar plants found the heat transfer fluid's flow meters the one noted exception to otherwise reliable instruments. Plants estimate the flow from pump current, drive speed and pressure drop instead, two plants were pleased with newer schemes for automatic flow control, and the study asks for performance forecasts from irradiance data at 5 to 15 min rather than hourly, because hourly data miss short cloud transients [14]. The outlet loop of Chapter 6 acts through that same field flow, and its passing clouds are the transients the study names. In the editor's view such an estimate will be judged, like the charger's anode map, by the error it carries. Figure 14.3 sets developments of this section and of the volume's chapters on three lanes.

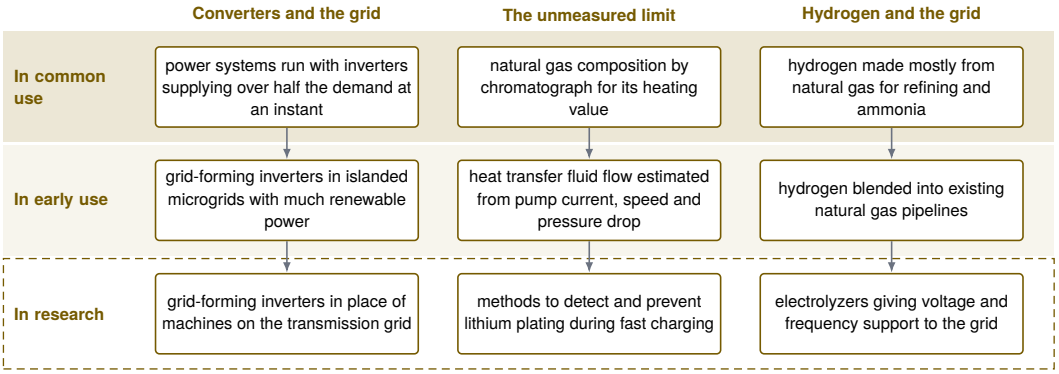


Figure 14.3 Developments in three threads. The top lane holds what the cited source calls standard or typical practice, or reports in common operation, the middle lane what it reports in service at some sites or in early use, and the bottom lane what it reports from simulation, from laboratory projects or as a need, in the editor's reading. Sources by thread: converters and the grid [1, 4], the unmeasured limit [14–16] and hydrogen and the grid [13, 17]. Notice that each thread moves from what a plant measures or follows toward what it must estimate or hold for others.

14.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [18]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

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

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

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

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a

change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

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

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [19]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

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

14.4.1 How This Volume Was Made

Each of the twelve 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. Four of the twelve referee reports call themselves self-reviews. The other eight describe a reviewer's stance or say only that they were written against the draft and then applied, and none names a reviewer independent of the writing. Six carry a second pass on predictive-control material added later. Eight of the papers were measured again in a consistency round and an app-fix round, and the other four were written after those rounds. The defect logs hold 131 entries with a severity. Of these, 42 are wrong output or behavior of an app or its model, 70 documentation or briefing content, 17 cosmetic or display, and 2 something else, a completion rule and a check. Some round notes list fixes without a severity, so the count is a floor. 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. Four cases show what went wrong and what caught it. In each a fast part met a slow clock or a slow layer misjudged a fast one.

A heater that never fired. The air cavern's supplemental heater had a lag of about six minutes, and the engine integrated it with a fourth-order Runge-Kutta step of thirty minutes. The step was 5.21 times the time constant, against a stability limit of 2.785, so the heater's state diverged on every step and a clip returned it to zero. Over the 150-day challenge the command reached 120 MW and integrated to 71 048 MWh while the heat delivered was exactly zero. The trend plotted the command, so the strip showed a working heater, and only a badge on the scene read 0 MW. The agent writing the paper found it, and a consistency round found the expander answering in about 1.8 h where its label said 12 min for the same reason. The app-fix round replaced the lags with their exact solution over the step, the challenge then burned 26.0 GWh of heat, and every run that moved a machine or fired the heater moved.

A program that counted its offsets twice. The feeder's central program fitted its voltage estimate with the inverter offsets in force, predicted the effect of new offsets on top of that estimate, and then replaced the old offsets with the new ones, so the old ones counted twice. On a cloudy day the summed offset alternated six times between 10:00 and 14:00. The stage that added the predictive controller beside the program found it, and a regression test now fails on the old code with six alternations and passes with none. The same round found that with the curves off the inverters run at

unity power factor while the program planned offsets for them. On the challenge as shipped it asked for up to 2.43 Mvar over 10 500 s, and the inverters delivered none. The first fix moved the program's case 9 from 27.3 to 21.2 min out of band, and the second moved the challenge as shipped from 143.7 to 30.1 min. The slow layer had modeled the fast devices doing what they could not.

A law that always matched the plant. The turbine's torque law below rated, $K\omega^2$, took its constant from the run's own air density and rotor radius. At 6 m/s the loop still settled at a tip-speed ratio of 7.53 to 7.54 for densities of 1.0 and 1.4 and radii of 60 and 70 m, so the Model dialog's warning that the law is right only if K is right could not be shown. A truth audit found it. The law now takes the declared design, a density of 1.0 gives a ratio of 7.02, a 70 m radius gives 8.87 and 3.9 % less power, and no shipped run moved. A test fails on the old code. Nine entries in six of the twelve logs record a controller or an operator aid that read the plant's truth or the run's hidden settings in place of a measurement or a declared value.

A re-plan that never saw the price. The nuclear plant's paper, and five texts of its app, said that re-planning every fifteen minutes puts the metered price into the objective, so that on the spike day the optimizer discharges at its limit once the meter shows the spike. The chapter agent compared the logged allocations of the spike day with those of the ordinary day and found them identical at every sample. The program builds its block prices from the forecast alone. The chapter now says that feedback corrects the state and not the price, and the five texts were logged with a recommendation to change the texts rather than the controller, so that the paper's numbers for the case stand. A repair round has since changed them.

The book round. The twelve chapter records list 156 rows of numeric verification, each a stated value or group of values set beside its reproduction. Seven records report no drift. The other five name a printed value that a run does not reproduce exactly, and all stand as printed. Every record tested its chapter's causal claims by run. Three explanations were contradicted and corrected without a new number. One is the nuclear case above. The flow battery's idle-day loss was said to be curtailment and nothing else, when curtailment is \$766 of the \$1,468 and the idle plant also buys 15.42 MWh for its pumps, auxiliaries and inverter. The cavern's heater was said to fire at its full 120 MW through the first days of the drought, when its command alternated on successive hourly scans, a limit cycle the repair round removed. Two summaries were aligned with their own tables, a feeder program said to win a challenge in which neither run holds the band and an economic schedule said to win where firm dispatch earns the most. The turbine's energy cost, one value in the paper's abstract, is the ten-seed range of 2.6 to 5.2 % in the chapter. Nine of the twelve chapter passes logged an app defect, most of them in briefings and helpers, such as a challenge hint that said tuning alone passes when it nets \$5,467 against a \$6,000 target. No chapter was found naming a loop's action against the convention of Section 1.6.2. Chapter 1 was written from the converted chapters, and seven of its statements, in its rows for five chapters, were corrected from the chapter reports, among them the turbine's fastest clock and the fluid in the trough field. One finding of the round is still open, and Section 14.5 describes it.

Every error above was caught by a check that could fail: a delivered heat compared with its command, a decision count on a cloudy day, a hidden value moved while the measurements were held, a logged plan compared with the plan of another day. None was caught by rereading the text that contained it. Four working rules follow. Integrate every fast lag at the step that will run it, and check each loop's gain at the scan it is given. Check that what a slow layer asks of a fast device the device can deliver. Give a controller only what the plant measures and the design declares. Test the event a plan claims to answer by comparing its moves with and without the event.

14.5 Gaps in the Control of Energy Conversion and Storage

Five gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 14.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 14.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *A fast loop on a slow clock.* In Chapter 11 the heater loop runs on the dispatch layer's hourly scan. The heater's six-minute lag is over within the scan, so the loop sees a static process. At the shipped gain of 1.6 MW/K its discrete loop gain was above the stable limit, and the heater command flipped between 0 and 120 MW every other scan. No tile and no number of the paper showed it. The chapter agent found it in the hourly rows of the exported record. A repair round set the gain to 0.2 MW/K, which settles, and every heater-armed number of the paper and the chapter was measured again. Supervisory and lower-level loops run continuously with cycle times from milliseconds to minutes [20], and a study of hybrid nuclear and renewable plants held that techniques are needed to coordinate linked but independent systems with differing dynamic timescales [21]. The gap is hard because a supervisory layer chooses its scan for its own decisions, and a loop placed on it inherits that period without a check of its gain. What stays open is the tooling. Tools that check every loop's gain at the scan a supervisory layer imposes on it are not routine.

2. *Inertia stated with its meter.* The microgrid's first swing read 2.48 at the shipped 20 ms step and 2.75 Hz/s at 2.5 ms, and the virtual inertia below which the solar trips moved from between 1.25 and 1.3 s to between 1.6 and 1.7 s once the meter was resolved (Section 7.6). European guidance states the rate of change of frequency that generating units must ride through separately for each averaging window, from 2 Hz/s over 500 ms to 1.25 Hz/s over two seconds [3], and a report on frequency measurement from the same body notes that the shorter the time for a frequency calculation, the noisier the measurement [22]. Reduced inertia gives a larger rate of change of frequency [4]. The gap is hard because an inverter's first swing is over within the window that a relay averages, so the same inertia passes or fails with the meter. A method that sizes virtual inertia against the protection's own meter is missing.

3. *Estimating the limit that has no instrument.* The margin of Chapter 9 covered a table 20 % optimistic and no more, and an aged pack plated as much as the reckless charge. In Chapter 8 the program's margins on the worst customer were a median of 6.4 times the error and too wide to steer by on a cloudy day. The fuel cell of Chapter 3 runs on an estimate of the gas arriving at its anode, and its governor matched the predictive controller only when given the same estimate, 0.805 against 0.807 (Table 3.8). Estimators that track an anode potential, an arriving fuel or a far-end voltage and state their own uncertainty online are not standard. The Open challenge box after this list takes the gap further.

4. *Forecasts with a stated uncertainty.* The plan of Chapter 13 priced every block from a forecast that missed a spike of 5.5 times, and no re-plan could recover a storage decision taken six hours earlier. A roadmap for grid-forming inverters lists the reserve requirements that keep frequency within bounds as an open question, and notes that capacity held back as headroom can be an opportunity cost [4]. A study of solar plants asks for forecasts from data at 5 to 15 min, because hourly data miss short cloud transients [14]. The gap is hard because a reserve against the tail costs money on every ordinary day and pays on a few. A dispatch layer that carries the forecast's error and prices a reserve against it is not standard.

5. *Degradation that no loop reports.* The flow battery of Chapter 10 lost 54 MWh of working band to an imbalance while its power and flow loops held their setpoints. In Chapter 12 a membrane defect raised its first alarm 4.6 hours after it appeared, because the crossover is a rising fraction of

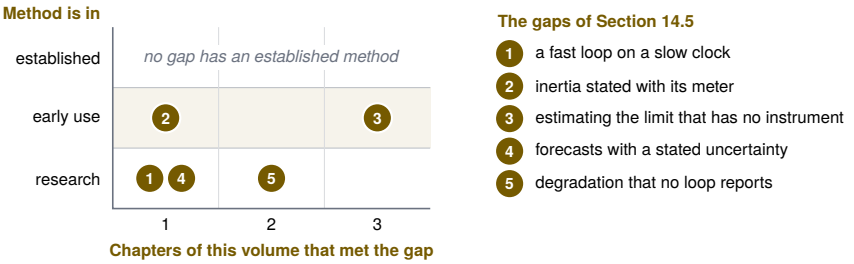


Figure 14.4 The five gaps of this section. Across is the number of chapters that each item names as evidence: one for gap 1 (Chapter 11), one for gap 2 (7), three for gap 3 (3, 8 and 9), one for gap 4 (13) and two for gap 5 (10 and 12). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that no gap has a method in established use.

a falling flow. In simulations of a wave energy converter, heavy biofouling cut the take-off power by 15 % in one sea and retuning the control recovered only about 6 % of the loss, and the report calls for adaptive control and models that account for environmental degradation [23]. Flow batteries promise lives of 20 years or more, yet few United States projects have stayed online for several years to prove it at scale [24]. The gap is hard because lost capacity drifts slowly, the loops absorb it, and its trend looks like a change of load. Monitoring that normalizes a reading by the load that the dispatch sets, and reports lost capacity as a measured quantity, is not routine.

⚠ **OPEN CHALLENGE** **Estimating the Limit That Has No Instrument**

Gap 3 is the gap most chapters met. A gap assessment of fast charging held that methods are needed to detect and prevent lithium plating, and that voltage and temperature sensors on every cell would identify the weaker and hotter cells [16]. Trough plants with unreliable flow meters estimate the flow from pump current, drive speed and pressure drop, with accuracy typically not nearly as good as a working meter's [14]. An online gas chromatograph can take minutes per analysis [15]. The gap is hard because the estimate is checked only when the limit is crossed, and its model ages with the plant. Progress would be an estimator that reports its own uncertainty online and a margin sized from that report.

14.5.1 Gaps Common to Agent-Era Practice

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

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [25]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [19]. In

the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [26]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [27]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [28]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A fast loop that works is necessary. It says little about whether the plant is worth running, which the slow layer decides.
- A store lets the layers run apart, and a dispatch comparison without the end state of its store is not a comparison.
- A forecast pays where feedback cannot see the cause, and is worth little where the price follows the load it measures.
- The limits that ended a run were rarely the ones on a gauge.
- Tuning repairs the fast layer, and structure repairs the slow one.

References

[1] J. Matevosyan, B. Badrzadeh, T. Prevost, E. Quitmann, D. Ramasubramanian, H. Urdal, S. Achilles, J. MacDowell, S. H. Huang, V. Vital, J. O'Sullivan, and R. Quint, "Grid-forming inverters: Are they the key for high renewable penetration?" *IEEE Power and Energy Magazine*, vol. 17, no. 6, pp. 89–98, 2019.

[2] D. B. Rathnayake, M. Akrami, C. Phurailatpam, S. P. Me, S. Hadavi, G. Jayasinghe, S. Zabihi, and B. Bahrani, "Grid forming inverter modeling, control, and applications," *IEEE Access*, vol. 9, pp. 114 781–114 807, 2021.

[3] ENTSO-E, "Rate of change of frequency (RoCoF) withstand capability: ENTSO-E guidance document for national implementation for network codes on grid connection," European Network of Transmission System Operators for Electricity, Brussels, Guidance document, Jan. 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20RfG/IGD_RoCoF_withstand_capability_final.pdf

[4] Y. Lin, J. H. Eto, B. B. Johnson, J. D. Flicker, R. H. Lasseter, H. N. Villegas Pico, G.-S. Seo, B. J. Pierre, and A. Ellis, "Research roadmap on grid-forming inverters," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5D00-73476, 2020.

[5] A. Scholbrock, P. Fleming, D. Schlipf, A. Wright, K. Johnson, and N. Wang, "Lidar-enhanced wind turbine control: Past, present, and future," in *Proceedings of the 2016 American Control Conference (ACC)*, 2016, pp. 1399–1406.

[6] J. V. Ringwood, "Wave energy control: Status and perspectives 2020," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 12 271–12 282, 2020.

[7] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2882–2894, 2018.

[8] Y. Liu, Y. Zhu, and Y. Cui, "Challenges and opportunities towards fast-charging battery materials," *Nature Energy*, vol. 4, no. 7, pp. 540–550, 2019.

[9] H. E. Perez, X. Hu, S. Dey, and S. J. Moura, "Optimal charging of li-ion batteries with coupled electro-thermal-aging dynamics," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 9, pp. 7761–7770, 2017.

[10] A. Buttler and H. Spliethoff, "Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: A review," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 2440–2454, 2018.

[11] S. J. Smith, A. Hubbard, A. Newkirk *et al.*, "United states data center energy usage report: 2025 update," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-2001758, 2026.

[12] Y. Li, D. Wang, S. Ghose, J. Liu, S. Govindan, S. James, E. Peterson, J. Siegler, R. Ausavarungnirun, and O. Mutlu, "Techniques for efficiently handling power surges in fuel cell powered data centers: Modeling, analysis, results," arXiv preprint arXiv:1806.04506, 2018. [Online]. Available: https://arxiv.org/abs/1806.04506

[13] N. Katsampes, D. Montgomery, G. Arney, and D. B. Olsen, "Hydrogen-natural gas fuel blending in a "rich burn" engine with 3-way catalyst," *Frontiers in Fuels*, vol. 2, p. 1416716, 2024.

[14] M. Mehos, H. Price, R. Cable, D. Kearney, B. Kelly, G. Kolb, and F. Morse, "Concentrating solar power best practices study," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5500-75763, 2020.

[15] 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.

[16] S. Ahmed, I. Bloom, A. N. Jansen, T. Tanim, E. J. Dufek, A. Pesaran, A. Burnham, R. B. Carlson, F. Dias, K. Hardy, M. Keyser, C. Kreuzer, A. Markel, A. Meintz, C. Michelbacher, M. Mohanpurkar, P. A. Nelson, D. C. Robertson, D. Scoffield, M. Shirk, T. Stephens, R. Vijayagopal, and J. Zhang, "Enabling fast charging – a battery technology gap assessment," *Journal of Power Sources*, vol. 367, pp. 250–262, 2017.

[17] U.S. Department of Energy, "Department of energy hydrogen program plan," U.S. Department of Energy, Washington, DC, Report, 2020.

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

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

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

[21] M. F. Ruth, O. R. Zinaman, M. Antkowiak, R. D. Boardman, R. S. Cherry, and M. D. Bazilian, "Nuclear-renewable hybrid energy systems: Opportunities, interconnections, and needs," *Energy Conversion and Management*, vol. 78, pp. 684–694, 2014.

[22] ENTSO-E RG-CE System Protection & Dynamics Sub Group, "Frequency measurement requirements and usage," European Network of Transmission System Operators for Electricity, Brussels, Report Final Version 7, 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf

[23] T. Brekken, "Applied research and development to support open-water testing at PacWave, task 7: WEC system control for marine growth and electrical faults," Oregon State University, Corvallis, OR, Final Technical Report DOE-OSU-EE0009969-7, 2025.

[24] A. Hollas, A. Tuan, V. Viswanathan, and I. Ragazzi, "Adoption readiness level assessment of redox flow batteries," Pacific Northwest National Laboratory, Richland, WA, Report PNNL-36780, 2024.

[25] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[26] International Electrotechnical Commission, “Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels,” IEC 62443-3-3:2013, Edition 1.0, 2013.

[27] W. Zhao, J. P. Queralta, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: A survey,” in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

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

14.6 Exercises

14.1 Discussion · 8 min

An agent reviewing Chapter 10 proposes to retire the dispatch layer and run the battery on the net-load controller alone, “because on the clean day it books the same margin and regulates twice as well”. Using Table 10.8 and Table 13.5, say what the proposal gets right, what it misses, and what you would keep.

14.2 Discussion · 8 min

An agent reports that the one-hour energy dispatch of Chapter 7 “saves 78.83 L of diesel against the rules on the added case”. Using Table 7.8 and the end-of-day tank in Table 13.5, explain what the claim leaves out and what a fair comparison of two dispatch policies on a plant with a store must report.

14.3 Discussion · 8 min

On the fast charger of Chapter 9 an agent proposes two simplifications: remove the voltage and power limits from the selector, “since they hold the current for at most 0.32 min”, and keep the 15 mV anode margin, “since it plated nothing in scenario 4”. Using Tables 9.5, 9.7 and 9.8, say what is wrong with each and what the charger’s design should state.

14.4 Discussion · 8 min

An agent proposes to close the air cavern’s heater defect of Section 14.5 by halving the shipped gain of 1.6 MW/K. Using Section 1.4.2, explain the cause, judge the proposal and say what else must happen first.

14.5 Discussion · 7 min

Chapter 8 closes with “the slow device should set the level and the fast ones should trim the spread”. Using Tables 8.4, 8.5 and 8.6, explain the sentence, give a result where the layers fought, and say what the central controllers add and cannot override.

14.6 Discussion · 9 min

An agent reports that its new dispatch layer for one plant of this volume “passes all scenarios and raises the margin 4 %”. Choose the plant and, using Section 14.4.1, name four checks to run before accepting it and the result of each that would reject it.