INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 17 Control of Energy Conversion and Storage

CHAPTER 13

Supervisory Economic Dispatch of Nuclear Heat, Storage and Hydrogen Production

Pages 323–351

Edited by John D. Hedengren

How to cite this chapter. "Supervisory Economic Dispatch of Nuclear Heat, Storage and Hydrogen Production," 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. 13, pp. 323–351.

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

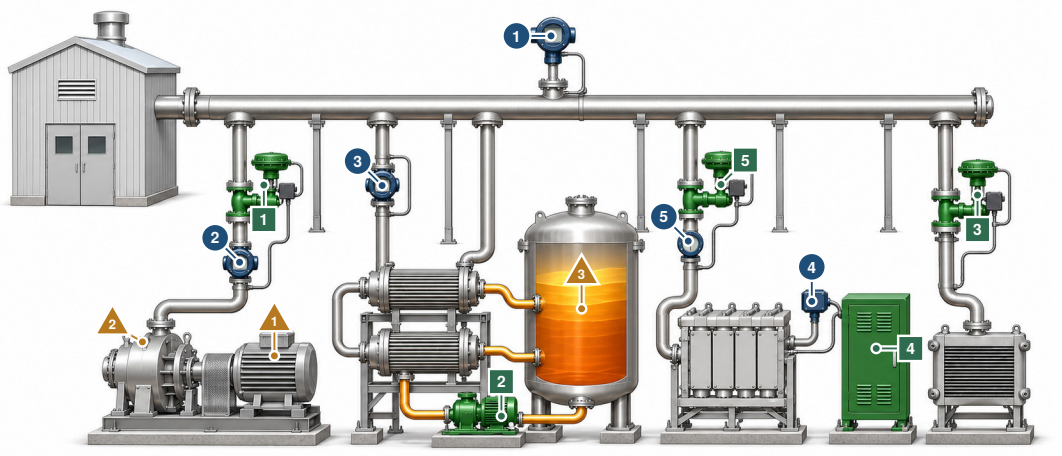
The app for this chapter
https://apmonitor.com/apps/nuclear-hydrogen/

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

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

CHAPTER 13

Supervisory Economic Dispatch of Nuclear Heat, Storage and Hydrogen Production



The controller moves 1 turbine throttle valve 2 salt pump, charging or discharging the store 3 dump valve to the condenser 4 electrolyzer power supply 5 superheater on the stack steam

The controller reads 1 steam header pressure 2 steam flow to the turbine 3 heat into or out of the store 4 electrolyzer power 5 stack temperature

What disturbs it 1 electricity price and its forecast 2 turbine outage 3 smaller usable tank

Notice that every branch leaves one header fed by steady reactor heat, so whatever the turbine, the salt store and the electrolyzer do not take goes to the dump condenser at the right. Illustration, not to scale.

IN PLAIN TERMS **Steady Reactor Heat, Sold at the Best Hour**

The main idea. A nuclear reactor makes the same heat all day, but the price of electricity changes every hour. The plant can sell the heat as electricity, keep it in a tank of hot melted salt, or use it to make hydrogen fuel.

The challenge. All of the heat must go somewhere at every moment, and heat that nobody buys is thrown away. On a normal day, a plan that looks a whole day ahead earns only a little more than three simple price rules.

Why it matters. Knowing on which days a smarter planner pays off keeps a plant from buying one for the wrong reason.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of the Introduction, the two history boxes of Section 13.2, the fact and try-it boxes of Section 13.6, and the key points
Study	55 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 13.6, Section 13.7, Section 13.8 and the open challenges

The control question A reactor that cannot be turned down sells into a market whose price moves sixteen to one across a day. How much does a supervisory economic optimizer earn over three price thresholds an engineer can defend, and on which days does its horizon pay for itself?

Learning objectives

- Write the heat balance of the steam header and compute how fast a residual moves the header pressure.
- Derive the electrolyzer switching price and the arbitrage value of a thermal megawatt-hour from the plant's parameters.
- Explain how the dispatch program eliminates the turbine, keeps its Hessian constant and carries the contract as a one-sided penalty.
- Compare fixed, threshold and optimizing supervisors on one day, with the heat left in the tank at midnight counted.
- Identify the constraint or forecast error that makes a horizon pay on a hard day.

Prerequisites Sections 1.4.2 (the header as stored inertia), 1.4.4 (the store balance), 1.4.5 (constant heat downstream), 1.4.6 (water value), 1.6.2 and 1.6.4. Chapter 10 for storage dispatch on a forecast.

The app https://apmonitor.com/apps/nuclear-hydrogen/

See also Volume 9, Chapter 6 for a price of heat set by predictive control and Volume 9, Chapter 7 for a release planned on an ensemble forecast. Volume 18, Chapter 9 for one header shared by three reactor modules and Volume 18, Chapter 5 for the steam dump after a load rejection.

13.1 Introduction

The usual justification for a supervisory optimizer is that it beats the heuristic underneath it. That claim is rarely presented with both numbers attached, and on the plant of this chapter it is only half true. Over an ordinary day, three price thresholds an engineer can defend at a meeting book \$516,300 against a fixed allocation's \$470,311, and a 48-variable economic optimization over a twenty-four hour horizon books \$529,303. The heuristic captures 78 % of the available improvement and the optimization collects the remaining \$13,003. Raise the evening price to five and a half times the day-ahead curve, leave the forecast blind to it, and the same two controllers separate by \$85,447. A horizon is insurance, and insurance looks expensive on the days it is not needed.

The plant that produces that result is a reactor whose output cannot move. Nuclear heat is capital rather than fuel, so no price justifies turning a reactor down, and the flexibility has to be built downstream of it [1, 2]. Four sinks take 860 MWth of reactor heat in this application. A steam turbine rated at 1180 MWth of swallowing capacity can absorb more than the reactor makes, which is what gives a store somewhere to discharge into. A two-tank molten salt store holds 2400 MWh at a 95.0 % thermal round trip [3]. A solid-oxide electrolyzer draws 180 MWe of stacks at 800 °C and takes header steam as well as electricity [4]. A dump condenser takes whatever is left and is worth nothing. Grid export swings from −102.7 MWe to 431.7 MWe off a heat source that never moves.

> **DID YOU KNOW? A Quarter of a Million Tonnes From One Reactor**
>
> A feasibility study for an existing United States plant found that a standard four-loop pressurized water reactor of about 1095 MWe could make at most about 240 thousand tonnes of hydrogen a year by high-temperature steam electrolysis. Piped to some six users, that output needed about 150 miles of pipeline [5].

What a simulator of that plant offers over a block diagram is the constraint itself. The header stores 3.0 MWh of steam and drum inventory per bar against 860 MW of supply, which is thirteen seconds of reactor output per bar, so the allocation equation is not a modeling convenience but a physical identity that closes every instant. An operator moving three sliders discovers that they are not three degrees of freedom, that turning the turbine down saves nothing because the heat still has to go somewhere, and that whatever the three do not account for leaves through a valve nobody is paid for.

Neither the identity nor its consequences appear in a transfer function from one input to one output, and the misconception the plant corrects is the one stated above, that a better optimizer is always worth having. A three-threshold rule is very hard to beat on a smooth day with no active constraints, and published assessments of control practice make the same point about advanced control in general, which is that the value comes from the constrained and the coupled cases rather than from the average one [6, 7].

The class of problem generalizes. A supervisory layer allocates one inflexible resource among competing consumers, against a forecast that is wrong, over a horizon long enough that terminal constraints bind, with a fast regulatory layer underneath. A refinery hydrogen network distributes make-up hydrogen from a reformer that does not turn down among hydrotreaters that compete for it [8]. A combined heat and power plant allocates one boiler between a steam host and an electricity market whose price moves [9]. A data center shifts a flexible fraction of its compute load in time against a carbon or price signal it must forecast first [10]. A reservoir operator spends a finite inflow across a season and prices the water left in storage against what the rest of the horizon will pay for it [11], which is exactly the marginal value that sets the charge and discharge thresholds here. In every one of them the supervisory problem is a linear or quadratic program over hours to months and the regulatory problem is PID over seconds [6, 12].

However, the layer that decides the allocation is normally met as a formulation on paper, with no means of watching a plan being made, watching it be wrong, and watching feedback recover part of it. This chapter treats a browser simulation of an integrated nuclear energy system in enough detail to be rebuilt from the text, and measures what a supervisory economic optimizer is and is not worth against the heuristic it replaces. Section 13.2 traces nuclear cogeneration and flexible operation. Section 13.3 collects the equations and constants that reproduce the simulator, and Section 13.4 derives the three prices that decide what a good allocation looks like. Section 13.5 describes the six regulatory loops and the three supervisors, among them an economic quadratic program whose Hessian is constant, whose turbine allocation is eliminated rather than optimized, and whose fixed eighteen projected-gradient iterations cost 23.4 µs per execution. Section 13.6 measures the three supervisory architectures on identical deterministic days, decomposes the improvement into what the horizon buys and what is left sitting in the tank at midnight, and puts a wrong forecast in front of a controller that acts on it with complete confidence. Section 13.7 names the instruments, final elements and control platform behind the loops, and Sections 13.8 and 13.9 state the limitations and what the results mean for practice.

13.2 Nuclear Cogeneration and Flexible Operation

Using reactor heat for something other than electricity is as old as civil nuclear power, and desalination was the first serious proposal [13], as the history segment after this paragraph tells. Chemical process heat followed the same reasoning at higher temperature, and the high-temperature gas-cooled reactor was developed in part to reach the 700 °C to 950 °C that industrial reforming, gasification and hydrogen production need. The economic argument in every one of those studies was the same as the one this application makes, which is that a capital-intensive heat source is cheapest when its output is constant and therefore wants a second customer whenever the first one does not want the heat.

⧖ HISTORY **Reactor Heat for Water and Industry (1973 to 2008)**

On the eastern shore of the Caspian Sea a 750 MWth sodium-cooled fast reactor started up in 1973 and, until 1999, supplied electricity, drinking water and heat to the city of Aktau. It sent up to 125 MWe to the grid and made 80 000 m^3 of potable water a day [14]. In Switzerland a 970 MWe pressurized water reactor has sent process steam to a nearby cardboard factory since 1979, about 2 % of its thermal power. At ten coastal sites in Japan, desalination units beside the reactors have made make-up and drinking water for more than 30 years, and in 2008 a heavy water reactor in India was retrofitted with a desalination plant [14].

Light-water plants were nevertheless built and licensed to run flat out, for reasons that are partly economic and partly physical. With fuel a small fraction of the busbar cost, every hour spent below full power is revenue given away against a cost that does not fall, so the merit order places nuclear at the bottom and leaves it there. The physical reasons are the ones that make load following an engineering exercise rather than a dispatch instruction. Reactivity feedback, control rod worth, xenon transients after a power reduction and fatigue on thick-walled components all restrict how far and how often a plant may be moved, and a profitability evaluation that carries those constraints explicitly finds the flexible operating envelope narrower and more expensive than a purely economic model suggests [1]. Fleets in France and Germany did operate flexibly at scale, and did so within exactly those limits.

The economic pressure changed when zero-marginal-cost generation reached a large share of the market. The value of a megawatt-hour is no longer a single number, because the average price a technology captures falls as its own penetration rises, and the effect is strong enough to be measured across European markets [15]. A firm generator therefore faces a price that varies by more than an order of magnitude inside one day and goes negative when solar output exceeds demand. Two answers appeared, and both put storage or a second product between the reactor and the market rather than asking the reactor to move. Thermal storage on the steam side lets a constant-output reactor sell a variable amount of electricity [2], and sizing that store for assured peak production rather than for arbitrage alone changes both the store and the turbine [16], which is why the turbine in this application swallows 1180 MWth against a reactor that makes 860 MWth.

The current programs make this a coordinated system problem. Nuclear-renewable hybrid energy systems are defined by the coupling between the reactor, the grid and an industrial process, and the research needs identified for them are led by the supervisory control question of how to dispatch the flexible elements rather than by the reactor itself [17]. Hydrogen is the industrial product that has attracted most of the effort, because high-temperature steam electrolysis consumes both electricity and heat and its electrical demand falls as its operating temperature rises [4]. Feasibility evaluations for hydrogen production at existing light-water plants set the scales, the electrolyzer specific energies and the hydrogen prices that this application uses [5]. What those studies leave to the control engineer is the layer above the loops, which is where a day-ahead price curve, a state estimate and a contract quantity have to become setpoints.

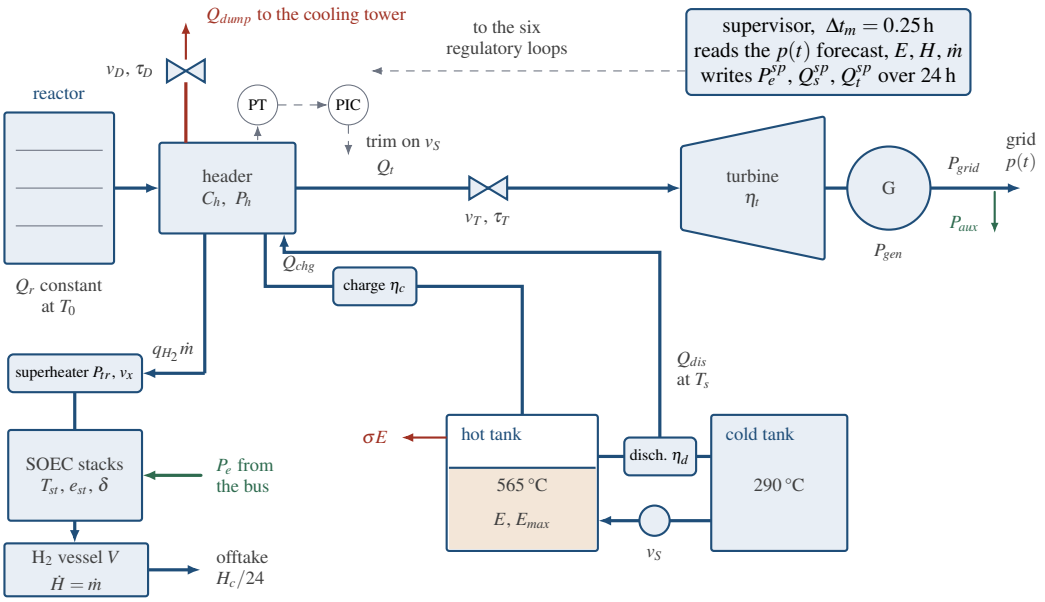


Figure 13.1 The integrated energy system, its four heat sinks and the two control layers. Constant reactor heat Q_r enters a steam header of inventory C_h holding pressure P_h, and leaves through the turbine throttle v_T, the charging exchanger, the electrolyzer steam draw $q_{H_2}\dot{m}$ and the dump valve v_D. Hot salt at E raises steam again through the discharge exchanger at T_s and returns it to the header, where the same throttle passes it to the turbine, which is why the store can lift generation above the reactor's own output. PIC trims the salt pump v_S to close the heat balance, and the supervisor writes the three allocation setpoints every 0.25 h from a price forecast and the measured state.

> ⧖ HISTORY **A Test Reactor Built for High Heat (1998 to 2010)**
>
> Light water reactors give steam at 200 to 300 °C, while many industrial processes want 500 to 900 °C [17]. Japan built a 30 MWth engineering test reactor of the helium-cooled, graphite-moderated type to reach those temperatures. It reached first criticality in 1998, ran at full power with an 850 °C outlet from 2001 and at 950 °C from 2004, and in a 50-day test in 2010 it held full power above 930 °C [14]. Its tests showed that it rode through a loss of coolant circulation without active control or emergency measures, and a thermochemical hydrogen process was to be connected to it. Today's designs of the type are passively safe modules of under 600 MWth each, added several to a site as capacity is needed [14], as the four modules of this chapter are.

13.3 Plant Model

The plant is four high-temperature reactor modules delivering a combined 860 MWth as 565 °C steam into one header. Steam leaves the header through four paths. The turbine throttle takes 840.8 MWth at the design point for 333.0 MWe gross at 39.60 % cycle efficiency. A charging heat exchanger moves heat into a hot salt tank at 565 °C over a cold tank at 290 °C, and a discharge exchanger raises 550 °C steam back out of it. The electrolyzer island draws 106.5 MWe and 19.2 MWth of header steam for 3.003 t/h of hydrogen, which is 72.07 t a day against a 72 t contract. The dump condenser takes the remainder. Station service is 12.0 MW, so 214.5 MWe reaches the grid. Figure 13.1 carries the layout and every symbol used in the equations, and Section 1.4.1 places the plant in the volume's flow scheme. Time is in hours, heat and power in megawatts, stored heat in megawatt-hours, hydrogen in tonnes, header pressure in bar and price in dollars per megawatt-hour.

Table 13.1 classifies what the two control layers interact with, and twelve of those entries are

Table 13.1 Classification of Variables. Nominal Values Are Those of the Design Allocation at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
Q_t (FIC)	CV, measured	840.8 MWth	250–1180 MWth
P_h (PIC)	CV, measured	120.0 bar	120 bar SP
Q_s (FIC)	CV, measured	0 MWth	−340–320 MWth
P_e (PIC)	CV, measured	106.5 MWe	18–180 MWe
T_{st} (TIC)	CV, measured	800 °C	800 °C SP
v_T throttle	MV	71.3 %	0–100 %
v_S salt pump	MV	0 %	−100–100 %
v_D dump valve	MV	0 %	0–100 %
P_e^{cmd}	MV	106.5 MWe	0–180 MWe
v_x superheater	MV	34.9 %	0–100 %
p price	DV, measured	table	−25–$528/MWh
$\hat{p}$ forecast	DV, measured	table	spike fraction f
H_c contract	DV, measured	72 t/day	30–115 t/day
a_T turbine avail.	DV, measured	1.00	0–1.00
r reactor trim	DV, measured	1.00	0.85–1.03
c_E tank derate	DV, measured	1.00	0.20–1.00
δ_0 stack wear	DV, unmeasured	0.00	0–1.00
E	state	600 MWh	0–2400 MWh
P_h	state	120.0 bar	20–220 bar
v_T, v_S, v_x, v_D	states	see above	clipped
P_e	state	106.5 MWe	0–180 MWe
T_{st}	state	800 °C	300–950 °C
V	state	14.0 t	0–40 t
H, H_d	states	0 t	≥ 0
δ	state	0.00	0–1.00

integrated states. Operation is bounded by five latching annunciations: a dump condenser passing more than 25 MWth, a hot tank at its capacity, a hot tank empty, a header below the 104 bar low alarm, and a declaration that the hydrogen contract can no longer be reached even with the electrolyzer at nameplate for every hour that remains. None of the five alters the model, which leaves the transient that raised it readable to the end of the run.

13.3.1 The Heat Balance and the Header

The reactor cannot store heat and the header holds very little, so the allocation is an identity rather than a design target, as shown in Equation (13.1).

$$C_h \frac{\mathrm{d}P_h}{\mathrm{d}t} = r\,Q_r + Q_{dis} - Q_t - Q_{chg} - q_{H_2}\dot{m} - Q_{dump} \tag{13.1}$$

where C_h is the header and drum steam inventory in MWh/bar, P_h is header pressure in bar, Q_r is reactor thermal power, r is a reactor trim multiplier, Q_{dis} and Q_{chg} are the salt discharge and charge duties, Q_t is heat to the turbine, $q_{H_2}\dot{m}$ is the electrolyzer steam draw and Q_{dump} is the dump condenser duty, all in MWth. Both sides carry MWth/bar times bar/h. With $C_h = 3.0\,\mathrm{MWh/bar}$ a residual of 100 MWth moves the header at 33 bar/h, so nothing on the right of (13.1) can be left unaccounted for. The state is clipped to 20–220 bar after each step. Section 1.4.2 sets this header beside the other stored inertias of the volume, and Volume 18, Chapter 9 shares one header of the same kind among three small reactor modules.

The three duties in (13.1) come from three actuators, and each carries the pressure or tank end effects that a linear dispatch model cannot see, as shown in Equations (13.2) and (13.3).

$$Q_t = \frac{v_T}{100}\,Q_{t,\max}\,\phi\,a_T, \quad \phi = \mathrm{clamp}\big(P_h/P_{h,0}, 0.25, 1.20\big) \tag{13.2}$$

$$\begin{aligned} Q_{chg} &= \tfrac{1}{100}\max(0, v_S)\, P_{c,\max}\, f_{hi} \\ Q_{dis} &= \tfrac{1}{100}\max(0, -v_S)\, P_{d,\max}\, f_{lo} \end{aligned} \tag{13.3}$$

where v_T is throttle position in percent, $Q_{t,\max}$ is the swallowing capacity, a_T is turbine availability, v_S is the signed salt pump command, and $f_{hi} = \text{clamp}((c_E E_{max} - E)/40, 0, 1)$ and $f_{lo} = \text{clamp}(E/40, 0, 1)$ throttle the exchangers over the last 40 MWh of tank travel so that a full tank cannot be charged and an empty one cannot be discharged. The throttle gain in (13.2) falls with header pressure, which is why a starved header settles rather than running away.

13.3.2 The Two-Tank Store

The store is one energy state with fixed exchanger efficiencies and a standing loss, as shown in Equation (13.4), which is the inventory balance of Equation (1.4) written in stored heat rather than in state of charge.

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \eta_c Q_{chg} - \frac{Q_{dis}}{\eta_d} - \sigma E, \qquad 0 \le E \le c_E E_{max} \tag{13.4}$$

where E is stored heat in MWh, η_c and η_d are the charging and discharge exchanger efficiencies, σ is the standing heat loss per hour and c_E derates the usable capacity. Charge and discharge are kept as separate non-negative flows so that the round-trip loss stays linear in the allocation, and no optimum ever asks for both at once because both cost money. The thermal round trip is $\eta_c \eta_d = 95.04\,\%$.

13.3.3 Turbine Cycle and the Electric Balance

The gross cycle carries two penalties, one for part load and one for salt steam that arrives at 550 °C rather than 565 °C with no reheat behind it, as shown in Equation (13.5).

$$\eta_t = \eta_0 \left[1 - c_p (1 - x)^2\right] \left[1 - c_d f_d\right], \qquad x = \frac{Q_t}{Q_{t,\max}} \tag{13.5}$$

where x is clipped to 0–1.05 and $f_d = Q_{dis}/(Q_r + Q_{dis})$ is the salt-raised fraction of the steam, taken against the untrimmed reactor rating. At the design allocation (13.5) returns 39.60 % against a rated 40.00 %. The electric balance is Equation (13.6).

$$\begin{aligned} P_{gen} &= \eta_t Q_t, \qquad P_{aux} = P_{a,0} + k_p \left(Q_{chg} + Q_{dis}\right) \\ P_{grid} &= P_{gen} - P_e - P_{aux} \end{aligned} \tag{13.6}$$

where $P_{a,0}$ is station service, k_p is salt pumping power per unit of salt duty and P_e is the whole electrolyzer island load. A negative P_{grid} means the plant is importing, which at a negative price is a revenue stream rather than a fault.

13.3.4 The Electrolyzer Island

The electrolyzer is represented by a specific-energy correlation rather than an electrochemical model, and the island load covers three consumers, as shown in Equations (13.7) to (13.9).

$$P_{bop} = b P_e, \qquad P_{tr} = \frac{v_x}{100} P_{tr,\max} \tag{13.7}$$

$$P_{st} = \max\left(0, P_e - P_{tr} - P_{bop}\right)$$

$$e_{st} = \max(24,\ e_0 + e_T\,(T_{sp} - T_{st}) + k_\delta \delta) \tag{13.8}$$

$$\dot{m} = P_{st}/e_{st} \tag{13.9}$$

where b is the balance-of-plant share, P_{tr} is the electric superheater duty, P_{st} is stack DC power, $\dot{m}$ is hydrogen production in t/h, e_{st} is stack DC energy per tonne, e_T is its temperature coefficient, T_{sp} is the 800 °C stack setpoint and δ is a scalar degradation state. At the design allocation (13.7) splits 106.5 MWe into 97.90 MWe of stack, 4.88 MWe of superheater and 3.73 MWe of balance of plant,

giving 3.003 t/h at 32.60 MWh/t of DC and 35.47 MWh/t of plant electricity. Holding the same load with the stacks at 750 °C raises that figure to 36.50 MWh/t, which is why the temperature loop is not decorative.

Stack temperature follows the enthalpy the feed steam and the superheater supply plus the stack's own irreversible heat, as shown in Equation (13.10).

$$\begin{aligned} T_f &= \frac{rQ_rT_0 + Q_{dis}T_s}{rQ_r + Q_{dis}}, \qquad \dot{m}_s = \max(1, s_{H_2}\dot{m}) \\ T^{tgt} &= \min\left(950,\ T_f + \frac{1636\,P_{tr}}{\dot{m}_s} + T_J\frac{P_{st}}{P_{e,\max}}\right) \\ \tau_{st}\frac{\mathrm{d}T_{st}}{\mathrm{d}t} &= T^{tgt} - T_{st} \end{aligned} \tag{13.10}$$

where T_f is the mixed feed steam temperature, $\dot{m}_s$ is steam flow in t/h at s_{H_2} tonnes of steam per tonne of hydrogen, T_J is stack self-heating at full load, and the constant 1636 °Ct/MWh is the inverse specific heat of steam at 2.2 kJ/(kg K). Turning the electrolyzer down reduces $\dot{m}_s$ faster than the superheater duty falls, so the stack runs hot unless the trim backs off.

13.3.5 Hydrogen Inventory, Degradation and Actuators

Production accumulates against a daily contract, and the offtake draws from a vessel that can run dry, as shown in Equation (13.11).

$$\frac{\mathrm{d}V}{\mathrm{d}t} = \dot{m} - \dot{m}_o, \quad \dot{m}_o = \frac{H_c}{24}\mathrm{clamp}\left(\frac{V}{0.5}, 0, 1\right), \quad \frac{\mathrm{d}H}{\mathrm{d}t} = \dot{m} \tag{13.11}$$

where V is vessel inventory in tonnes, H_c is the contract quantity in tonnes per day, $\dot{m}_o$ is the offtake rate and H is cumulative production. A third accumulator H_d integrates $\dot{m}_o$ and is reported but not used by any controller. Stack degradation accrues from three mechanisms, as shown in Equation (13.12).

$$\frac{\mathrm{d}\delta}{\mathrm{d}t} = g_0 + g_c\frac{\max(0, T_c - T_{st})}{100} + g_r\frac{\left|P_e^{cmd} - P_e\right|}{P_{e,\max}} \tag{13.12}$$

where T_c is the cold limit below which degradation accelerates and g_0, g_c and g_r are the baseline, cold and ramp coefficients. Thermal cycling and hard ramps dominate the total, which is the reason the dispatch layer carries an electrolyzer ramp limit it does not strictly need. The five actuators are first-order, as shown in Equation (13.13).

$$\tau_i\frac{\mathrm{d}u_i}{\mathrm{d}t} = u_i^{cmd} - u_i, \qquad u_i \in \{v_T,\ v_S,\ P_e,\ v_x,\ v_D\} \tag{13.13}$$

with time constants τ_T, τ_{pump}, τ_e, τ_{tr} and τ_D. Every state is clipped to its bound after each step, and the electrolyzer command additionally carries the 72 MW/h ramp limit on the controller output rather than on the actuator.

13.3.6 Price and the Money Accumulators

Price is an exogenous hourly table with the shape a market carrying a large solar share now has, cheap overnight, a morning ramp, a midday trough and an evening peak, as shown in Equation (13.14).

$$p(h) = [(1-g)\,p_b(h) + g\,p_{neg}] \times m(h) + p_{bias} \tag{13.14}$$

where p_b is the base curve, the solar glut g applies to hours 11 through 14 only and pulls them toward p_{neg} of −\$25/MWh, and $m(h) = s$ over the three hours starting at hour 17 and unity elsewhere. The base curve runs from \$6/MWh at hour 13 to \$96/MWh at hour 18, a ratio of sixteen to one before any

spike. The forecast the controller is given replaces s by $1+f(s-1)$, so $f=1$ is a perfect forecast and $f=0$ leaves the controller planning an ordinary evening. Money accrues as Equation (13.15).

$$\begin{aligned} \Pi_e &= \int p\,P_{grid}\,dt, \qquad \Pi_v = \int v\,P_{gen}\,dt \\ \Pi_H &= \int \left[c_c\,\dot{m}_c + c_m\,(\dot{m}-\dot{m}_c)\right] dt, \quad \dot{m}_c = \dot{m}\big|_{H\le H_c} \\ \Pi &= \Pi_e + \Pi_H - \Pi_v - c_s S\,\mathrm{clamp}(1-h_\ell/h_r, 0, 1) \end{aligned} \tag{13.15}$$

where v is variable operating cost on generation, c_c and c_m are the contract and merchant hydrogen prices, c_s is the take-or-pay penalty, $S=\max(0,H_c-H)$ is the shortfall, h_ℓ is the hours left before midnight and h_r is the window over which the accrual phases in. The accrual is what makes a missed contract visible as a step down in the margin trend during the last six hours rather than as a surprise at midnight.

13.3.7 Parameters, Initialization and Numerics

The fifty-six constants are listed in Table 13.2. Fifty-two are adjustable from the model dialog and four are hidden, and the four hidden ones are the standing tank loss and the three degradation coefficients, none of which a learner can read off a trend inside one day.

Every scenario opens at an allocation that closes (13.1) exactly. The three allocation controls set the throttle, the electrolyzer load and the salt duty, the electrolyzer island is evaluated through (13.7) and (13.8) at the stack setpoint, and whatever residual remains is placed on the dump valve so that the header does not have to swing on the first step, as shown in Equation (13.16).

$$\begin{aligned} \varepsilon_0 &= rQ_r + Q_{dis} - Q_t - Q_{chg} - q_{H_2}\dot{m} \\ v_D(0) &= \mathrm{clamp}\big(100\,\varepsilon_0/Q_{d,\max},\ 0,\ 100\big) \end{aligned} \tag{13.16}$$

The tank opens at $E(0)=\mathrm{soc}_0 c_E E_{max}$, the header at $P_{h,0}$, the stacks at T_{sp} and the vessel at V_0. A run that opens at a clock hour $h_0>0$ opens on a day already partly delivered, so cumulative production is initialized to $H(0)=(H_c/24)h_0 k_s$ where k_s is the fraction of schedule already met, and $k_s<1$ starts the day behind. Verification of the design point is one of fifteen self-checks the application runs, and it confirms 840.8 MWth, 214.5 MWe of export, 3.003 t/h and zero dump duty from (13.5) to (13.9).

The twelve states advance under fourth-order Runge-Kutta at a fixed 0.002 h step, which is 7.2 s of plant time. Six regulatory loops run on a 0.01 h (36 s) period, the supervisory quadratic program runs on a 0.25 h period, and the historian retains twenty logged samples per simulated hour. Not one disturbance in this application is stochastic, so the seed 20261118 never comes into play and each dollar figure quoted below is exactly repeatable. The speed selector sets simulated hours per second of wall clock and touches nothing else.

13.3.8 Assumptions

Detail is spent on whatever a controller has to be designed around and withheld elsewhere. Each of the seven simplifications below has a place where it stops holding.

The turbine is one efficiency curve. Equation (13.5) has no shaft, no inertia, no frequency, no governor droop and no minimum up or down time, so the turbine outage in scenario 10 is a heat-balance event rather than the mechanical and grid event it would really be. The 250 MWth minimum stable load is honored by the supervisor and by the operating map but is not enforced on the plant, so a manual allocation may carry the turbine below it.

The store is a single energy state. Equation (13.4) has no salt temperature, no tank level, no thermocline and no freeze protection, and the discharge steam temperature T_s is a constant rather than a function of the hot tank's condition. Two-tank molten salt systems in service depart from that on both counts [3]. The assumption understates how the last of a tank behaves.

Table 13.2 Model Parameters. Hidden Constants, Marked †, Are Never Displayed. Every One of the Fifty-Six Is Read by the Model, the Controller or Both.

Sym.	Meaning	Value	Unit	Sym.	Meaning	Value	Unit
Q_r	reactor thermal power	860	MWth	q_{H_2}	header steam per tonne H_2	6.4	MWh/t
h_0	clock hour at start	0	h	b	balance-of-plant share	0.035	—
T_0	header steam temperature	565	°C	$P_{tr,\max}$	superheater rating	14	MWe
$Q_{t,\max}$	turbine swallowing capacity	1180	MWth	s_{H_2}	steam fed per tonne H_2	12.0	t/t
$Q_{t,\min}$	turbine minimum stable load	250	MWth	T_J	stack self-heating at full load	25	°C
η_0	cycle efficiency at rated flow	0.400	—	T_{sp}	stack temperature setpoint	800	°C
c_p	part-load penalty	0.12	—	T_c	cold degradation limit	720	°C
c_d	salt-steam penalty	0.055	—	τ_e	rectifier response	0.008	h
τ_T	throttle travel time	0.020	h	τ_{tr}	superheater response	0.005	h
E_{max}	hot tank capacity	2400	MWh	τ_{st}	stack thermal time constant	0.15	h
$P_{c,\max}$	charging rate limit	320	MWth	$g_0{}^\dagger$	baseline degradation rate	8×10^{-6}	1/h
$P_{d,\max}$	discharge rate limit	340	MWth	$g_c{}^\dagger$	degradation, 100 °C cold	1.2×10^{-4}	1/h
η_c	charging exchanger efficiency	0.99	—	$g_r{}^\dagger$	degradation at a full-range ramp	9×10^{-5}	1/h
η_d	discharge exchanger efficiency	0.96	—	$P_{h,0}$	header pressure setpoint	120	bar
$\sigma^\dagger$	tank standing loss	1.458×10^{-4}	1/h	C_h	header steam inventory	3.0	MWh/bar
T_s	steam temperature from salt	550	°C	$P_{h,D}$	dump condenser opens at	126	bar
τ_{pump}	salt pump response	0.015	h	$P_{h,L}$	low header pressure alarm	104	bar
k_p	pumping power per salt duty	0.010	MWe/MWth	$Q_{d,\max}$	dump condenser capacity	1200	MWth
$P_{e,\max}$	electrolyzer nameplate	180	MWe	τ_D	dump valve travel time	0.010	h
$P_{e,\min}$	electrolyzer turndown	18	MWe	$P_{a,0}$	station service load	12	MWe
R_e	electrolyzer ramp limit	72	MW/h	$V_{\max}$	H_2 vessel capacity	40	t
e_0	stack DC energy at 800 °C	32.6	MWh/t	V_0	vessel inventory at midnight	14	t
e_T	stack energy per °C below T_{sp}	0.020	MWh/t/°C	c_c	contract hydrogen price	4200	$/t
k_δ	extra stack energy, full wear	3.0	MWh/t	c_m	merchant hydrogen price	1550	$/t
c_s	take-or-pay penalty	6000	$/t	h_r	shortfall phase-in window	6	h
v	variable O&M on generation	1.2	$/MWh	h_s	first hour of the evening peak	17	h
n_s	length of the evening peak	3	h	p_{neg}	midday floor at full solar glut	−25	$/MWh
τ_{Pm}	header pressure transmitter	0.004	h	τ_{Qm}	steam flow transmitter	0.006	h

High-temperature electrolysis is a correlation. Equation (13.8) has no polarization curve, no sweep gas, no pressurization and no stack-to-stack variation, and degradation is one scalar. The 35.47 MWh/t of plant electricity at 800 °C is a defensible figure for a nuclear-coupled solid-oxide unit [4, 5] and is not a measurement.

The dispatch model is deliberately a linearization of the plant beneath it, with a constant cycle efficiency, a constant electrolyzer specific energy, a linear steam draw per unit of load and no header dynamics at all. The mismatch is measured every execution and enters the linear cost vector as a bias, which is what keeps the layer offset free and the Hessian constant.

Price is an exogenous table and the plant is a price taker. There is no bid curve, no ancillary-service revenue, no imbalance settlement and no uncertainty band, and the forecast error is a fixed fraction of the spike chosen with a slider rather than a distribution.

The hydrogen contract is a daily quantity with a phased penalty. Real offtake agreements carry hourly delivery pressure obligations, monthly true-ups and force-majeure terms, none of which are here.

The numerics are more repeatable than a control system. A frozen step, a frozen scan and a frozen dispatch period give the identical trajectory everywhere, which is convenient and false, because scan order and execution time move on a real distributed control system and optimization server.

13.4 What an Hour Is Worth

Before any controller is chosen, three prices decide what a good allocation looks like, and all three are arithmetic on the parameters of Section 13.3.

13.4.1 The Electrolyzer Switching Price

Merchant hydrogen beyond the contract is worth $c_m = \$1{,}550/t$ against 35.46 MWh/t of plant electricity, which is \$43.71/MWh of electrolyzer load. That figure alone is not the switching price, because the electrolyzer also diverts header steam that the turbine would have converted, as shown in Equation (13.17).

$$p^* = \frac{c_m / e_{tot}}{1 + \kappa \eta_t}, \qquad \kappa = \frac{q_{H_2}}{e_{tot}} = 0.1805 \tag{13.17}$$

where $e_{tot} = 35.46\,\mathrm{MWh/t}$ is the linear specific energy the dispatch layer carries and κ is the thermal megawatts diverted per electrical megawatt of load. At the design cycle efficiency of 39.60 %, Equation (13.17) returns \$40.80/MWh, and at the rated 40.00 % it returns \$40.77/MWh. The price-threshold architecture is offered a default of \$44/MWh, which is deliberately the wrong number by \$3.20/MWh and is one of the three thresholds scenario 3 asks the learner to tune.

13.4.2 Storage Arbitrage in Thermal Terms

A megawatt-hour charged into the tank is a megawatt-hour the turbine did not convert, so the arbitrage has to be evaluated at the cycle efficiency on each side of the trade rather than at a single number. Charging at 320 MWth drives the turbine to 520.8 MWth where Equation (13.5) gives 38.50 %, and discharging at 340 MWth drives it to 1180 MWth where the same equation gives 39.38 % once the salt-steam penalty is applied. Charging one thermal megawatt-hour during the \$6/MWh hour therefore forgoes 0.385 MWh of generation and costs \$2.31, and returning $\eta_c \eta_d = 0.9504$ of it into the \$96/MWh hour yields 0.374 MWh worth \$35.93, a net \$33.62 per thermal megawatt-hour. Between \$24/MWh and \$35/MWh the same trade is worth \$3.86 and does not pay for the cycle. Salt pumping is left out of both figures.

The two efficiencies in that calculation are not equal, and the inequality runs the useful way. Evaluated at those endpoints the electrical round trip is $0.9504 \times 0.3938/0.3850 = 97.2\,\%$, which is higher than the 95.04 % thermal round trip, because the part-load penalty makes the displaced heat worth less per megawatt-hour than the heat the discharge adds. Measuring instead with the marginal efficiency over the full 320 MWth charge and 340 MWth discharge steps from the design allocation

gives 0.414 MWh forgone per thermal megawatt-hour charged and 0.387 MWh returned per thermal megawatt-hour discharged, an electrical round trip of 88.9 %. Charging the 0.010 MWe/MWth of salt pumping to both legs moves those to 0.424 MWh and 0.377 MWh, a round trip of 84.6 % and an arbitrage of $31.87 per thermal megawatt-hour. The dispatch layer of Section 13.5 uses one filtered efficiency for both directions, so within its own model the electrical round trip is the thermal one less pumping.

13.4.3 The Water Value of Stored Heat

Charge and discharge thresholds are not chosen in the optimizing architecture. They follow from a merit order over the forecast horizon, as shown in Equation (13.18).

$$\begin{aligned} E^{*} &= \min\left(c_E E_{max},\; E + \eta_c P_{c,\max} \Delta_b N_b\right) \\ n &= \operatorname{clamp}\left(\left\lceil \frac{\eta_d E^{*}}{P_{d,\max}\Delta_b} \right\rceil, 1, N_b\right), \qquad p_{ref} = \hat{p}_{[n]} \end{aligned} \tag{13.18}$$

where $\hat{p}_{[1]} \geq \hat{p}_{[2]} \geq \cdots$ are the forecast block prices sorted downward, Δ_b is the block length, N_b is the number of blocks, E^{*} is the fullest the tank could be made over the horizon, and n is how many blocks of full discharge that inventory would supply. The price at which the tank runs out is the marginal value of stored heat, and dividing the exchanger efficiencies into it produces the thresholds directly, charge below $p_{ref}\eta_c$ and discharge above p_{ref}/η_d. This is the same reservoir water value that season-long hydro scheduling is built on [11], evaluated over twenty-four hours instead of a year. It is the merit order of Equation (1.9) with two differences. The tank is first made as full as the cheap blocks of the horizon allow, and p_{ref} is the block price itself rather than that price multiplied by η_d, so both thresholds sit a factor $1/\eta_d$ above those of Chapter 1.

13.5 Two Control Layers

13.5.1 The Regulatory Layer

Six loops run on a 0.01 h scan and never see a price. Each is the velocity-form PID of Section 1.6.2 with back-calculation anti-windup, written with this chapter's symbols in Equation (13.19).

$$\begin{aligned} e_k &= a\left(\mathrm{SP}_k - \mathrm{PV}_k\right), \qquad \Delta u_k^{P} = K_c\left(e_k - e_{k-1}\right) \\ \Delta u_k^{I} &= \frac{K_c}{\tau_I} e_k \Delta t_c, \qquad s_k = s_{k-1} + \Delta u_k^{P} + \Delta u_k^{D} \\ c_k &= u_b + s_k + I_{k-1} \\ u_k &= \operatorname{slew}\left(\operatorname{clamp}\left(c_k, u_{\min}, u_{\max}\right),\; u_{k-1},\; R\,\Delta t_c\right) \\ I_k &= I_{k-1} + \Delta u_k^{I} + \frac{\Delta t_c}{T_t}\left(u_k - c_k\right) \end{aligned} \tag{13.19}$$

where $a = \pm 1$ carries the loop action, s_k is the unlimited running sum of Equation (1.11), c_k is the output the algorithm demands, u_k is that demand clipped to the output span and then rate limited at R, and T_t is the tracking time of the back-calculation in Equation (1.13). Only the electrolyzer loop has a finite R, its 72 MW/h ramp limit. No loop in this application uses derivative action, so $\tau_D = 0$ and $\Delta u_k^{D} = 0$ in all six, because every measured variable here is either an actuator-following signal or a pressure on a thirteen-second inventory and neither rewards anticipation. The six loops never leave automatic. Manual operation changes only where the three allocation setpoints come from. The one use of the tracking form is the stack trip, during which the rectifier loop's integral is reset every scan so that its output reads zero, and the load then returns from zero at the ramp limit when the trip clears.

Table 13.3 lists the six as they are configured. Two of them are worth separating from the other four. The header pressure controller is *direct* acting because charging takes steam out of the header, so more charge means less pressure, and its output straddles zero, which is why it can be left in

Table 13.3 The Six Regulatory Loops as Implemented. Action +1 Is Reverse Acting and −1 Is Direct Acting in the ISA Sense.

CV	MV	K_c	τ_I (h)	Act.	Output span
Q_t	throttle v_T	0.30	0.040	+1	0–100 %
P_h	salt trim	45	0.10	−1	−340–320 MWth
Q_s	salt pump v_S	0.30	0.030	+1	−100–100 %
P_h	dump valve v_D	12	0.060	−1	0–100 %
P_e	rectifier P_e^{cmd}	1.6	0.030	+1	0–180 MWe
T_{st}	superheater v_x	1.2	0.15	+1	0–100 %

service all day without parking against a limit. It does not drive an actuator directly. Its output is a signed trim added to whatever the supervisor asked the salt loop for, as shown in Equation (13.20).

$$Q_s^{eff} = \mathrm{clamp}\big(Q_s^{sp} + t_h,\ -P_{d,\max},\ P_{c,\max}\big) \tag{13.20}$$

where t_h is the pressure controller output in MWth. That arrangement is feedforward plus feedback with the store as the buffer. When the supervisor's plan is right the trim sits near zero, and when it is wrong the store silently carries the difference, which is what a thermal buffer exists to do. The dump condenser controller is the second. Its setpoint is 126 bar, six bar *above* the header setpoint, so it is an override rather than a regulator, it sits clamped at zero with a large negative error for most of every day, and the back-calculation term in (13.19) with $T_t = 0.02$ h is what stops it winding down while it waits. It holds the demand c_k near $(K_c/\tau_I)\,e\,T_t$, which is −24 % at a −6 bar error, rather than letting the integral run away. The tracking constants are 0.02 h on the throttle, dump and salt loops, 0.03 h on the header loop, 0.01 h on the electrolyzer and 0.05 h on the stack temperature loop.

Pause. Try Exercise 13.1 before reading on.

13.5.2 Three Supervisors

The regulatory layer is identical under all four architectures, and only the source of the three allocation setpoints changes. Manual operation passes the operator's three allocations. The fixed split passes one allocation chosen once and held for twenty-four hours, 840.8 MWth and 106.5 MWe by default, which meets the contract exactly and never dumps. The price-threshold architecture is three comparisons and a catch-up term, as shown in Equation (13.21).

$$\begin{aligned}
P_e^{thr} &= \begin{cases} P_{e,\max}, & p < p_{ely} \\ P_{e,\min}, & \text{otherwise} \end{cases} \\
P_e^{sp} &= \max\left(P_e^{thr},\ \mathrm{clamp}\left(\frac{H_c - H}{h_\ell} e_{tot}, P_{e,\min}, P_{e,\max}\right)\right) \\
Q_s^{sp} &= \begin{cases} +P_{c,\max}, & p < p_{chg},\ E < 0.995\,c_E E_{max} \\ -P_{d,\max}, & p > p_{dis},\ E > 0.005\,c_E E_{max} \\ 0, & \text{otherwise} \end{cases} \\
Q_t^{sp} &= \mathrm{clamp}(Q_r - Q_s^{sp} - \kappa P_e^{sp},\ Q_{t,\min}, Q_{t,\max})
\end{aligned} \tag{13.21}$$

with defaults $p_{chg} = \$22/MWh$, $p_{dis} = \$62/MWh$ and $p_{ely} = \$44/MWh$. The salt setpoint is then recomputed from the clipped turbine allocation so that the store absorbs whatever the turbine could not take. Equation (13.21) reads only the price on the meter at this instant and the contract position, and it has no view of any future hour at all. Its Q_r and κ are the declared values of Table 13.2.

13.5.3 The Dispatch Quadratic Program

The optimizing architecture is this application's own program rather than the framework option of Section 1.6.3, and its objective is money rather than a tracking error. It divides twenty-four hours into

$N_b = 12$ blocks of $\Delta_b = 2\,\mathrm{h}$ and carries four manipulated variables per block, the electrolyzer load, the charge duty, the discharge duty and the dump duty, for $N_u = 48$ decision variables. Turbine heat is not among them. It is eliminated through the heat balance, as shown in Equation (13.22).

$$\begin{aligned} \tilde{Q}_{t,k} &= rQ_r + Q_{dis,k} - Q_{chg,k} - \kappa P_{e,k} - b_Q - Q_{dump,k} \\ Q_{t,k} &= \mathrm{clamp}\big(\tilde{Q}_{t,k},\ Q_{t,\min},\ Q_{t,\max}\big) \end{aligned} \tag{13.22}$$

where Q_r and κ are the declared values of Table 13.2, r is the trim the reactor was dispatched to, and $b_Q = q_{H_2}\dot{m} - \kappa P_e$ is the measured error in the linear steam draw, with $q_{H_2}\dot{m}$ the metered steam to the stacks. Eliminating the turbine rather than optimizing it removes the only equality constraint in the problem, which leaves a genuine box-constrained program, and it makes every economic gradient exact instead of weighted. Charging costs the turbine heat it displaces at $p\eta_t$, discharging earns the same, dumping costs $p\eta_t$ and therefore *earns* that at a negative price, and the electrolyzer costs $p(1+\kappa\eta_t)$ while earning c_m/e_{tot}. No weight has to be chosen for any of the four.

Each decision is scaled to the unit interval, $u_{k,m} = o_m + s_m z_{k,m}$ with $z \in [0,1]^{48}$, where the offsets and spans are $(P_{e,\min}, P_{e,\max} - P_{e,\min})$, $(0, P_{c,\max})$, $(0, P_{d,\max})$ and $(0, 950\,\mathrm{MWth})$. The objective is $J(z) = \frac{1}{2} z^{\mathsf{T}} H z + g^{\mathsf{T}} z$ with the Hessian of Equation (13.23).

$$H = 2\left(w_H\, hh^{\mathsf{T}} + w_E\, ss^{\mathsf{T}} + \Lambda D^{\mathsf{T}} D + \rho I\right) \tag{13.23}$$

where h has the constant $\Delta_b s_{P_e}/(e_{tot} H_n)$ in each of the twelve electrolyzer positions and zero elsewhere, because cumulative hydrogen is a pure integrator of the electrolyzer decisions and H_n normalizes it, s carries $\Delta_b \eta_c s_{chg}/E_{max}$ in the charge positions and $-\Delta_b s_{dis}/(\eta_d E_{max})$ in the discharge positions, D is a first difference between consecutive blocks, Λ weights move suppression at λ on the electrolyzer and 0.4λ on the other three, and $\rho = 0.02$ is a ridge. Nothing in (13.23) depends on the price, the state or the contract, so H is assembled once at reset and rebuilt only when a weight or a plant limit moves. Default weights are $w_H = 8$, $w_E = 1.2$ and $\lambda = 0.06$.

Everything that changes every execution enters the linear term, as shown in Equation (13.24).

$$\begin{aligned} g_{k,P_e} &= 2w_H h_k c_H - 2\lambda\,[k{=}0]\, z_{P_e}^{now} \\ &\quad + \mu s_{P_e}\left[\hat{p}_k\,(1+\kappa\eta) - c_m/e_{tot} - v\eta\,\kappa\right] \\ g_{k,chg} &= 2w_E s_k c'_E + \mu s_{chg}\left[\hat{p}_k\,(\eta + k_p) - v\eta - v_E\eta_c\right] \\ g_{k,dis} &= 2w_E s_k c'_E + \mu s_{dis}\left[-\hat{p}_k\,(\eta - k_p) + v\eta + v_E/\eta_d\right] \\ g_{k,dmp} &= \mu s_{dmp}\left[\hat{p}_k\eta - v\eta\right] \end{aligned} \tag{13.24}$$

where $\mu = \Delta_b/M$ with $M = \$40{,}000$ scales the economic term against the dimensionless penalties, η is the filtered measured cycle efficiency, $v_E = w_V\,\eta\,p_{ref}$ is the water value of Equation (13.18) in dollars per thermal megawatt-hour, c_H is the offset of the hydrogen terminal state from its target and c'_E the offset of the terminal tank from its reference. The move-suppression terms carry the allocation actually in force, and the other three decisions carry $-0.8\lambda\,[k{=}0]\,z^{now}$ in the same way. The hydrogen target itself absorbs the urgency of the approaching deadline, as shown in Equation (13.25).

$$\begin{aligned} \gamma &= 1 + 2\,\mathrm{clamp}(1 - h_\ell/24,\ 0,\ 1) \\ H^{tgt} &= \frac{H_c \Delta_b N_b}{24} + \gamma\left(\frac{H_c h}{24} - H\right) \end{aligned} \tag{13.25}$$

Putting the urgency in the target rather than in w_H is what allows the Hessian to stay constant as midnight approaches. The measured model error in the hydrogen rate, $b_H = \dot{m} - P_e/e_{tot}$, enters the same target, which is the bias correction that makes the layer offset free.

Each execution runs a fixed $N_{it} = 18$ projected-gradient iterations, as shown in Equation (13.26).

$$z \leftarrow \Pi\big(z - \alpha\left[Hz + g - 2w_H h\,\max\big(0,\ h^{\mathsf{T}} z + c_H\big)\right]\big) \tag{13.26}$$

where Π is the feasibility restoration described below and the third term inside the bracket makes the contract penalty one sided. Falling short costs \$6,000/t while going over is merchant hydrogen already priced in the economic term, so when the plan is above target the terminal gradient is dropped by twelve subtractions while the Hessian keeps it, because the step size still has to be safe for the case where it binds.

The step size is not the textbook $1/\max_i H_{ii}$. Both terminal states in Equation (13.23) are sums over every block and enter the Hessian as dense rank-one blocks whose largest eigenvalue is N_b times their diagonal entry, so the row-sum bound of Equation (13.27) is used instead.

$$\alpha = 1 \Big/ \max_i \sum_j \left| H_{ij} \right| \tag{13.27}$$

Equation (13.27) is a Gershgorin bound on the spectrum, it is one extra loop, and it is computed once with the Hessian. At the default weights it returns $\alpha = 0.2066$ against $1/\max_i H_{ii} = 1.860$, so the diagonal step is 9.00 times longer than the safe one, and Section 13.6 measures what taking it costs.

Feasibility restoration Π runs in a fixed order after every gradient step. The four boxes are applied first. The electrolyzer ramp limit is then applied forward from the load in force, at $R_e \Delta_b$ per block. Simultaneous charge and discharge is netted, because it is always strictly worse than doing neither yet a gradient step can ask for it. The state of charge is restored in two passes, a merit-order trim that removes the excess from the blocks that valued it least, the dearest charging and the cheapest discharging, followed by a causal forward pass because energy cannot be borrowed from a later block. The turbine limits of (13.22) are restored last, by moving the dump valve and then the charge duty. That is greedy feasibility restoration rather than the exact Euclidean projection onto the chain, which is what a production dispatch code does, it is $O(N_b)$, and doing the merit trim before the causal pass is what stops the tank being spent on the first profitable block instead of the best one.

Warm starting slides the trajectory forward by the fraction of a block that has elapsed, $z_k \leftarrow (1-\beta) z_k + \beta z_{k+1}$ with $\beta = \Delta t_m / \Delta_b = 0.125$, so eight executions make one whole block shift and a block is refined ninety-six times before it becomes the present. Whenever another architecture is driving the plant, the first block of the trajectory is overwritten with the allocation actually in force before the solve, which is the same substitution dynamic matrix control makes when it charges its prediction with the move that was delivered rather than the move that was requested, and it is what makes selecting the optimizer in mid-run bumpless. One execution is a shift, a linear-term build and eighteen matrix-vector products of size 48×48. Timed over 400 executions after a 200-execution warm-up, which the application does at reset when its solver benchmark is switched on, the cost is 23.4 µs with a spread of 23.2 to 26.0 µs over six independent measurements on an Apple M-series laptop. The timing runs in the client, so the microsecond figure the application reports belongs to whatever machine it was opened on. A fixed iteration count with a precomputed step is the same design choice that makes first-order predictive control certifiable in real time [18], and the economic objective and terminal penalty follow the standard economic model predictive control formulation [19]. Every coefficient of the dispatch model, Q_r, κ, η_c, η_d, k_p and the rated η_0 that replaces the measured efficiency while the turbine is off, is the declared value of Table 13.2, frozen when the application loads. A change to one of them in the model dialog is therefore plant and model mismatch, absorbed by the biases b_Q and b_H, the measured tank level and the header pressure trim, while the plant limits, the contract, the prices, the forecast and the three measured disturbances a_T, r and c_E of Table 13.1 are read as the plant has them.

13.6 Results

Table 13.4 lists the eleven scenarios in the order a laboratory would work through them, the mechanism each exercises and what completes it. Each one pins the starting clock hour, the supervisory architecture, the three price thresholds, the four optimizer weights, the disturbance settings and any timed events. A case is credited when the listed actions are taken and the simulation reaches the listed

Table 13.4 Scenarios and Completion Requirements

#	Scenario (kind)	What it exercises	Required to complete
1	Allocate by hand at hour 11 (exploration)	the heat balance as an identity	In Manual, open the dump condenser past 25 MWth (a turbine near its 250 MWth minimum does it), raise the electrolyzer load to at least 150 MWe, and find an allocation that takes the header below its 104 bar low alarm.
2	Fixed split, one whole day (experiment)	a good plant and a poor merchant	Run the 24 h day on the fixed split at 106.5 MWe, then a second day with the fixed electrolyzer load set at least 5 MWe away before the run, and compare the hydrogen delivered with the 72 t contract.
3	Price thresholds, same day (experiment)	a heuristic with no horizon	Run the day on the shipped thresholds ($22, $62 and $44/MWh), then a second day with at least one threshold changed before the run and held for all 24 h.
4	Economic MPC, same day (experiment)	what a horizon is worth	Run the optimizer to 06:00 and read the plan ahead of the present, then run on to midnight and compare what was done with what was planned.
5	Afternoon spike, $f = 0.35$ (experiment)	three constraints binding at once	Run the optimizer from 08:00 to midnight with the 4.5× peak and an 82 t contract, then the same window on the price rules, the architecture changed before the run.
6	Negative midday prices (experiment)	every sign in the objective flips	On the optimizer, import more than 20 MW with the dump valve more than 5 % open for 0.5 h and run from 09:00 to midnight, then repeat on the price rules chosen before the run.
7	Spike at 5.5×, $f = 0$ (demonstration)	acting on a wrong forecast	Run the whole day on the optimizer with the 5.5× evening peak absent from the forecast, then the same day on the price rules or the fixed split chosen before the run.
8	Stack trip, contract behind (experiment)	a terminal penalty, not a constraint	On the optimizer, see the electrolyzer drop below 1 MWe after the 13:00 trip and pass 50 MWe again after the stacks return at 14:30, then run to midnight (12 h).
9	Derated tank, thresholds (demonstration)	a capacity limit felt hours later	On the price rules with an 840 MWh tank, run until the hot tank reports full and on to midnight (23 h), then repeat the day on the optimizer chosen before the run.
10	Turbine outage for 6 h (demonstration)	the dump condenser earns its name	On the optimizer, see the dump valve open past 10 % after the 07:00 trip, an hour of import above 100 MW with more than 300 MWth dumped, and after 13:00 the turbine above 500 MWth with the store discharging more than 50 MWth.
11	Graded challenge (challenge)	the whole day under a target	Over 24 h deliver at least 90 t, hold the header within ±6 bar of 120 bar whenever the dump valve is shut, and finish above $975,000, with the model parameters and disturbances as given.[a]

[a] As loaded (a fixed split at 780 MWth and 140 MWe, optimizer weights deliberately wrong) the day books $641,543 with 94.6 t delivered and the header within 1.07 bar while the dump valve is shut, so it fails on margin alone, by $333,457. The suggested configuration (economic MPC with $\lambda = 0.06$, $w_V = 1$, $w_H = 8$, $w_E = 1.2$ and a terminal state of charge of 0.25) books $1,003,720 with 94.3 t and 5.66 bar on the logged record, a pass by $28,720 and 0.34 bar, and the same whether it is revealed before Run or carried through a Reset.

state. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt. Section 1.6.4 describes the briefings, the credit and the certificates common to every application. Reset restarts a run from its starting clock hour and keeps the learner's own settings, while reloading a scenario restores its shipped ones, so a comparison of two architectures on one day needs only the architecture changed between runs. A run exports as a CSV record with 108 columns at twenty samples per simulated hour, holding the realized allocation beside the plan, the settled price beside the forecast, and the money accumulators beside the heat balance residual.

13.6.1 The Heat Balance Before Any Controller

Scenario 1 makes the point with arithmetic. At the design allocation the turbine takes 840.8 MWth, the electrolyzer steam 19.2 MWth and the store nothing, which sums to 860.0 MWth. Asking the store to charge at its 320 MWth limit forces the turbine to 520.8 MWth, and turning the turbine down

Table 13.5 Measured Outcomes at Midnight. The Ordinary Day Is Scenarios 2, 3 and 4 and the Spike Day Is Scenario 7 With the Architecture Changed. Every Figure Is Deterministic.

	Fixed	Thresholds	Economic MPC
Ordinary day			
Net margin ($)	470 311	516 300	529 303
Electricity ($)	177 392	178 678	203 243
Hydrogen ($)	302 508	346 641	335 312
H_2 delivered (t)	72.07	100.55	93.24
Heat dumped (MWh)	0.0	14.7	0.0
Tank charged (MWh)	0	1896	2976
Tank at midnight (%)	24.9	59.3	35.1
Spike day, peak 5.5×*, forecast blind*			
Net margin ($)	709 818	924 749	1 010 196

to its 250 MWth minimum without moving anything else places 590.8 MWth on the dump condenser. No combination of the three allocations avoids the identity, and the pressure controller closes it whether or not the operator intended the closure.

Over an economically optimized day the residual in (13.1) never exceeds 100.8 MWth, is larger than 40 MWth for 0.310 h out of twenty-four, holds the header within 2.02 bar of setpoint, and integrates to 25.3 ppm of the heat supplied over the day. A residual during a setpoint move is real and is what the buffer exists for, because 100 MWth held for one 36 s regulatory scan is 0.33 bar. What has to be true is that it is bounded, brief and integrates to nothing.

13.6.2 Three Architectures on One Day

Table 13.5 reports the same twenty-four hour price curve dispatched three ways, and then the same comparison with the evening peak at five and a half times the day-ahead curve and a forecast that saw none of it.

Three thresholds book $45,989 more than the fixed split and the optimization books a further $13,003, so the heuristic captures 78 % of the available improvement and the optimization is worth 2.5 % on top of it. That is the honest result on an ordinary day, and it is worth stating plainly, because an ordinary day's price curve is smooth and none of its constraints bind, and smooth unconstrained problems are what heuristics are good at.

🕮 **Pause.** Try Exercise 13.2 before reading on.

The comparison also has to account for what is left in the tank, because a run that ends with the store fuller than it started has put some of its value into inventory rather than into cash. The threshold architecture finishes the ordinary day at 59.3 % state of charge against the optimizer's 35.1 %, which is 579 MWh of stored heat. Discharged at η_d and the 39.4 % high-load efficiency into the overnight $26/MWh hour, that inventory is worth $5,693, which would cut the optimizer's advantage on this day from $13,003 to about $7,300 and raise the heuristic's share of the improvement from 78 % to 88 %. Carried into the following evening's $96/MWh peak instead, the same inventory is worth $21,000 and the ranking reverses. The ordering of the three architectures is robust and the size of the gap is not, which is why the application reports the state of charge at midnight beside the margin.

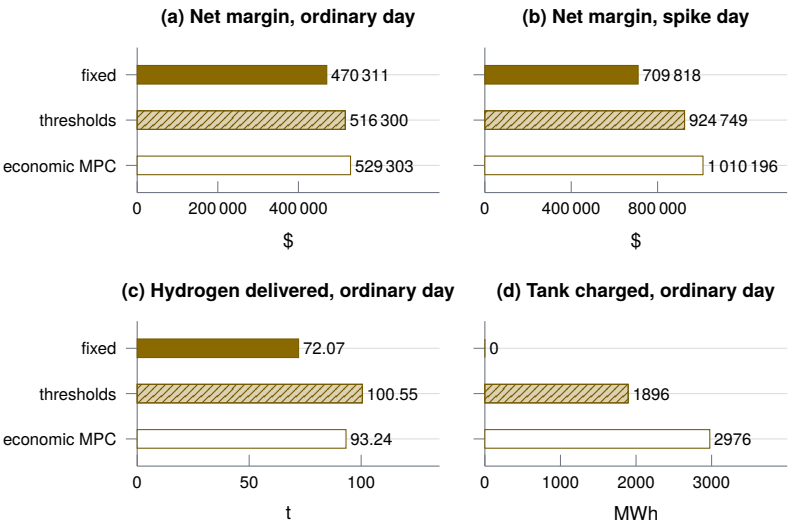


Figure 13.2 The three supervisory architectures of Table 13.5: (a) net margin on the ordinary day, (b) net margin on the spike day with a blind forecast, (c) hydrogen delivered and (d) heat charged into the tank on the ordinary day. Solid: fixed split. Hatched: price thresholds. White: economic MPC. Notice how close the thresholds come to the optimizer in (a) and how far apart they stand in (b), and that the thresholds make the most hydrogen while the optimizer moves the most heat through the store.

> ✦ **DID YOU KNOW?** **More Hydrogen, Less Money**
>
> From the simulator: on the ordinary day the threshold rule delivers 100.55 t of hydrogen and books $346,641 from it, more than the optimizer's 93.24 t and $335,312. The optimizer still books more in all, because its electricity earns $203,243 against $178,678 (Table 13.5).

On the spike day the arithmetic changes character. The optimization books $85,447 more than the thresholds, six and a half times its ordinary-day advantage, against a fixed split that books $300,378 less than the optimizer. Neither controller saw the spike in its forecast, and neither changes its dispatch when the meter shows it. Both allocations are those of the ordinary day to the megawatt, because the rules' discharge and electrolyzer thresholds are already crossed at the forecast peak and the optimizer prices every block, the present one included, from the forecast alone. The whole difference is therefore what each was already doing through the three peak hours, and the optimizer arrives at them with the tank it built for the forecast peak and spends it with the electrolyzer at turndown. The larger margin gap does not establish forecast-error recovery by MPC. Figure 13.2 draws Table 13.5 as bars.

The mechanism behind the gap is visible in the state of charge rather than in the allocation, as shown in Figure 13.3. The threshold rule charges whenever the price is below $22/MWh, which saturates the tank at hour 14 while the price is still $11/MWh, and it discharges only above $62/MWh, which begins at hour 17 and immediately meets the 340 MWth rate limit. Three hours of full discharge leave 58 % in the tank at hour 20 with the peak over, and the rule moves 1008 MWh out of the store over the whole day. The optimizer charges overnight to 71.6 %, discharges into the morning ramp, recharges through the midday trough to 93.4 %, begins the evening discharge at 14:31, two and a half hours before the rule, and empties the tank to 10.0 % at 21:33, moving 2591 MWh against the rule's 1008 MWh. It also holds the electrolyzer at its 18 MWe turndown from hour 17 to hour 20 while the threshold rule takes until 19:17 to get there, and it puts the turbine on its 1180 MWth swallowing

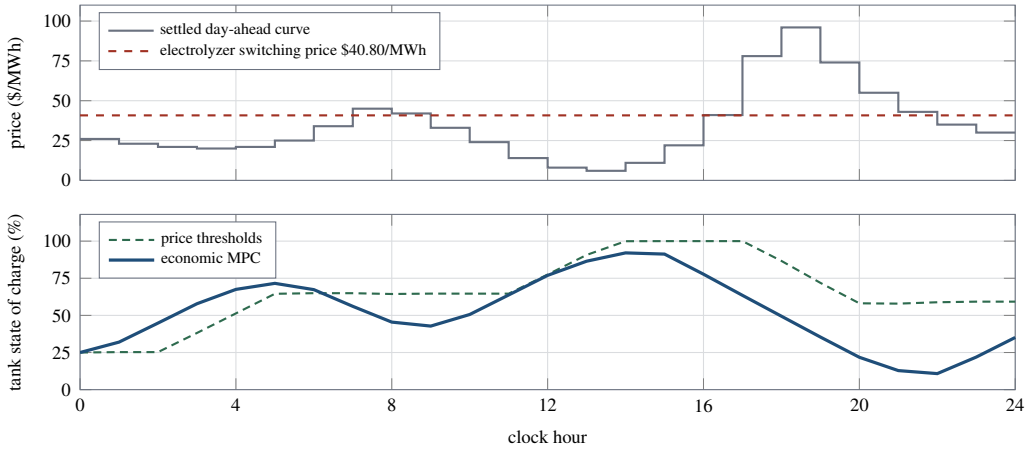


Figure 13.3 Where the twenty-four hour horizon shows up. The upper panel carries the settled price (solid, stepped) and the switching price (dashed), the lower panel the tank under the threshold rule (dashed) and the optimizer (solid). The threshold rule saturates the tank at hour 14 while the price is still $11/MWh, waits for a price above $62/MWh at hour 17 and then meets the 340 MWth discharge limit, so it moves 1008 MWh out of the store over the day and still holds 58 % at hour 20. The optimizer begins its evening discharge at 14:31, empties the tank to 10.0 % at 21:33 and moves 2591 MWh. Hourly samples of the physics-step state.

limit for 431.8 MWe of export at hour 19.5. Nothing told it to do any of that. There is no hydrogen setpoint anywhere either, and 93.24 t against a 72 t contract falls out of the terminal penalty of (13.25) and the merchant price of (13.17), with the 21.24 t of surplus made in the hours below $40.80/MWh.

> **TRY IT IN THE APP** **4 · Linear Economic MPC, the Same Day**
>
> Load scenario 4 and press Run. Watch the hot tank state of charge. The optimizer charges overnight to 71.6 %, spends some of it into the morning ramp, refills through the midday trough to 93.4 %, starts its evening discharge at 14:31 and empties the tank to 10.0 % at 21:33.

13.6.3 Two Formulation Details That Cost Money

Two lines of the dispatch layer are easy to write the obvious way and expensive to get wrong, and both are measurable on the ordinary day with every other setting held. Neither is a setting a learner can change, so the figures in this subsection come from counterfactual model runs on a modified copy of the controller.

Replacing the row-sum step of (13.27) with the textbook $1/\max_i H_{ii}$, a step 9.00 times longer, drops the day from $529,303 to $515,156, a loss of $14,147 that is larger than the entire advantage the optimizer had over the thresholds. The mean value of J after each of the 96 executions of the day rises from 2.36 to 3.06, so the iteration is not converging within its budget. What it does not do is oscillate between turndown and nameplate on alternate iterations. The projection absorbs most of the overshoot, and the largest back-and-forth move in a normalized electrolyzer decision (the smaller of two consecutive iteration moves of opposite sign) grows only from 0.0112 to 0.0274, so the symptom is a worse plan rather than a visible instability, which is the harder failure to notice.

Making the contract penalty symmetric, by keeping the terminal gradient in (13.26) when the plan is already above target, costs $20,463 on the same day and drops production from 93.24 t to 70.45 t. The symmetric penalty forbids profitable surplus, so the optimizer stops making hydrogen in the hours where electricity is below the switching price and then finishes 1.55 t short of the contract it was trying to hit exactly.

Table 13.6 Effect of the Iteration Count on the Ordinary Day. The Shipped Value Is 18.

N_{it}	2	4	10	18	120	600
Net margin (k$)	531.5	533.4	531.1	529.3	526.0	525.9
H_2 (t)	90.90	92.66	93.16	93.24	93.57	93.59

The third detail runs the other way and is worth reporting because it is counterintuitive. Table 13.6 sweeps the iteration count with everything else fixed.

Solving the quadratic program thirty-three times harder makes the day $3,396 *worse*, and the whole range from two iterations to six hundred spans $7,476, or 1.4 %. The reason is that the program being solved is a linearization whose bias is corrected every fifteen minutes, so converging it exactly buys accuracy against a model that is already approximate while the move-suppression term keeps pulling a fully converged plan closer to the previous one. A capped iteration count is therefore not a compromise on this plant, and the honest reading of Table 13.6 is that eighteen iterations are sufficient and that no iteration count recovers a bad forecast.

🕮 **Pause.** Try Exercise 13.3 before reading on.

13.6.4 The Hard Days

Scenario 6 inverts every sign in the objective. With the midday hours at −$25/MWh the optimizer holds the turbine at its 250 MWth minimum, the electrolyzer at its 180 MWe nameplate and the plant at −102.7 MW at the meter, so it is paid to import, and it sends 1529 MWh to the cooling tower over the run on purpose because heat that cannot be used is better dumped than converted at a negative price. It finishes $13,719 ahead of the threshold rule, which reaches only 502.5 MWth on the turbine, never imports more than 1.1 MW and dumps 15 MWh.

Scenario 5 binds three constraints inside one hour, the 72 MW/h electrolyzer ramp, the 340 MWth discharge limit and the turbine's swallowing capacity. The optimizer gets the stacks to their 18.0 MWe turndown where the thresholds reach only 42.6 MWe, and it holds 431.7 MWe of export against the rule's 407.4 MWe. It books $734,488 against $706,281, so the horizon is worth $28,207 here, and it is not free, because the optimizer finishes on 80.15 t against an 82 t contract and pays take-or-pay on the last 1.85 t. Two tonnes of hydrogen were worth less than the peak hours they would have cost, which is the trade that writing the contract as a penalty rather than a constraint was for.

Scenario 8 trips the electrolyzer for ninety minutes with the day already behind. Against a companion model run of the identical day with the scripted trip removed, which a learner cannot do, the trip costs 2.90 t of hydrogen and $46,437 of margin. The deficit enters the target of (13.25) rather than the weight, so the Hessian never changes and the solve stays exactly as cheap as it was, and the optimizer buys the deficit back in whichever remaining hours are cheapest instead of declaring the problem infeasible.

Scenario 9 derates the tank to 840 MWh and runs the threshold rule. Charging at 320 MWth through every hour under $22/MWh brings the exchanger derate factor f_{hi} of (13.3) into action, the dump valve opens at 03:10, the dump annunciation follows at 03:11, the tank reaches its capacity at 03:32, the header climbs to 126.54 bar and the dump condenser peaks at 293 MWth, sending 51.4 MWh of reactor heat to the cooling tower in the cheapest hours of the day. The same day on the optimizer never opens the dump valve at all and never takes the header past 120.80 bar, because the capacity constraint is enforced over the whole horizon rather than over the present instant. That is also the clearest statement in the application of what a horizon does not buy, because the optimizer ends the run holding 91.1 % of a small tank it never got to spend and books $482,537 against the rule's $501,272. A horizon bought feasibility here and not money.

> **TRY IT IN THE APP** **9 · The Hot Tank Hits Its Limit**
>
> Load scenario 9 and press Run on the price rules it opens with. Watch the dump valve and the header pressure. The dump valve opens at 03:10, the tank is full at 03:32 and the header climbs to 126.54 bar. Run the day again with the optimizer chosen before Run: the dump valve never opens and the header never passes 120.80 bar.

Scenario 10 opens the generator breaker at 07:00 with the reactor at full power. The header carries the whole imbalance while the salt loop and the dump valve travel, and because the store happened to be discharging 148.5 MWth at that instant the residual reaches 985.4 MWth rather than 860 MWth, which against 3.0 MWh/bar is 5.47 bar/min, and the header peaks at 129.22 bar at 07:03. The reactor alone would give 4.78 bar/min. The plant then settles into the only configuration available, the store charging at its limit, the electrolyzer at nameplate on imported power at −195.2 MW at the meter, and 507 MWth straight to the cooling tower, and 3242 MWh of reactor heat is thrown away over the six-hour outage. That figure is the argument for building the store larger than arbitrage alone would justify. Volume 18, Chapter 5 treats the same event on a pressurized water reactor with the reactor kinetics and the secondary side in the model, where a secondary shut in at full power reaches its first safety valve within seconds.

The graded challenge asks for the full 90 t contract, the header held inside ±6 bar whenever the dump condenser is shut, and a margin above $975,000 through a midday collapse to −$18.8/MWh and an evening spike five times the day-ahead curve that the forecast saw a third of. As loaded, on a fixed split with the optimizer weights deliberately wrong, the day books $641,543. The suggested configuration books $1,003,720 and lands 94.33 t, and the threshold architecture with the case's own thresholds books $935,748 on the same day. The suggested configuration reaches 126.35 bar at its worst, 6.35 bar above setpoint, at 10:32 with the dump condenser 21.1 % open, and stays inside 5.66 bar of setpoint whenever that valve is shut. The carve-out is what makes the target consistent with its own answer. The dump controller holds its setpoint at 126 bar, so deliberate dumping puts the header 6 bar high by construction, and a flat band would ask the student to undo the very action this scenario teaches. The application judges the band on its logged record of header pressure and dump valve position, one sample every three simulated minutes, on which the suggested configuration also reads 5.66 bar.

Figure 13.4 shows the Dispatch view during scenario 4. The price band runs across the top with the forecast the controller was given dashed over it, the realized allocation is drawn to the left of the vertical marker and the optimizer's plan for the rest of the day is ghosted to the right of it, and the state of charge appears realized and planned below. That ghosted trajectory is the controller's intention, it is visible before it happens, and it is the single most useful object on the screen for the purpose this application exists to serve.

13.7 Technology in Practice

The symbols of Figure 13.1 are, in a plant, a pressure transmitter on the steam header, flow measurements on the turbine steam and on the salt circuit, a thermometer in the stacks, a power measurement on the electrolyzer feed, and five final elements. Figure 13.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The measurement the whole allocation answers to is the header pressure, because Equation (13.1) closes through it. Both the pressure trim and the dump override read it. The chapter gives its transmitter the 0.004 h lag τ_{Pm} and the steam flow transmitters the 0.006 h lag τ_{Qm}. In a pressurized water plant the turbine control valve is set from the turbine impulse pressure, a surrogate for turbine power [5], where the chapter's throttle loop reads the heat to the turbine directly.



Figure 13.4 The Dispatch view of scenario 4 at $t = 6.00$ h, with the optimizer's plan for the rest of the day ghosted ahead of the vertical marker, cut to the plant view and the heat readout of the control panel. Callouts: (1) the price ribbon, (2) the present, (3) the planned allocation, (4) the tank state of charge, realized and planned, (5) the heat balance readout. The price ribbon at the top carries the settled hourly curve with the forecast the controller was given dashed over it, with the dearest hours shaded amber and the cheapest violet on screen. Beneath it the allocation shows what has happened as solid traces to the left of the marker and what the controller intends as dashed traces to the right: turbine heat is planned down through the midday trough and up to its 1180 MWth swallowing limit from 15:45 to 21:45, the electrolyzer is planned in the opposite sense, down to its 18 MWe turndown over the peak, and the dump trace is flat on zero because at a positive price dumping costs the electricity the heat would have made. The hot tank state of charge below that stands at 67 % and is planned to fall through the morning ramp, refill to 96 % by mid-afternoon and empty into the evening peak. At this instant the heat balance reads 1000 MWth to the turbine, 170 MWth out of the store, 31 MWth of electrolyzer steam and nothing dumped, against the constant 860 MWth the reactor supplies, and the application reads $131,365 booked with 29.4 t of the 72 t daily contract delivered, 212 MW of export at the $25/MWh price of the hour then ending and the header on 120.1 bar.

TECHNOLOGY **Sensor: The Header Pressure Transmitter**

A common electronic pressure transmitter senses with a taut metal diaphragm between two fixed plates, so that a pressure difference flexes the diaphragm and changes two capacitances [20]. Hot steam would damage the sensing element, so the impulse line gets a low point or a pigtail siphon where condensate collects as a liquid barrier. A transmitter that follows the NAMUR signal levels drives its current to 3.6 mA or below, or to 21 mA or above, when its sensing element fails, and the control system can then force the loop into manual [20]. On a trend that is a jump off scale, not a drift.

What Moves It

Five final elements act on the plant: the turbine throttle, the salt pumps, the dump valve, the rectifier of the electrolyzer and its electric superheater. The salt pumps carry the chapter's lesson, because the pressure controller trims their duty around whatever the supervisor asked for. A review of hybrid systems lists remote flow monitoring and control valves and pumps for molten salts among the hardware such plants will need [17]. The electrolyzer moves faster than any of them, because 75 to

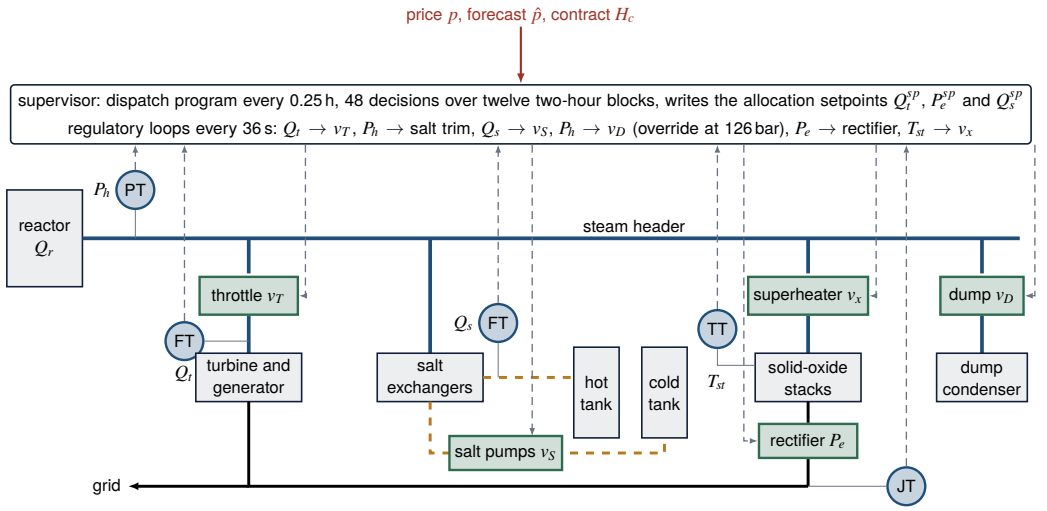


Figure 13.5 The chapter's loops with their hardware. Sensors are the round tags: the pressure transmitter PT on the steam header, flow transmitters FT on the turbine steam and on the salt duty, the thermometer TT in the stacks and the power measurement JT on the electrolyzer feed. Final elements are the square tags: the turbine throttle, the salt pumps, the dump valve, the rectifier and the electric superheater. Thick solid lines carry steam, dashed lines salt and plain black lines electricity. Notice that the supervisor reads only the price, the forecast, the contract and the plant state, and acts through three setpoints, while the header pressure reaches the salt pumps through the regulatory layer.

90 % of its power is electricity for electrolysis, which can be changed almost immediately [5].

> **TECHNOLOGY** **Final Element: Long-Shafted Salt Pumps**
>
> No mechanical seal suitable for molten salt has yet been identified, so salt pumps are vertical machines sealed by a throttle bushing, with the leakage running back to the suction by gravity [21]. In storage plants they hang from a structure above each tank and reach through its roof on shafts of 12 to 15 m, with bearings lubricated by the salt, and the tank keeps a minimum heel of salt over the impeller. Variable-frequency drives on the salt pumps allow the salt flow to be trimmed finely, and plants report few major problems apart from installation and alignment [21]. The chapter's one signed command v_S stands for the pumps of both tanks, charging one way and discharging the other, behind the 0.015 h lag τ_{pump}.

Where the Controller Runs

In a distributed control system a supervisory controller sends setpoints to the field controllers and requests data from them, and the field controllers move their actuators from those commands and their own sensors [22]. Supervisory and lower-level loops run with cycle times from milliseconds to minutes. The chapter's six loops run on a 36 s scan above a 7.2 s physics step, and the dispatch program writes its three allocation setpoints every 0.25 h. The dispatch program is the supervisory controller of that description, the six loops are its field controllers, and the pressure trim of Equation (13.20) is a summing block between the two layers.

IN PRACTICE **Keep the Stacks Hot**

Solid-oxide cells crack under steep thermal gradients. The rule of thumb that makers recommend for start-up and shutdown is no more than 2 °C per minute, so reaching 800 °C takes about 6.5 hours, too slow for a plant that moves between markets [5]. Makers therefore offer a hot standby mode. The electrolysis power is cut while the steam heat and the electric topping heat stay on, so no gradients form and production can resume at any moment [5]. In the chapter the dispatch never takes the stacks below their 18 MWe turndown, and the degradation term grows whenever they run colder than 720 °C.

What Is Modernizing It

Three developments touch this plant. The first is grid service from the electrolyzer. Nuclear plants have historically been unable to sell ancillary services because of their ramp limits, but a plant coupled to steam electrolysis could, because the electrical part of the electrolyzer's power can be moved back and forth almost at once [5]. The second is online optimization itself. More data generated, collected and used now makes continuous online optimization of a coupled plant possible [17], and the dispatch layer of this chapter is such an optimization. The third is the network. Operational systems increasingly resemble information systems and are less isolated from the outside world, so they need more security [22], and a supervisor that reads market prices sits at that boundary.

13.8 Limitations

The application is not a techno-economic study and cannot be used as one. With a surrogate steam cycle, a specific-energy electrolyzer and one deterministic price curve it cannot support a claim about whether a real integrated energy system pays for itself, what tank size is optimal, or what a hydrogen contract should be priced at, and the published feasibility work that can support such claims carries a full plant model and a market simulation [5, 17]. The reactor is a constant with a trim multiplier, so none of the neutronics, xenon or thermal-fatigue constraints that actually decide whether a nuclear unit may be moved appear anywhere [1]. The dispatch layer is a linearization with a measured bias, the feasibility restoration is greedy rather than an exact projection, and eighteen iterations return a good feasible plan rather than a solved program, so the trajectory drawn on the screen should be read as an intention and not as an optimum. Two-hour blocks smear one-hour prices, and a block averaging $39.50/MWh receives a partial electrolyzer allocation where hourly blocks would have given a full one in the cheap hour and none in the dear one. Hourly blocks over the same horizon would be ninety-six decision variables and roughly four times the solve cost, which is the trade the application makes on purpose. Section 13.3.8 states where each individual assumption stops holding.

OPEN CHALLENGE **Several Plants, One Supervisor**

A reactor, a grid connection and a hydrogen plant are separate systems with their own rules and time scales. A review of nuclear-renewable hybrid systems calls for techniques that coordinate linked but independent systems with differing dynamic time scales, and for new schemes of supervisory monitoring and control [17]. In 2019 high-temperature steam electrolysis had yet to be validated with a large generator at commercial scale, and how its stacks should be ramped thermally was still undecided [5]. Progress would be a coupled plant run for a season under a supervisor like the one of this chapter, with its constraints and its value measured.

13.9 Discussion

The result that matters for practice is the one in Table 13.5 read across both halves. An optimizer justified on average-day performance is justified on the wrong evidence, because the average day is where a heuristic is closest to it. The case for the horizon is built on the days where a constraint binds hours after the decision that caused it, which is scenario 9, where three constraints bind inside one hour, which is scenario 5, and where the price does something the rule was not written for, which is scenario 6. A plant that cannot articulate which of those it expects to meet, and how often, cannot

size the benefit of the layer it is buying, and this is the same conclusion that assessments of real-time optimization practice reach from industrial experience [6].

The second result is about forecasts. Scenario 7 gives the controller a day-ahead curve saying the evening peak is \$96/MWh when it will be \$528/MWh, and the controller acts on that curve with complete confidence. Re-planning every fifteen minutes rebuilds the plan from the state the plant is actually in, but the plan prices every block from the forecast, so the settled price never reaches it and the day is dispatched exactly as the ordinary day was. Feedback here corrects the state and not the price, and no re-plan can recover a storage decision taken six hours earlier. Chapter 10 measures the same pattern on a flow battery, with a forecast worth little on an ordinary day and a great deal on a day with one scarce hour. The engineering answer is not a better optimizer, and Table 13.6 shows that solving the same program harder makes the day slightly worse. The answer is a forecast with an uncertainty band and a policy that keeps some tank in reserve against the tail, which costs money on every ordinary day and pays for itself on a handful. That is a risk decision rather than a control decision, and it is the point at which a control engineer should say so.

Methods other than the implemented layer fit the same documented model. The twelve states admit a nonlinear economic formulation carrying the part-load curve and the temperature-dependent stack energy inside the optimization rather than as a measured bias [19], and a scenario-based formulation over the same horizon would price the uncertainty that scenario 7 exposes, as the scenario-tree program of Volume 9, Chapter 7 does for a reservoir on an ensemble forecast. Moving horizon estimation on (13.4) and (13.8) would recover the standing tank loss and the degradation state that the parameter table hides, turning two hidden constants into a soft sensor, and the block structure transfers directly to the cogeneration scheduling formulations the same problem class uses at larger scale [8, 9].

> ⚠ **OPEN CHALLENGE** **Hydrogen Beside a Licensed Reactor**
>
> Hydrogen has its own codes and standards, and a flammable gas near a reactor brings closer scrutiny under the United States nuclear power rules 10 CFR 50 and 52. The hydrogen plant may have to stand outside the exclusion area boundary, which adds security problems, heat lost on the way and concern over losing a heat sink [17]. One path keeps it a non-nuclear-grade chemical plant: the reactor stays safe whatever the hydrogen plant does, events there cannot affect the reactor, and the heat transfer loop between them carries no safety function [14]. The chapter's model has neither the distance nor the safety case.

13.10 Conclusion

This chapter has treated an integrated nuclear energy system as an interactive laboratory for supervisory economic dispatch. Twelve states are integrated, with a steam header of 3.0 MWh/bar that closes the reactor's heat balance continuously, a two-tank store at a 95.04 % thermal round trip, a solid-oxide electrolyzer taking 35.47 MWh/t of plant electricity and 6.4 MWh/t of header steam, and a dump condenser sized for the whole reactor. Six PI loops close the plant on a 36 s scan, and the header pressure controller is direct acting and trims the salt loop around whatever the supervisor asked for. Above them, a 48-variable box-constrained quadratic program over twelve two-hour blocks eliminates the turbine allocation through the heat balance, carries the hydrogen contract as a one-sided terminal penalty, and costs 23.4 µs per execution because its Hessian is constant and its eighteen projected-gradient iterations use a precomputed row-sum step. Measured on one deterministic day, a fixed split books \$470,311, three price thresholds book \$516,300 and the optimization books \$529,303, so the heuristic captures 78 % of the improvement and the optimizer earns \$13,003, while on a day whose evening peak is five and a half times the forecast the same two controllers separate by \$85,447. Replacing the row-sum step with the textbook diagonal step costs \$14,147 and making the contract penalty symmetric costs \$20,463, both larger than the advantage the optimizer had. Extensions of the simulator would add a price forecast with an uncertainty band, hourly blocks with the solve cost measured against them, a reserve policy scored on the tail rather than the mean, and

a reactor model carrying the constraints that decide whether a nuclear unit may be moved at all. Measuring what a supervisory optimizer is worth against the heuristic it would replace is what the simulator makes accessible.

KEY POINTS

- A reactor's heat cannot be turned down, so at every instant all of it goes to the turbine, the salt store, the electrolyzer or the dump condenser.
- On an ordinary day three price thresholds capture 78 % of the improvement over a fixed split, and the optimizer adds $13,003.
- The horizon pays where a constraint binds hours after the decision, where several bind at once, or where the price does something the rule was not written for.
- Re-planning every fifteen minutes corrects the state and not the price, so a spike the forecast missed is dispatched exactly as an ordinary day.

Where else this applies. A plan that prices a store against a forecast also sets the heat price of Volume 9, Chapter 6 and the reservoir release of Volume 9, Chapter 7.

References

[1] R. Ponciroli, Y. Wang, Z. Zhou, A. Botterud, J. Jenkins, R. B. Vilim, and F. Ganda, "Profitability evaluation of load-following nuclear units with physics-induced operational constraints," *Nuclear Technology*, vol. 200, no. 3, pp. 189–207, 2017.

[2] P. Denholm, J. C. King, C. F. Kutcher, and P. P. H. Wilson, "Decarbonizing the electric sector: Combining renewable and nuclear energy using thermal storage," *Energy Policy*, vol. 44, pp. 301–311, 2012.

[3] U. Herrmann and D. W. Kearney, "Survey of thermal energy storage for parabolic trough power plants," *Journal of Solar Energy Engineering*, vol. 124, no. 2, pp. 145–152, 2002.

[4] M. A. Laguna-Bercero, "Recent advances in high temperature electrolysis using solid oxide fuel cells: A review," *Journal of Power Sources*, vol. 203, pp. 4–16, 2012.

[5] K. Frick, P. Talbot, D. Wendt, R. Boardman, C. Rabiti, S. Bragg-Sitton, M. Ruth, D. Levie, B. Frew, A. Elgowainy, and T. Hawkins, "Evaluation of hydrogen production feasibility for a light water reactor in the midwest," Idaho National Laboratory, Idaho Falls, ID, Tech. Rep., 2019.

[6] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[7] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[8] J. J. Alves and G. P. Towler, "Analysis of refinery hydrogen distribution systems," *Industrial & Engineering Chemistry Research*, vol. 41, no. 23, pp. 5759–5769, 2002.

[9] S. Mitra, L. Sun, and I. E. Grossmann, "Optimal scheduling of industrial combined heat and power plants under time-sensitive electricity prices," *Energy*, vol. 54, pp. 194–211, 2013.

[10] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care, S. Talukdar, E. Mullen, K. Smith, M. Cottman, and W. Cirne, "Carbon-aware computing for datacenters," *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2023.

[11] J. W. Labadie, "Optimal operation of multireservoir systems: State-of-the-art review," *Journal of Water Resources Planning and Management*, vol. 130, no. 2, pp. 93–111, 2004.

[12] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[13] B. M. Misra and J. Kupitz, "The role of nuclear desalination in meeting the potable water needs in water scarce areas in the next decades," *Desalination*, vol. 166, pp. 1–9, 2004.

[14] International Atomic Energy Agency, *Opportunities for Cogeneration with Nuclear Energy*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2017, no. NP-T-4.1.

[15] L. Hirth, "The market value of variable renewables," *Energy Economics*, vol. 38, pp. 218–236, 2013.

[16] C. Forsberg, "Variable and assured peak electricity production from base-load light-water reactors with heat storage and auxiliary combustible fuels," *Nuclear Technology*, vol. 205, no. 3, pp. 377–396, 2018.

[17] 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.

[18] S. Richter, C. N. Jones, and M. Morari, "Computational complexity certification for real-time MPC with input constraints based on the fast gradient method," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1391–1403, 2012.

[19] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 3851–3861.

[20] 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.

[21] 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.

[22] 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.

13.11 Exercises

13.1 Concept check · 2 min

Table 13.3 gives the header pressure controller the action -1, direct acting. Why is that the right action for this loop?

(a) A rising pressure calls for more charging, and charging takes steam out of the header, so the output must rise with the measurement

(b) The salt pump fails open, so a larger signal closes it and the sign of the loop has to be turned around

(c) Direct action keeps the loop out of windup, because its output then never reaches either end of its span

(d) The dump controller is reverse acting, so the trim must act the opposite way to avoid fighting it

13.2 Concept check · 1 min

On the ordinary day the threshold rule captures 78 % of the improvement the optimizer makes over the fixed split. Which reason does the chapter give?

(a) The default thresholds were tuned on this day, so the rule starts close to the optimum

(b) The price curve is smooth and no constraint binds, which is the case a heuristic handles well

(c) The two-hour blocks smear the prices so badly that the optimizer loses most of its advantage

(d) Eighteen iterations stop the optimizer far short of the optimum of its own program

13.3 Concept check · 2 min

Raising the iteration count of the dispatch program from 18 to 600 makes the ordinary day \$3,396 worse. What explains it?

(a) The row-sum step becomes unstable when it is repeated that many times

(b) Six hundred iterations overrun the fifteen-minute period, so each plan arrives late

(c) The program is a linearization corrected every execution, and move suppression pulls a converged plan toward the last one

(d) The one-sided contract penalty turns symmetric once the iteration has converged

13.4 Theory and calculation · 5 min

Merchant hydrogen sells at $c_m = \$1{,}550/\mathrm{t}$. The dispatch layer carries $e_{tot} = 35.46\,\mathrm{MWh/t}$ of plant electricity and $q_{H_2} = 6.4\,\mathrm{MWh/t}$ of header steam per tonne.

(a) Compute the hydrogen value per megawatt-hour of electrolyzer load, c_m/e_{tot}, and the steam ratio $\kappa = q_{H_2}/e_{tot}$. *Checkpoint:* \$43.71/MWh and 0.1805.

(b) Use Equation (13.17) to compute the switching price at the design cycle efficiency of 39.60 % and at the rated 40.00 %. By how much is the default threshold of \$44/MWh too high? *Checkpoint:* \$40.80/MWh and \$40.77/MWh, so \$3.20/MWh too high.

(c) Explain in one sentence why the switching price is lower than c_m/e_{tot}. *Checkpoint:* Each megawatt of load also diverts κ megawatts of heat that the turbine would have converted.

13.5 Theory and calculation · 6 min

The header stores $C_h = 3.0\,\text{MWh/bar}$, Equation (13.1).

(a) A residual of 100 MWth persists. How fast does the header move, and how far does it move in one regulatory scan of 0.01 h? *Checkpoint:* 33.3 bar/h, 0.33 bar per scan.

(b) When the turbine trips at full power the residual is the reactor's 860 MWth, or 985.4 MWth with the store discharging at the instant of the trip. Convert both to bar per minute. *Checkpoint:* 4.78 bar/min and 5.47 bar/min.

(c) The dump controller acts from 126 bar and the dump valve has a travel time $\tau_D = 0.010\,\text{h}$. Estimate how far above 126 bar the header rises during one travel time at the larger rate, and compare with the measured peak of 129.22 bar. *Checkpoint:* About 3.3 bar, to about 129.3 bar.

13.6 Theory and calculation · 7 min

Use Equation (13.5) with $\eta_0 = 0.400$, $c_p = 0.12$, $c_d = 0.055$, $Q_{t,\max} = 1180\,\text{MWth}$ and $Q_r = 860\,\text{MWth}$.

(a) Charging at 320 MWth puts the turbine at 520.8 MWth, and discharging at 340 MWth puts it at 1180 MWth with $f_d = 340/(860+340)$. Compute the cycle efficiency at each point. *Checkpoint:* 38.50 % and 39.38 %.

(b) With these endpoint efficiencies and the thermal round trip $\eta_c\eta_d = 0.9504$, compute the net value of one thermal megawatt-hour charged in the \$6/MWh hour and returned in the \$96/MWh hour, ignoring salt pumping. *Checkpoint:* \$2.31 forgone, \$35.93 returned, \$33.62 net.

(c) Repeat for a charge at \$24/MWh and a return at \$35/MWh, and compute the electrical round trip on the same convention. *Checkpoint:* \$3.86 net, 97.2 %.

13.7 App challenge · 8 min

Load scenario 2, Fixed allocation, one whole day, change no setting, and run at speed 6 (six simulated hours per wall second) with the default seed past 24 h. Pause and export the CSV file. Do the same for scenario 3, Price-threshold rules, the same day, and scenario 4, Linear economic MPC, the same day. In each file use the row at `time_h` = 24.00.

(a) Read `metric_margin` and `h2_cum_t` in each file. *Checkpoint:* \$470,311, \$516,300 and \$529,303, within \$1, and 72.07, 100.55 and 93.24 t.

(b) Compute the share of the optimizer's improvement over the fixed split that the rules capture, as a percentage of that improvement. *Checkpoint:* 78.0 %.

(c) Read `soc_pct` in the rules and optimizer files. Convert the difference to megawatt-hours of a 2400 MWh tank, value it discharged at $\eta_d = 0.96$ and a cycle efficiency of 39.4 % into the \$26/MWh hour, and recompute the optimizer's advantage and the rules' share. *Checkpoint:* 59.25 and 35.13 %, 579 MWh, \$5,693, an advantage of about \$7,300 and a share of 88 %.

13.8 App challenge · 8 min

Load scenario 7, The forecast is wrong and the controller acts on it anyway, and run at speed 6 with the default seed past 24 h. Pause and export the CSV file. Press Reset, set Dispatch architecture to Price rules before pressing Run, and repeat. Press Reset again, set Fixed split, and repeat. In each file use the row at `time_h` = 24.00.

(a) Read `metric_margin` in the three files. *Checkpoint:* $1,010,196, $924,749 and $709,818, within $1.

(b) Compute the optimizer's advantage over the rules, and divide it by its ordinary-day advantage of $13,003. *Checkpoint:* $85,447, about 6.6 times.

(c) Read `h2_cum_t` and `metric_charged_MWh` in the optimizer and rules files and compare them with the ordinary-day files of scenarios 4 and 3. Explain what you find. *Checkpoint:* 93.24 t and 2976 MWh for the optimizer, 100.55 t and 1896 MWh for the rules, both identical to the ordinary day.

13.9 App challenge · 8 min

Load scenario 9, The hot tank hits its limit, change no setting, and run at speed 6 with the default seed past 23 h, which is midnight on this run. Pause and export the CSV file. Press Reset, set Dispatch architecture to Economic MPC before pressing Run, and repeat. In each file use the rows up to `time_h` = 23.00.

(a) In the rules file read `metric_dumped_MWh`, `metric_margin` and `soc_pct` at $t = 23.00$ h, and the largest `hdr_bar`. *Checkpoint:* 51.4 MWh, $501,272, 0.0 % and 126.42 bar, within 0.05 bar.

(b) Read the same four values in the optimizer file. *Checkpoint:* 0.0 MWh, $482,537, 91.09 % and 120.80 bar.

(c) The optimizer books less money. Say what its horizon bought instead, and how much stored heat it holds at midnight in a tank derated to 840 MWh. *Checkpoint:* No dumping and a header near setpoint, with about 765 MWh left in the tank.

13.10 Application area · 4 min

An engineer at a utility asks why the plant needs a salt store, an electrolyzer and a turbine larger than the reactor when the reactor itself could follow the price. Give two economic and two physical reasons from the chapter, and say what the 1180 MWth swallowing capacity has to do with the answer.

13.11 Application area · 3 min

During the six-hour turbine outage of scenario 10 the plant throws 3242 MWh of reactor heat away with 507 MWth settled on the dump condenser. Why is the dump condenser sized for the whole reactor, and what does the assumption that the turbine is one efficiency curve leave out of this event?

13.12 Review the agent · 5 min

An AI assistant reviews scenario 4 and writes: "The optimizer delivers 93.24 t against a 72 t contract. The 21.24 t of surplus is a tracking error that wastes electricity. Make the contract penalty symmetric, so that the plan lands on 72 t exactly and the margin rises." Find the flaw and say how you would check it with the app.