INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 6

Boiling Water Reactor Simulator for Coupled Loops With Sign-Reversing Dynamics

Pages 142–173

Edited by John D. Hedengren

How to cite this chapter. "Boiling Water Reactor Simulator for Coupled Loops With Sign-Reversing Dynamics," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 142–173.

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

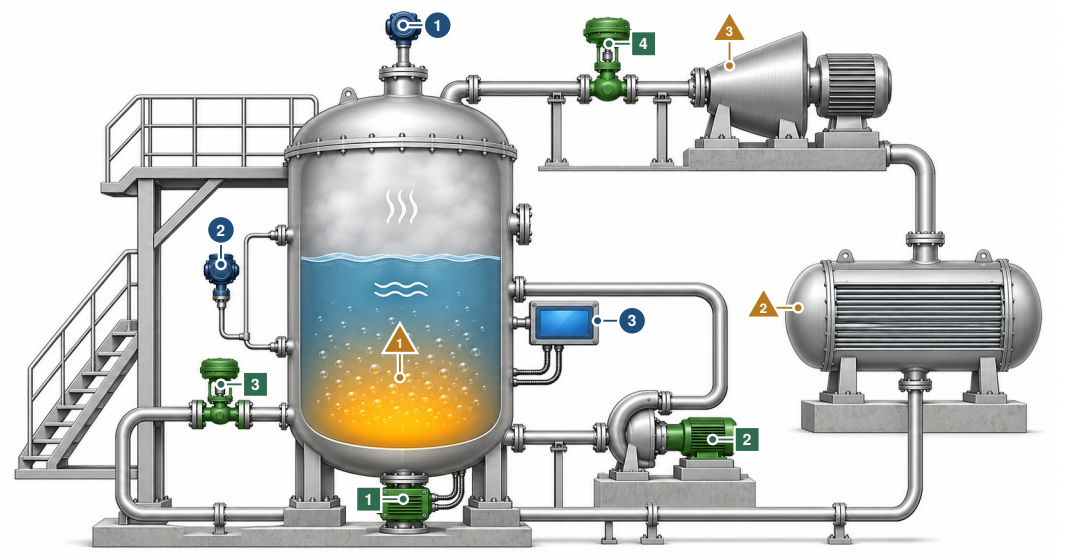
The app for this chapter
https://apmonitor.com/apps/bwr-reactor/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v18-nuclear-fusion/#bwr-reactor

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

CHAPTER 6

Boiling Water Reactor Simulator for Coupled Loops With Sign-Reversing Dynamics



The controller moves 1 control rod drives 2 recirculation pump speed 3 feedwater valve 4 turbine control valves

The controller reads 1 steam dome pressure 2 vessel water level 3 reactor power

What disturbs it 1 boiling noise in the core 2 feedwater heating 3 turbine steam demand

Notice that all four controls, from the rod drives under the vessel to the turbine valve above it, act on the same boiling water whose bubbles also lift the level the transmitter reads. Illustration, not to scale.

IN PLAIN TERMS **Boiling Water That Steers Its Own Reactor**

The main idea. In this reactor, water boils inside the core and the steam drives the turbine directly. The steam bubbles slow the chain reaction that makes the heat, so anything that changes the boiling changes the power.

The challenge. Adding cooler water collapses some bubbles, so the level gauge first reads low while the vessel is filling. A controller that trusts it feeds the wrong way, and the cure is to measure the steam leaving and replace it as it goes.

Why it matters. Four controls push on the same bubbles, so how they are wired together matters more than how hard each one pushes.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 6.2, the try-it and fact boxes of Section 6.4, the try-it and fact boxes of Section 6.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 6.6, Section 6.7, Section 6.8 and the open challenge

The control question The steam voids in a boiling water reactor core are at once the missing moderator, the strongest reactivity feedback and a false part of the level reading, so every actuator moves every loop. How should the four loops be paired, and how much of the result comes from structure rather than from tuning?

Learning objectives

- Compute the rod worth, the integrating gain of the pressure path and the decay ratio from the model's equations.
- Explain shrink and swell from the split of the vessel inventory into liquid and sub-surface steam.
- Pair the four loops from the measured self-regulation, pressure coefficient and flow-to-power gain.
- Compare single-element and three-element feedwater on one load change and say why structure wins.
- Judge a gain-scheduling rule against a measured process gain.

Prerequisites Sections 1.3, 1.4.2, 1.4.4, 1.6.2 and 1.6.4, and Chapter 4. Section 1.6.3 for the predictive controller only.

The app https://apmonitor.com/apps/bwr-reactor/

See also Volume 6, Chapter 2 (a pairing decided before any loop is tuned) and Volume 6, Chapter 7 (a structure choice that changes what the plant does). Volume 7, Chapter 2 for a self-regulating plant that ignites, and Volume 7, Chapter 13 for a delay-driven instability whose boundary moves with the operating point.

6.1 Introduction

A block diagram hides the one fact that decides how a boiling water reactor is controlled. Water boils in the fuel channels themselves, and the resulting steam voids are simultaneously the missing moderator, the strongest reactivity coefficient in the core, the mechanism by which recirculation flow sets power, and the volume that a differential-pressure level transmitter counts as though it were water. Four transfer functions drawn side by side present that as four plants. It is one state. Moving the turbine valves, the feedwater valve, the recirculation pumps or the control rods moves the void fraction, so each of the four loops is a load on the other three, and the coupling runs through the physics rather than through a matrix that could be inverted away.

The misconception this simulator corrects is specific and common. When the vessel level indication falls after feedwater is added, or rises after feedwater is lost, the reading is taken for an instrument fault. It is not. A differential-pressure cell straddling a two-phase mixture reports the height of the mixture, and in the first seconds of a feedwater change the mixture and the water inventory move in opposite directions. The application computes both from the same inventory and draws them together, so the wrong-way motion appears as a measured swollen surface against a computed collapsed surface rather than as a numerator coefficient chosen by the author of a textbook problem. A measured feedwater step of +10 % drives the indication −40.9 mm while the vessel is gaining 193 kg, and no gain and no filter changes that.

The class of plant is wider than reactors. Three features define it: a measurement that is the sum of two states responding to the same actuator with opposite signs, a feedback path that the operator never closed and cannot open, and loop gains that migrate with operating point. A natural-circulation steam drum has the first [1], distillation composition control with an inverse-responding tray has the first and the third, and any plant with a right-half-plane zero has the bandwidth limit that follows from it [2, 3]. What a boiling water reactor adds is that the second feature dominates the other two. The plant is its own controller, at a measured feedback gain of 30.6 pcm per percent of power, so interaction between the loops is not a nuisance to be decoupled but the mechanism that delivers most of the control action [4, 5].

Reduced-order boiling water reactor models exist in quantity, written to answer stability questions [5] or optimal control questions [6], and feedwater control algorithms for these plants have been documented against measurements [7]. However, a model published as a set of state equations gives the student the couplings already identified and the signs already assigned, which is the part of the

problem that decides the control structure. This chapter builds the void fraction into an interactive boiling water reactor as a state rather than a parameter, so that the four loop pairings, the two sign reversals and the migration of the gains can each be measured on the running plant, and it quantifies what each of five control architectures achieves on the same disturbances. The point kinetics, the drift-flux void correlation [8, 9] and three-element feedwater control are all standard. The last has been ordinary practice on drum boilers and on these reactors since long before either was simulated [7]. The simulator carries five loop architectures over one actuator set, one sample period and one set of limits, with the switch made by re-seeding the controllers rather than by a transfer routine, so a comparison between structures compares information flow and nothing else. The scenarios are ordered so that the plant, not the controller, is the thing that goes unstable. Every balance, property fit, constant and numerical choice needed to rebuild the simulator appears in Section 6.3.

Section 6.2 traces how boiling water reactor control arrived at its loop assignments. Section 6.3 states the plant model, its parameters and its assumptions. Section 6.4 measures the couplings and the two sign reversals on the running plant and places them on the operating map. Section 6.5 gives the loop pairing, the PID loops, the six architectures, the gain scheduling, the protection and the predictive controller. Section 6.6 compares single-element and three-element feedwater, follows a runback into the exclusion region, takes the graded challenge apart one change at a time and sets the predictive controller against the coordinated loops. Section 6.7 sets the loops in their plant hardware. Sections 6.8 and 6.9 state where the model stops and what the results mean beyond this plant, and Section 6.10 sums up. The exercises follow in Section 6.11.

6.2 Control of Boiling Water Reactors

Boiling in a reactor core was originally regarded as a hazard rather than a design feature, on the argument that a moving steam-water interface would make reactivity unpredictable. The experimental reactors of the 1950s settled the question the other way. Boiling proved to be self-limiting because the voids themselves remove moderator, and the control system that followed was built around a single-loop direct cycle in which the steam raised in the core drives the turbine with no secondary side. The earliest control account of such a plant describes the arrangement that survived: the turbine admission valves are assigned to reactor pressure rather than to power, because the void feedback already regulates power and the pressure is what must be held for that feedback to be predictable [10]. The assignment is still counterintuitive to anyone who has run a fired boiler, where the firing rate answers load and the throttle answers steam demand.

> ⧖ HISTORY **Boiling on Purpose, From Test Reactor to Power Plant (1954 to 1956)**
>
> The first test boiling water reactor showed that a reactor with water boiling in its core regulates itself and is inherently safe. It was destroyed on purpose in 1954, to show that extraordinary measures were needed to produce a damaging reaction. A pressurized successor was given a turbine generator, and on the night of July 17, 1955, its steam lit a nearby town. The Experimental Boiling Water Reactor followed, built in 19 months between May 1955 and December 1956 to prove the direct cycle, in which steam raised in the vessel goes straight to the turbine. It made 20 megawatts of heat and 5 of electricity and ran in parallel with the utility grid. Swings in demand went to an automatic steam valve that bypassed the turbine, so the reactor power did not have to follow them [11].

The void coefficient is the reason and it cuts both ways. It gives the plant strong inherent load-following, and it makes every quantity that changes the boiling into a reactivity input. In this model the coefficient is -1.2×10^{-3} in $\Delta k/k$ per percent of void, which at a delayed-neutron fraction of 0.0056 is 21.4 ¢ per percent of void. Because the void fraction is set by the flow through the channels, recirculation flow became the load-following actuator. Slowing the recirculation pumps raises the void fraction, subtracts reactivity and takes the reactor down with no rod motion at all, which covers roughly the top third of the power range in seconds rather than in the minutes a rod bank needs. Optimal-control studies of boiling water reactor load following are written around exactly that

actuator and the rod bank behind it [6]. The cost is that the recirculation system rather than the rod drive becomes the primary control element, and that the flow-to-power gain depends on where the plant sits on the operating map.

The same coefficient produced the industry's defining control event. A boiling channel is a feedback loop in its own right, because the pressure drop across it depends on the void fraction, the void fraction depends on the enthalpy the fluid absorbed one transit time earlier, and the neutronics close the loop through the void coefficient. Density-wave oscillations of this kind were catalogued for two-phase systems generally in the early 1970s and analyzed for the coupled neutronic case through the 1980s [5]. The event of March 1988 at LaSalle Unit 2 is told in the history box below [12]. The response was an exclusion region drawn on the operating map itself, entered by reducing flow rather than by raising power, and a requirement to leave it immediately by adding flow. Because the mode announces itself as ordinary neutron noise that stops settling rather than as an alarm, monitors that estimate the decay ratio from the neutron noise time series were developed for the purpose [13].

> ⧖ HISTORY **Two Pumps Trip, and the Core Oscillates (1988)**
>
> In March 1988 a boiling water reactor in the United States was running at about 85 percent power when a valving error in a routine test tripped both recirculation pumps. Power fell sharply as the core went over to natural circulation, to 40 percent within half a minute [14, 15]. Less steam meant less extraction steam for the feedwater heaters, some heaters were isolated, and the colder feedwater added reactivity [15]. Within five minutes the power began to oscillate, and about seven minutes after the pump trip the reactor scrammed on overpower at about 118 percent of full power [14]. A laboratory simulation traced the instability to low core flow, moderately high power, low feedwater temperature and a power shape peaked toward the bottom of the core [15].

Vessel level control took a separate path and arrived at the same place as the fossil boiler industry. Level alone onto the feedwater valve is the obvious loop and it is adequate at steady load. On a load change it is not, because the first several seconds of the measurement have the wrong sign, and a controller that reacts to them makes the excursion worse. The answer is three-element control: the steam flow measurement is fed forward to a feedwater flow controller so that every kilogram leaving the vessel is replaced as it leaves, and the level controller is reduced to a slow trim on the residual inventory error. Feedwater control algorithms for boiling water reactors have been documented in that form and compared against plant data [7], and drum boiler practice reached the same conclusion once the steam held below the water surface was recognized as the source of the reversal [1]. The steam generator of Chapter 4 carries the same three structures on a drum. The fix is structural rather than a matter of tuning because the right-half-plane zero the balances produce cannot be canceled, and it fixes a ceiling on closed-loop bandwidth that raising the gain only worsens [3]. Section 1.4.4 derives that ceiling for a drum.

6.3 Reactor Model

The unit is a BWR/4-class plant rated at 3330 MW thermal and 1100 MW electric. Saturated water leaves the steam separators at the top of the vessel, flows down the downcomer annulus where it mixes with the feedwater, is driven through the core by the jet pumps, boils in the fuel channels, and the steam is separated and dried and passes to the turbine through the control valves. Figure 6.1 shows the arrangement, the four manipulated variables and every symbol used in the equations that follow. The narrow-range level transmitter straddles the water surface, so it reports the swollen level while the water inventory sits at the collapsed level, and both are drawn.

Table 6.1 classifies what the four loops interact with. Seventeen states are integrated. The plant is expressed in seconds, kilograms, cubic meters, MPa, kJ/kg and °C, with level carried in meters and displayed in millimeters.

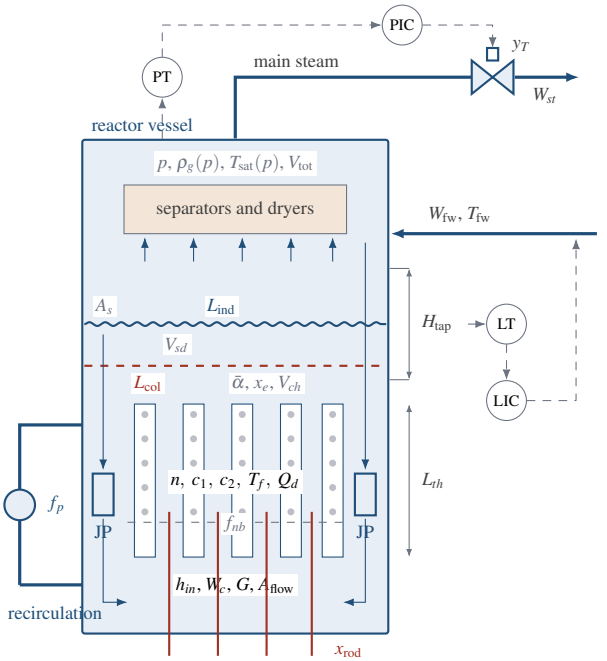


Figure 6.1 Reactor vessel, core channels and the four manipulated variables. Saturated water separated at the top of the vessel flows down the downcomer annulus, mixes with feedwater W_{fw} at temperature T_{fw}, enters the core at enthalpy h_{in} and is driven at flow W_c and mass flux G through channels of thermal height L_{th} by jet pumps at drive fraction f_p. Boiling begins above the non-boiling fraction f_{nb} and produces core-average void $\bar{\alpha}$ and exit quality x_e against the point-kinetics states n, c_1, c_2, the fuel temperature T_f, the decay-heat fraction Q_d and the rod bank x_{rod}. Steam leaves through the turbine control valves at position y_T, which PIC holds the dome pressure p with. The transmitter LT spans H_{tap} across the surface and reports the swollen level L_{ind}, while the water inventory sits at the collapsed level L_{col} and the difference is the sub-surface steam volume V_{sd}.

Table 6.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range	Symbol	Role	Nominal	Range
p dome pressure	CV, meas.	7.170 MPa	6.85–7.45 MPa SP	$\bar{\alpha}$ core void	state	40.32 %	0–97 %
L_{ind} level	CV, meas.	0 mm	±250 mm SP	h_{in} core inlet	state	1227.7 kJ/kg	—
Q_{th} power	CV, meas.	100 %	40–105 % SP	W_c core flow	state	13 100 kg/s	600–24 000 kg/s
y_T turbine valves	MV	85.0 %	0–100 %	f_p drive fraction	state	0.950	0–1.02
W_{fw} demand	MV	100.0 %	0–130 %	p dome pressure	state	7.170 MPa	4.6–9.4 MPa
f_p pump speed	MV	95.0 %	0–100 %	M_v vessel inventory	state	227 209 kg	—
x_{rod} bank	MV	35.0 % in	0–100 %	W_{fw} flow	state	1800.2 kg/s	—
T_{fw}	DV, meas.	216 °C	150–225 °C	T_{fw} state	state	216 °C	—
FW pump capacity	DV, meas.	130 %	0–140 %	y_T valve position	state	0.850	0–1
turbine admission	DV, meas.	1.00×	0.30–1.10×	x_{rod} state	state	35.0 %	0–100 %
ρ_{ext}	DV, unmeas.	0 pcm	±200 pcm	n neutron density	state	1.000	—
boiling noise	DV, unmeas.	0×	0–4×	c_1, c_2 precursors	state	124.4, 1716.5	—
T_f fuel temperature	state	699.5 °C	—	z_o void oscillation	state	0	±0.35
Q_d decay heat	state	0.0650	—	z_v its rate	state	0	±3

6.3.1 Saturated Properties

Saturation properties are quadratics in pressure fitted through IAPWS values at 5, 7 and 9 MPa, as shown in Equation (6.1).

$$\begin{aligned} \rho_g &= 0.13825\,p^2 + 3.9285\,p + 2.25225 \\ \rho_f &= -0.04875\,p^2 - 18.24\,p + 869.78875 \\ T_{\text{sat}} &= -0.54625\,p^2 + 17.5\,p + 190.09625 \\ h_f &= -2.1125\,p^2 + 81.9\,p + 797.9125 \\ h_{fg} &= 1.075\,p^2 - 80.15\,p + 2013.575 \end{aligned} \tag{6.1}$$

where p is in MPa, densities are in kg/m^3, temperature in °C and enthalpies in kJ/kg. The fits reproduce 6 and 8 MPa to better than 0.1 % on every property. The flashing coefficient $\mathrm{d}h_f/\mathrm{d}p = -4.225\,p + 81.9$ follows by differentiation and is used for the swell term. Feedwater enthalpy is $h_{\text{fw}} = 4.19\,T + 3.4 \times 10^{-4}\,T^2$ with T in °C, giving 920.9 kJ/kg at the design 216 °C. Property arguments are clamped to 4.5–9.5 MPa rather than extrapolated. At the design pressure of 7.170 MPa the fits give $\rho_g = 37.53\,\text{kg/m}^3$, $\rho_f = 736.50\,\text{kg/m}^3$ and $T_{\text{sat}} = 287.49\,°\text{C}$, so the steam is a twentieth as dense as the water it displaces and carries essentially none of the moderation.

6.3.2 Boiling Channel and Void Fraction

The channel closes an enthalpy balance on the heat the fuel is handing over, as shown in Equation (6.2).

$$h_e = h_{in} + \frac{Q_c}{W_c}, \qquad x_e = \frac{h_e - h_f}{h_{fg}} \tag{6.2}$$

where h_e is the exit enthalpy in kJ/kg, Q_c is the heat crossing into the coolant in kW, W_c is the core flow in kg/s and x_e is the exit quality. Voids appear before the bulk reaches saturation, so the non-boiling fraction of the channel is shortened from the thermodynamic value by the factor K_{nb}, as shown in Equation (6.3), which is the device that carries the subcooled-boiling behavior [9].

$$f_{nb} = \frac{K_{nb}\,\Delta h_{\text{sub}}}{\Delta h_{\text{sub}} + x_e h_{fg}}, \qquad \Delta h_{\text{sub}} = h_f - h_{in} \tag{6.3}$$

where Δh_{sub} is the core inlet subcooling in kJ/kg and f_{nb} is clamped to 0–0.95. The local void fraction at quality x follows the Zuber-Findlay drift-flux correlation [8, 16] and is averaged in closed form over a boiling length whose quality runs linearly from 0 to x_e, as shown in Equation (6.4).

$$\begin{aligned} \alpha(x) &= \frac{x}{C_0\left[x + \frac{\rho_g}{\rho_f}(1-x)\right] + \frac{\rho_g V_{gj}}{G}} \\ \bar{\alpha}_b &= \frac{1}{A} - \frac{B}{A^2 x_e}\ln\left(1 + \frac{A x_e}{B}\right) \\ A &= C_0\left(1 - \frac{\rho_g}{\rho_f}\right), \qquad B = C_0\frac{\rho_g}{\rho_f} + \frac{\rho_g V_{gj}}{G} \end{aligned} \tag{6.4}$$

where C_0 is the distribution parameter, V_{gj} the drift velocity in m/s and $G = W_c/A_{\text{flow}}$ the mass flux in kg/m^2/s. The core-average void fraction that the neutronics see is the boiling-length average diluted over the whole channel, as shown in Equation (6.5).

$$\alpha_{\text{ss}} = (1 - f_{nb})\,\bar{\alpha}_b, \qquad \tau_v \frac{\mathrm{d}\bar{\alpha}}{\mathrm{d}t} = \alpha_{\text{ss}} - \bar{\alpha} \tag{6.5}$$

where τ_v is the void relaxation time in seconds. At the design point (6.2) to (6.5) return $x_e = 13.742\,\%$, an exit void fraction of 65.23 % from (6.4), a non-boiling fraction of 10.57 % and a core-average void fraction of 40.32 %.

6.3.3 Point Kinetics and the Reactivity Balance

Reactivity is the sum of a rod term on an S-shaped integral worth curve, a Doppler term on the fuel temperature and a void term, all referenced to the design point so that the design state closes exactly, as shown in Equation (6.6).

$$\begin{aligned} \rho &= \rho_{\text{rod}} + \rho_D + \rho_\alpha + \rho_{\text{ext}} \\ \rho_{\text{rod}} &= W_{\text{rod}} \left[S(x_{\text{rod},0}) - S(x_{\text{rod}}) \right] \\ S(u) &= \frac{u}{100} - \frac{1}{2\pi} \sin\left(\frac{2\pi u}{100} \right) \\ \rho_D &= \alpha_D \left(T_f - T_{f,0} \right), \qquad \rho_\alpha = \alpha_v \left(\bar{\alpha}_{\text{eff}} - \bar{\alpha}_0 \right) \end{aligned} \tag{6.6}$$

where W_{rod} is the total bank worth in $\Delta k/k$, S is the normalized integral worth, α_D is the Doppler coefficient in $\Delta k/k$ per °C, α_v is the void coefficient in $\Delta k/k$ per unit void, $\bar{\alpha}_{\text{eff}} = \bar{\alpha} + z_o$ adds the density-wave oscillation and ρ_{ext} is an external reactivity disturbance. Equation (6.6) is the reactivity ledger of Equation (1.3) with the moderator term carried by the voids. Differentiating S gives a differential rod worth of $W_{\text{rod}} \left[1 - \cos(2\pi u/100) \right] /100$, which at the design insertion of 35 % is 47.6 pcm per percent of bank travel, or 8.51 ¢ per percent. Two delayed-neutron groups carry the kinetics, as shown in Equation (6.7). It keeps the prompt term that the prompt jump approximation of Equation (1.1) drops, and it carries two precursor groups in place of one.

$$\begin{aligned} \Lambda \frac{\mathrm{d}n}{\mathrm{d}t} &= (\rho - \beta_1 - \beta_2)\, n + \Lambda (\lambda_1 c_1 + \lambda_2 c_2) \\ \Lambda \frac{\mathrm{d}c_1}{\mathrm{d}t} &= \beta_1 n - \Lambda \lambda_1 c_1, \qquad \Lambda \frac{\mathrm{d}c_2}{\mathrm{d}t} = \beta_2 n - \Lambda \lambda_2 c_2 \end{aligned} \tag{6.7}$$

where n is the neutron density normalized to unity at rated power, c_i are the precursor concentrations, Λ is the prompt neutron generation time in seconds, β_i are the delayed fractions and λ_i their decay constants in s. One dollar is $\beta_1 + \beta_2 = 0.0056$. Heat passes through a single average fuel node and a one-group decay-heat state, as shown in Equation (6.8).

$$\begin{aligned} Q_{\text{th}} &= Q_0 \left[n \left(1 - f_{dk} \right) + Q_d \right], \qquad Q_c = \frac{T_f - T_{\text{sat}}}{R_f} \\ \frac{\tau_f}{R_f} \frac{\mathrm{d}T_f}{\mathrm{d}t} &= Q_{\text{th}} - Q_c, \qquad \tau_{dk} \frac{\mathrm{d}Q_d}{\mathrm{d}t} = f_{dk}\, n - Q_d \end{aligned} \tag{6.8}$$

where Q_0 is the rated thermal power in kW, f_{dk} is the decay-heat fraction, $R_f = \Delta T_{f,0}/Q_0$ is the fuel-to-coolant thermal resistance in °C/kW, τ_f is the fuel time constant and τ_{dk} the short-term decay-heat time constant. Thermal power is the reported controlled variable and neutron flux is reported separately, because the flux answers a reactivity insertion in milliseconds while the thermal power follows it through $\tau_f = 6.0\,\text{s}$.

6.3.4 Recirculation Flow

Core flow is driven by the jet-pump head plus the buoyancy of the voids themselves against a quadratic loop resistance, as shown in Equation (6.9).

$$\begin{aligned} \Delta p_{\text{nat}} &= \bar{\alpha} \left(\rho_f - \rho_g \right) g L_{th} \\ W_{c,\text{ss}} &= \sqrt{\frac{f_p^2\, \Delta p_{d,0} + \Delta p_{\text{nat}}}{K_{\text{loop}}}} \\ \tau_p \frac{\mathrm{d}f_p}{\mathrm{d}t} &= f_{p,\text{cmd}} - f_p, \qquad \tau_W \frac{\mathrm{d}W_c}{\mathrm{d}t} = W_{c,\text{ss}} - W_c \end{aligned} \tag{6.9}$$

where f_p is the pump drive fraction, $\Delta p_{d,0}$ is the jet-pump head at design drive, K_{loop} is the loop resistance in Pa/(kg/s)^2, L_{th} is the thermal driving height in m, τ_p is the drive time constant and τ_W the flow inertia time constant. The two constants are not free. The natural-circulation flow fraction f_{nat} that the plant is quoted at fixes the total loop resistance and the jet pumps supply the remainder, as shown in Equation (6.10).

$$K_{\text{loop}} = \frac{\Delta p_{\text{nat},0}}{f_{\text{nat}}^2 W_{c,0}^2}, \qquad \Delta p_{d,0} = \frac{\Delta p_{\text{nat},0}\left(f_{\text{nat}}^{-2} - 1\right)}{f_{p,0}^2} \tag{6.10}$$

where $\Delta p_{\text{nat},0}$ is (6.9) evaluated at the design void fraction. Feedwater mixes into the recirculating saturated water in the downcomer, and the transport of that mixture to the core inlet is the dead time in every shrink transient, as shown in Equation (6.11).

$$\tau_{DC}\frac{\mathrm{d}h_{in}}{\mathrm{d}t} = h_f - \frac{W_{\text{fw}}}{W_c}\left(h_f - h_{\text{fw}}\right) - h_{in} \tag{6.11}$$

where τ_{DC} is the downcomer mixing transport time in seconds. At the design point (6.11) is stationary at $h_{in} = 1227.7\,\text{kJ/kg}$, which is 48.87 kJ/kg of subcooling and 9.97 °C below saturation.

6.3.5 Pressure, Inventory and the Two Levels

The vessel carries one energy balance and one mass balance, as shown in Equation (6.12).

$$\begin{aligned} C_e\frac{\mathrm{d}p}{\mathrm{d}t} &= Q_{\text{th}} + W_{\text{fw}}h_{\text{fw}} - W_{\text{out}}\left(h_f + h_{fg}\right) \\ \frac{\mathrm{d}M_v}{\mathrm{d}t} &= W_{\text{fw}} - W_{\text{out}} \\ W_{\text{out}} &= K_t\, y_T\, p\, c_{\text{turb}}\, m_{\text{MSIV}} + K_{\text{SRV}}\max(0,\, p - p_{\text{SRV}}) \end{aligned} \tag{6.12}$$

where C_e is the lumped energy storage of water, steam and metal in kJ/MPa, M_v is the vessel inventory in kg, K_t is the choked turbine nozzle coefficient, c_{turb} is the admission capacity multiplier, m_{MSIV} is 1 with the main steam isolation valves open and 0 once they have latched shut, and K_{SRV} is the relief valve capacity in kg/s/MPa. The nozzle coefficient is calibrated so that the design valve position passes the design steam flow, $K_t = W_{s,0}/(y_{T,0}\,p_0) = 295.4\,\text{kg/s}$ per unit valve per MPa. Differentiating the first line of (6.12) at the design point gives an integrating process gain of –4.19 kPa/s per percent of valve travel, so pressure has no self-regulating term of its own in the valve path.

Pause. Try Exercise 6.3 before reading on.

The level follows from splitting the inventory by volume. Sub-surface steam has three parts, the void in the channels and upper plenum, the steam the separators are holding in transit, and the steam that a falling pressure is flashing out of the liquid, as shown in Equation (6.13).

$$\begin{aligned} V_{sd} &= \bar{\alpha}_{\text{eff}}V_{ch} + \frac{\tau_{\text{sep}}W_{sg}}{\rho_g} + \frac{\tau_{\text{sep}}W_{\text{flash}}}{\rho_g} \\ W_{\text{flash}} &= -\frac{V_l\rho_f}{h_{fg}}\frac{\mathrm{d}h_f}{\mathrm{d}p}\frac{\mathrm{d}p}{\mathrm{d}t}, \qquad W_{sg} = W_c\,x_e \end{aligned} \tag{6.13}$$

where V_{ch} is the two-phase volume above the boiling boundary in m^3, τ_{sep} is the separator and standpipe holdup time in seconds, W_{sg} is the boiling rate and W_{flash} the flashing rate, both in kg/s. The inventory then divides exactly between liquid, sub-surface steam and dome steam, as shown in Equation (6.14).

$$\begin{aligned} V_{wr} &= \frac{\left(M_v - \rho_g V_{sd} - \rho_g V_{\text{tot}}\right)/\rho_f + V_{sd}}{1 - \rho_g/\rho_f} \\ L_{\text{sw}} &= \frac{V_{wr} - V_{wr,0}}{A_s}, \qquad L_{\text{col}} = \frac{V_{wr} - V_{sd} - V_{l,0}}{A_s} \end{aligned} \tag{6.14}$$

where V_{wr} is the water-region volume in m^3, V_{tot} is the vessel free volume, A_s is the free surface area in m^2, L_{sw} is the swollen surface height and L_{col} the collapsed level, both in meters from normal water level. The two differ by V_{sd}/A_s, which at the design point is $80.71\ \mathrm{m}^3$ over $32\ \mathrm{m}^2$, or 2522 mm of apparent level that is not water.

The transmitter does not report L_{sw} either. Its measured pressure difference is converted to a height using the two densities present when the cell was ranged, not the two densities present in the vessel, as shown in Equation (6.15).

$$L_{\mathrm{ind}} = \frac{\left(\frac{H_{\mathrm{tap}}}{2} + L_{\mathrm{sw}}\right)\left(\rho_f - \rho_g\right) + H_{\mathrm{tap}}\left(\rho_g - \rho_{g,0}\right)}{\rho_{f,0} - \rho_{g,0}} - \frac{H_{\mathrm{tap}}}{2} \tag{6.15}$$

where H_{tap} is the separation between the taps in m and a zero subscript marks a density at the calibration pressure. Equation (6.15) returns L_{sw} exactly when $p = p_0$ and departs from it as the vessel moves off the ranging pressure, so the same physical surface reads differently at a different dome pressure.

6.3.6 Density-Wave Stability

The decay ratio of the density-wave mode is a fitted function of the power-to-flow ratio and, independently, of absolute flow, which is what places the exclusion region in the upper-left corner of the operating map [5], as shown in Equation (6.16).

$$\mathrm{DR} = \mathrm{DR}_0 \exp\left[a_{\mathrm{DR}}\left(\frac{P}{F} - 1\right) + b_{\mathrm{DR}}\left(1 - F\right)\right] \tag{6.16}$$

where P is thermal power and F core flow, both as fractions of rated, and P is taken from the fuel node rather than from the flux so that the decay ratio does not itself oscillate. The mode is a driven second-order oscillator on the void fraction whose damping comes from the decay ratio, as shown in Equation (6.17).

$$\begin{aligned} \frac{\mathrm{d}z_o}{\mathrm{d}t} &= z_v, \qquad \frac{\mathrm{d}z_v}{\mathrm{d}t} = -\omega^2 z_o - 2\zeta_{\mathrm{eff}}\,\omega\, z_v + d_o \\ \zeta &= \frac{-\ln \mathrm{DR}}{\sqrt{4\pi^2 + \ln^2 \mathrm{DR}}}, \qquad \zeta_{\mathrm{eff}} = \zeta + k_{\mathrm{NL}} z_o^2 \end{aligned} \tag{6.17}$$

where $\omega = 2\pi/T_{\mathrm{osc}}$, d_o is the boiling-noise drive and k_{NL} is the amplitude-dependent term that bounds the limit cycle. The oscillation enters the reactivity balance through $\bar{\alpha}_{\mathrm{eff}}$ in (6.6), so a decay ratio above unity appears on the neutron flux, which is where a stability monitor reads it [13].

6.3.7 Parameters, Initialization and Numerics

All sixty-two constants are listed in Table 6.2. Three are hidden and never displayed, and the remaining fifty-nine are adjustable from the model panel. The three control bands and the exclusion-region boundary appear in no derivative and are read only by the scorecard accumulators, which total the seconds each controlled variable spends outside its band.

Each scenario begins from an equilibrium the application solves for. Three residuals in three unknowns close the algebra, because the void fraction depends on the flow through (6.4), the flow depends on the void fraction through the buoyancy term in (6.9), and the power depends on both through (6.6). The residuals are total reactivity, the pressure or valve constraint, and the flow balance, and they are driven to zero by damped Newton iteration on a numerical Jacobian. The precursor states are then initialized to $c_i = \beta_i n/(\lambda_i \Lambda)$ and the inventory to the value that puts the transmitter of (6.15) on the operator's level setpoint. The equilibrium test in the application evaluates all seventeen derivatives at that state and requires the largest magnitude, scaled by the state, to be below 1×10^{-7}. The measured value at the design point is 1.4×10^{-14}, on the neutron density.

Table 6.2 Model Parameters. The Three Marked h Are Hidden and Never Displayed, and the Rest Are Adjustable.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
Λ	prompt generation time	4.5×10^{-5}	s	p_0	design dome pressure	7.17	MPa
β_1	fast delayed fraction	0.002 24	—	C_e	energy storage	1.40×10^7	kJ/MPa
λ_1	fast decay constant	0.40	s	V_{tot}	vessel free volume	480	m^3
β_2	slow delayed fraction	0.003 36	—	$V_{wr,0}$	water-region volume	380	m^3
λ_2	slow decay constant	0.0435	s	A_s	free surface area	32	m^2
α_v	void coefficient	−0.120	$\Delta k/k$	τ_{sep}	separator holdup	0.8	s
α_D	Doppler coefficient	-2.3×10^{-5}	1/°C	H_{tap}	tap separation	1.6	m
W_{rod}	total bank worth	0.030	$\Delta k/k$	$y_{T,0}$	design valve position	0.85	—
$x_{\text{rod},0}$	design insertion	35	%	τ_{TCV}	valve stroke time	0.25	s
—	rod drive rate	0.55	%/s	k^h_{gen}	generator output	0.6111	MW/(kg/s)
—^{h}	scram insertion rate	35	%/s	$T_{\text{fw},D}$	design feedwater temp.	216	°C
Q_0	rated thermal power	3330	MW	τ_{fw}	feedwater flow constant	3.5	s
$W_{c,0}$	rated core flow	13100	kg/s	$\tau_{\text{fw}T}$	feedwater temp. constant	30	s
A_{flow}	core flow area	8.7	m^2	DR_0	decay ratio at rated	0.25	—
V_{ch}	two-phase volume	105	m^3	a_{DR}	power/flow sensitivity	1.20	—
L_{th}	thermal driving height	5.0	m	b_{DR}	low-flow sensitivity	0.77	—
C_0	distribution parameter	1.13	—	T_{osc}	oscillation period	2.2	s
V_{gj}	drift velocity	0.23	m/s	k^h_{NL}	limit-cycle saturation	400	—
K_{nb}	subcooled-boiling factor	0.55	—	L_{hi}	Level 8 trip	+0.560	m
$\Delta T_{f,0}$	fuel-to-coolant rise	412	°C	L_{lo}	Level 3 scram	−0.450	m
τ_f	fuel time constant	6.0	s	p_{MSIV}	low-pressure isolation	5.60	MPa
τ_v	void relaxation	1.6	s	p_{SRV}	relief valve setpoint	7.86	MPa
τ_{DC}	downcomer transport	3.0	s	K_{SRV}	relief valve capacity	3000	kg/s/MPa
f_{dk}	decay-heat fraction	0.065	—	p_{scram}	high-pressure scram	7.72	MPa
τ_{dk}	decay-heat constant	80	s	n_{scram}	APRM high flux scram	1.20	—
$f_{p,0}$	design pump drive	0.95	—	τ_{Pm}	pressure transmitter	0.15	s
f_{nat}	natural-circulation flow	0.33	—	τ_{Lm}	level transmitter	1.2	s
τ_p	pump drive constant	5.0	s	τ_{Nm}	APRM filter	0.6	s
τ_W	flow inertia constant	2.5	s	τ_{Wm}	flow transmitter	0.8	s
DR_{lim}	exclusion-region boundary	0.80	—	—	level control band	0.100	m
—	pressure control band	0.10	MPa	—	power control band	2.0	%

Integration is classical fourth-order Runge-Kutta on a fixed 0.005 s step. Controller execution is every 0.05 s, a 20 Hz scan chosen to match a real electrohydraulic pressure regulator, and ten samples of every simulated second reach the historian. The noise sequences are drawn from a mulberry32 generator seeded at 20 260 906, which makes a noisy run repeat bit for bit. The engine, the seed and the speed selector, which changes no trajectory, are those of Section 1.6.1.

6.3.8 Assumptions

Detail is spent where a controller has to be designed around the physics and withheld everywhere else. Each simplification fails somewhere.

One point-kinetics node with two delayed groups. A real core has a strongly bottom-peaked axial power shape and can oscillate with its two halves out of phase, which is the regional mode that made the LaSalle event difficult to diagnose [12]. A point model cannot represent a regional mode at all, so only the core-wide in-phase oscillation appears here. The two groups reproduce the prompt jump and the minute-scale precursor tail but not the intermediate decay constants.

One average channel with uniform axial power. Several hundred bundles at different powers are represented by one channel, so every thermal limit that an operator actually manages, the critical power ratio and the linear heat generation rate, is written against a hot channel the model does not have. The application therefore says nothing about margin to dryout.

Constant void and Doppler coefficients. Both vary with exposure, void history and rod pattern, and the void coefficient becomes markedly more negative late in a cycle. There is no xenon, no burnup and no rod pattern, only one ganged bank on one S-curve, with the external reactivity slider standing in.

Drift flux integrated in closed form. Equation (6.4) is sound for rod bundles at 4–9 MPa above roughly 300 kg/m^2/s [8, 16]. Below that the drift term dominates and the correlation degrades, which is precisely the low-flow corner where the instability region sits, so the void fractions reported there are indicative rather than predictive.

Density-wave instability is a calibrated oscillator. Equations (6.16) and (6.17) set the damping of a second-order mode from a fitted decay ratio and drive it with boiling noise. The onset boundary and the 0.45 Hz period are calibrated to published boiling water reactor behavior [5], but the channel momentum equations that produce the phenomenon are not solved, so the limit-cycle amplitude is plausible rather than predicted.

One exponential group of decay heat. Equation (6.8) carries 6.5 % of rated power with an 80 s time constant, which is fair for the first two or three minutes after a scram and increasingly wrong after that.

The turbine is a choked nozzle and the generator is a constant. There is no shaft, no inertia, no frequency, no overspeed protection and no turbine bypass valves, so a turbine trip here is a pressurization rather than the mechanical event it would be, and the bypass system a real plant uses to absorb one is absent. Chapter 5 treats the steam dump that takes that role on a pressurized water reactor.

One control volume for pressure and one for mass. Carry-under, downcomer subcooling gradients and the separator hydraulics are represented only by the holdup time τ_{sep}, so regional level effects and the difference between narrow-range and wide-range indication are outside the model.

A feedwater pump trip halves the capacity at once. A real pump set coasts down. The instant halving is a stylized element that isolates the swell transient from the coastdown, which would otherwise blur it.

No engineered safety systems. Once the plant has scrammed and isolated, the application integrates decay heat into a closed vessel with relief valves. Emergency core cooling, reactor core isolation cooling and standby liquid control are absent, and nothing about the sequences that follow a scram is represented.

6.4 Coupling, Sign Reversal and the Operating Map

The four couplings that decide the control structure are all measurable on the running plant, and the application is arranged so that each can be isolated. The Pressure-only architecture leaves the pressure regulator in service and every other station on the operator, which is the configuration a plant test would use.

6.4.1 Recirculation Flow as the Power Handle

Slowing the recirculation pumps from 95 % to 80 % takes the core flow from 100.0 % to 86.18 %, raises the core-average void fraction from 40.32 % to 41.04 %, which the reported reactivity balance shows as −15.4 ¢, and settles the power at 90.90 %. The gain is 0.659 % of rated power per percent of core flow and 63 % of the change arrives in 3.3 s. The rod bank position at the end of the run differs from its starting value by less than one part in 10^9, so none of it came from a rod. That is the normal load-following actuator on the plant [6].

The gain is not constant. Evaluated as a local derivative along the 35 % rod line it runs from 0.632 % of power per percent of flow at rated conditions to 0.820 % per percent at 52 % flow, because the same fractional flow change removes a larger absolute void inventory when the channels are already voided. The gain from pump speed to core flow moves the other way, from 0.932 % per percent at 95 % drive to 0.790 % per percent at 46.7 % drive, because jet-pump head goes as the

square of the drive. The two effects very nearly cancel: the measured gain from pump speed to power is 0.589 % per percent at 95 % drive and 0.627 % per percent at 46.7 % drive, a change of 6 % across half the pump range. That is not what the jet-pump-squared argument on its own predicts, and Section 6.6.3 reports what it costs the gain-scheduling argument.

6.4.2 What Rods Actually Buy

Driving the bank in 7 % from the design insertion inserts 366 pcm, which is 65.4 ¢. Power falls only 11.46 % once settled, or 1.64 % of rated per percent of insertion, because the feedbacks return almost all of it. The core-average void fraction falls from 40.32 % to 38.17 %, worth +46.0 ¢, the fuel cools by 47.2 °C, worth +19.4 ¢, and total reactivity settles within 0.02 ¢ of zero. The rods did not move the operating point so much as move the rod line the operating point sits on. Timing sharpens the comparison. The bank travels at 0.55 %/s, so 7 % of insertion is 11.5 s to 90 % of the move against the 3.3 s in which the recirculation pumps delivered 63 % of a comparable power change.

Sweeping an external reactivity bias with pressure and level held and the pumps and rods fixed measures the self-regulation directly. Biases of +10 pcm, −10 pcm, +20 pcm and −20 pcm settle at 30.59, 30.66, 30.55 and 30.69 pcm per percent of power, and the symmetric difference at either amplitude is 30.62 pcm per percent, so a controller acting on this plant works in parallel with a proportional regulator of about 30.6 pcm per percent that it did not install and cannot switch off. The pressurized water reactor of Chapter 2 has the same kind of regulator in its two negative temperature coefficients, and there it answers a change in steam demand with no rod moving.

6.4.3 Pressure as a Reactivity Signal

Raising the dome pressure makes the steam denser, collapses the voids and inserts reactivity. Holding power and flow and evaluating (6.1) to (6.5) at a higher pressure gives +21.8 pcm at +0.05 MPa, +43.4 pcm at +0.10 MPa and +86.3 pcm at +0.20 MPa. That is why the turbine control valves are assigned to pressure rather than to power and why the pressure loop is tuned tighter than anything else on the plant.

Taking the pressure regulator out of service with the valves parked four percent further open depressurizes the dome at up to 17.7 kPa/s and takes the power from 100 % to 88.8 % by the time the dome reaches 6.30 MPa, an average of 1.28 % of rated per 0.1 MPa. The excursion does not stop there. Falling pressure lowers the power, lower power makes less steam, and less steam lowers the pressure further, and in this scenario nothing arrests it. The run reaches 5.60 MPa at 243.44 s, the main steam isolation valves latch shut at the next controller scan (243.45 s), and the pressurization that follows collapses voids fast enough to scram the plant on low vessel level in the very next 5 ms integration step. The high pressure trip arrives 7.7 s after that and the dome peaks at 8.366 MPa on the relief valves.

Whether the excursion runs at all turns out to depend on the feedwater rather than on the valve. Scenario 9 holds level on three-element control, so the feedwater tracks the falling steam flow and the core inlet subcooling falls with it, from 43.9 kJ/kg at 50 s after the park to 38.9 kJ/kg at 100 s. Repeating the same parked-valve experiment in MANUAL, where the feedwater is pinned at 100 % of rated, the subcooling instead holds at 44.3 kJ/kg as the steam flow falls away from it, and that reactivity arrests the slide. Parked at 89 % the MANUAL dome settles at 6.552 MPa with the power at 93.9 %, while the same park with the level loop in service is at 6.317 MPa 100 s after the park, still falling at about 7 kPa/s, and diverges from there. Across parked positions from 82 % to 90 % the MANUAL dome reaches a stationary value, 7.694 MPa at 82 %, 7.338 MPa at 84 %, 7.007 MPa at 86 %, 6.774 MPa at 87.5 % and 6.410 MPa at 90 %, each stationary to 0.002 MPa or better over the 50 s preceding its level trip, and a ± 0.15 MPa perturbation at the design position decays back to 7.170 MPa with a fitted time constant of 28 s.

What is not self-regulating in MANUAL is the inventory. A fixed feedwater flow against a changed steam flow is a standing mass imbalance, so every off-design park ends on a level trip and not on a pressure trip: Level 8 at 331 s from an 87.5 % park and Level 3 at 569 s from an 84 % park. The

pressure loop on this plant is held up in part by the level loop, and taking the level loop out changes which of the two diverges.

The reverse direction is faster and is the clearest single demonstration in the application of what the level measurement is worth. Parking the valves at 62 % from full power pressurizes the dome, collapses voids, and drops the indicated level −400 mm within 0.5 s. The Level 3 low-level scram fires 0.835 s after the valve command changes, at which instant the dome has risen only 0.057 MPa, the collapsed level has risen +24 mm and the vessel has gained 276 kg. A reactor is tripped on low water level, in under a second, while the vessel is filling.

The same collapse limits how fast the pressure setpoint may rise. Scenario 12 steps it from 7.17 to 7.22 MPa at 420 s, as shown in Figure 6.3. Under Coordinated PIC takes the step at once and the indication falls to −114.9 mm at 420.95 s, 1.7 s outside the band, with no water lost, while the pressure enters ±0.005 MPa of the setpoint at 423.7 s, overshoots 0.0086 MPa and stays inside from 440.70 s. The MPC of Section 6.5.6 predicts the collapse, walks p^R up (half the step by 427 s, 90 % by 442 s) and trims the feedwater, so the level dips only −17.7 mm, at 425.4 s, and the pressure settles without overshoot at 440.02 s, about when Coordinated does. Its cost is the integrated pressure error, 0.551 against 0.224 MPa s from 420 to 600 s.

6.4.4 Feedwater as a Reactivity Actuator

Losing a feedwater heater string cools the feedwater by 28 °C with nothing on the reactor touched. Core inlet subcooling rises from 48.87 kJ/kg to 67.37 kJ/kg, the boiling boundary moves up the channel, the core-average void fraction falls from 40.32 % to 39.57 %, which is +16.0 ¢, and the power rises to 109.45 %. Half the rise arrives in 17.5 s, 99 % of it in 152 s, and the power first reaches its final value 237 s after the loss, limited by the 30 s feedwater temperature lag and by the fuel node. At the new steady state the Doppler term reads −16.0 ¢ against the void term's +16.0 ¢, the fuel having warmed 38.9 °C, and the rod position is unchanged to five decimal places. A failure two hundred meters from the reactor building raises core power by a tenth, delivered entirely through the condensate system.

6.4.5 Shrink and Swell, Measured

Stepping the feedwater +10 % with the pressure regulator in service and the level loop out admits an extra 180.0 kg/s to a vessel holding 227.2 t. The mass balance alone predicts the level rising at 7.64 mm/s. The indicated level instead falls, reaching −40.9 mm at 3.45 s and not returning through its starting value until 6.10 s. At the instant of the minimum the collapsed level of (6.14) is +12.1 mm and the inventory is 193 kg heavier than it was, and the collapsed level rose monotonically throughout. Nothing about the inventory went the wrong way. Only the measurement did, because the colder inlet enthalpy of (6.11) shortened the boiling length of (6.3) and condensed steam that had been holding the surface up. Section 1.4.4 writes the same split for a drum as a water term and a void term, Equation (1.7). Power rises at the same time and peaks 4.47 % higher 5.5 s after the step, so even the feedwater valve is a reactivity actuator on this plant.

> **TRY IT IN THE APP** **2 · Shrink: The Level Goes the Wrong Way**
>
> Load scenario 2 and press Run on the Vessel view. At 30 s the feedwater steps to 110 %. Watch the solid indicated level against the dashed collapsed level: the indication falls to −40.9 mm at 3.45 s after the step while the collapsed level rises the whole time, and the indication does not return through its starting value until 6.10 s.

The opposite disturbance is larger and more dangerous. Halving the feedwater pump capacity drops the flow to 1170 kg/s against 1800 kg/s of steam, a deficit of 630 kg/s. The indicated level rises to +212.9 mm at 3.15 s, at which instant the collapsed level is −58.8 mm: the two readings are 271.7 mm apart and pointing in opposite directions. Less feedwater is less subcooling, a longer boiling length and more void, and the extra steam lifts the surface faster than the deficit lowers it.

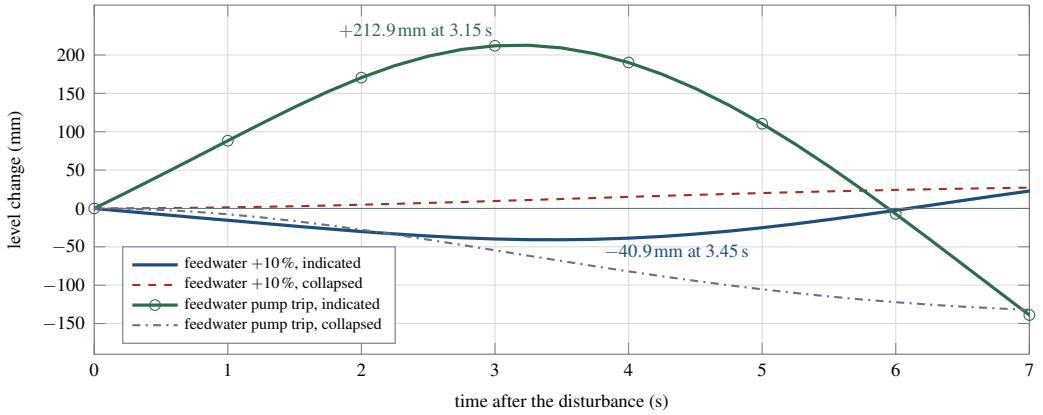


Figure 6.2 Indicated and collapsed level through both sign reversals, recorded from the running application and differenced against the pre-disturbance value. A +10 % feedwater step drives the indication −40.9 mm while the collapsed level rises monotonically, and a feedwater pump trip drives it +212.9 mm while the collapsed level is already 58.8 mm down. Both indications cross back through their starting value within 6.1 s, and in neither case does the inventory move in the direction the indication first takes. The feedwater step is drawn solid (indicated) and dashed (collapsed), the pump trip solid with circles (indicated) and dash-dotted (collapsed).

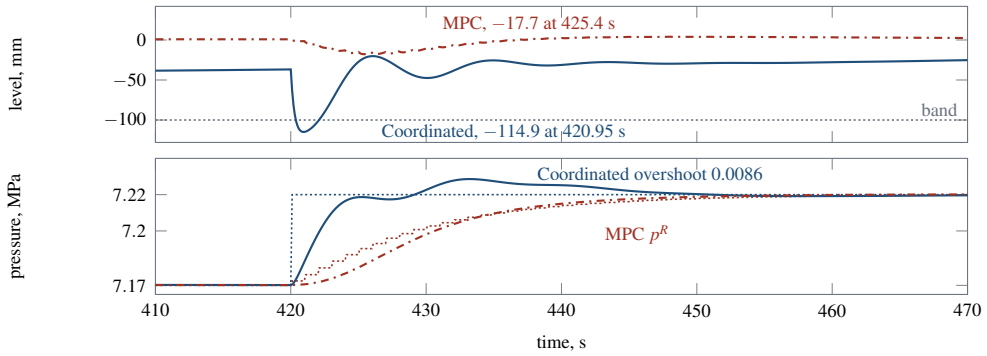


Figure 6.3 Scenario 12's pressure step under Coordinated (blue, solid) and the MPC (red, dash-dotted): level error, 36.8 mm low under Coordinated before the step, and pressure with the PIC setpoint each gives (dotted). The gray dotted line in the upper panel is the edge of the level band.

Then the mass wins. The indication crosses back through its starting value at 5.9 s, reaches −290 mm at 20 s, and the Level 3 scram fires 33.05 s after the trip.

Both disturbances are applied from an exactly stationary plant, so the pre-disturbance value serves as the held companion run and every trace in Figure 6.2 is one variable away from it.

Pause. Try Exercise 6.2 before reading on.

6.4.6 The Operating Map

Power and core flow are not independent, so the plant is licensed against a map rather than against a list of limits. Along the 35 % rod line the decay ratio of (6.16) rises from 0.250 at rated conditions to 0.283 at 90 % flow, 0.371 at 70 % and 0.495 at 52 % flow, at which point the power has fallen to 65.6 %. Stopping both recirculation pumps at the design rod line settles the plant at 50.4 % power on 34.6 % flow with a decay ratio of 0.717. That is just below the exclusion-region boundary $\mathrm{DR}_{\mathrm{lim}}$

of Table 6.2, in the configuration the LaSalle event was in [12]. Sweeping the rod bank from 70 % to 25 % inserted with both pumps stopped traces the natural-circulation line from 20.7 % power at 28.9 % flow to 59.5 % power at 35.5 % flow, so natural circulation alone supports well over half of rated power on about a third of rated flow.

> **DID YOU KNOW? If the Scram Had Failed**
>
> A laboratory simulation of the 1988 oscillation asked what would have happened had the reactor not scrammed. Doppler and void feedback limited the oscillations, but the power peaked at 1,300 percent of full power while its time average rose only to 86 percent [15].

6.5 Controller Design

6.5.1 Loop Pairing

Four loops close on this plant and the pairing follows from the physics rather than from a decoupling calculation. Dome pressure is paired with the turbine control valves because the valve path has the fastest and largest effect on pressure and because pressure is a reactivity term that must be held tightly. Vessel level is paired with feedwater because feedwater is the only inventory handle. Core power is paired with recirculation flow because that path is a factor of three faster than the rod path and costs no fuel duty. The rod bank is paired with the recirculation pump demand rather than with any plant variable, as a valve position controller that walks the pumps back toward the middle of their range so that the fast handle keeps authority in both directions. The couplings that remain are strong, which is why the comparison in Section 6.6 is between structures and not between tunings [4]. The reverse osmosis train of Volume 6, Chapter 2 is a plant on which the same pairing decision is made before any loop is tuned.

Pause. Try Exercise 6.1 before reading on.

6.5.2 Control Algorithm

All five controllers run the velocity-form PID of Section 1.6.2, executed every $\Delta t_c = 0.05\,\mathrm{s}$ and written here in the loops' own symbols. Each scan forms increments, as shown in Equation (6.18).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k) \\
\Delta u_k^{P} &= K_{c,k}\,e_k - K_{c,k-1}\,e_{k-1}, \qquad \Delta u_k^{I} = \frac{K_{c,k}}{\tau_I}\,e_k\,\Delta t_c
\end{aligned}
\tag{6.18}
$$

The output and the integral then follow the recursion of Equation (6.19).

$$
\begin{aligned}
s_k &= s_{k-1} + \Delta u_k^{P}, \qquad c_k = u_{\mathrm{bias}} + s_k + I_{k-1} \\
u_k &= \min(u_{\max},\ \max(u_{\min},\ c_k)) \\
I_k &= I_{k-1} + \Delta u_k^{I} + \frac{\Delta t_c}{T_t}\,(u_k - c_k)
\end{aligned}
\tag{6.19}
$$

where $a = +1$ is reverse and $a = -1$ direct action, $K_{c,k}$ is the proportional gain in force at scan k, τ_I the integral time, T_t the tracking time constant, $u_{\mathrm{bias}} = 0$, and the integral is seeded with the actuator position the plant starts from. The running sum s_k is the S_k of Equation (1.14), and the demanded output c_k is the v_k of Equation (1.15). Because each increment is the difference of two whole proportional terms, s_k is always exactly $K_{c,k}e_k$, so a gain that the scheduling of Section 6.5.4 moves acts on the whole proportional term at once. Anti-windup is the back-calculation in the last line, and s_k is never truncated. Derivative action is not exposed on any of the five loops, so $\tau_D = 0$. Manual and automatic transfer is bumpless both ways, as Section 1.6.2 describes. The one exception is the gain scheduling of Section 6.5.4, which acts only in AUTO, so a switch to AUTO first re-seeds the loops on the gains AUTO will use, as an architecture switch does.

Table 6.3 The Five Implemented Loops as Written, With the Default Settings Every Scenario Loads

Loop	PV → MV	a	K_c	τ_I (s)	T_t (s)
PIC	$p \to y_T$	−1	80 %/MPa	6	2
LIC	$L_{\text{ind}} \to W_{\text{fw}}$	+1	30 %/m	90	40
FIC	$W_{\text{fw}} \to$ valve	+1	1.5 %/%	3	5
RRC	$Q_{\text{th}} \to f_p$	+1	1.2 %/%	15	30
RWM	$f_{p,\text{cmd}} \to x_{\text{rod}}$	+1	0.05 %/%	400	240

Table 6.4 The Six Control Architectures and What Each One Closes

Architecture	Loops in service
Manual	none. Every station on the operator
Pressure only	PIC. The configuration a plant test uses
Single-element FW	PIC, and LIC writing the feedwater demand directly
Three-element FW	PIC, FIC on feedwater flow with steam flow fed forward, and LIC trimming the flow setpoint
Coordinated	the above plus RRC on power, RWM behind the pumps, and gain scheduling on power and pump speed
MPC	one plan each second for the pump and feedwater demands and the PIC setpoint, RWM behind the pumps, Coordinated as fallback

The five loops and their default settings are listed in Table 6.3. The pressure loop is direct acting in the ISA sense, action −1 in the convention of Section 1.3, because opening the valves lowers the pressure. The level, feedwater flow, power and rod loops are reverse acting, action +1. The rod controller reads the pump *demand* rather than the pump speed and is deliberately the slowest controller on the plant.

Output limits are part of the design. PIC spans 0–100 % of valve travel. FIC and LIC in single-element mode span 0–130 % of rated steam flow, further limited to the available feedwater pump capacity and forced to zero once the Level 8 feedwater trip has latched. In three-element mode LIC is instead clamped to ±75 %, because its output is then a trim on the feedforward rather than the demand. RRC spans 28 % to 100 % of pump drive, the lower bound being the speed below which the jet pumps stop controlling the flow, and RWM spans the full rod bank. Two limits act after a controller output, where its integral cannot see them. In three-element mode the sum of the feedforward and the level trim is clipped to 0–130 % and to the pump capacity outside LIC, which knows only its own ±75 %. Rate limiting lives in the plant rather than in the controllers: the rod bank of (6.6) is limited to 0.55 %/s and to 35 %/s on a scram, and the pump drive, feedwater flow and valve position reach their commands through the lags τ_p, τ_{fw} and τ_{TCV} of Table 6.2.

6.5.3 Six Architectures on One Actuator Set

The application selects among six structures, listed in Table 6.4, all of them writing to the same four actuators through the same limits at the same scan.

Three-element control writes the feedwater flow setpoint as shown in Equation (6.20).

$$W_{\text{fw,SP}} = K_{\text{ff}}\,\widehat{W}_{st} + u_{\text{LIC}}, \qquad u_{\text{FIC}} = \text{PID}\left(W_{\text{fw,SP}},\ \widehat{W}_{\text{fw}}\right) \tag{6.20}$$

where $\widehat{W}_{st}$ and $\widehat{W}_{\text{fw}}$ are the filtered steam and feedwater flows as percentages of the rated value, K_{ff} is the feedforward gain and u_{LIC} is the level controller's trim. At $K_{\text{ff}} = 1$ every kilogram leaving the vessel is replaced as it leaves. At $K_{\text{ff}} = 0$ the structure reverts exactly to single element, and any value in between leaves the level controller supplying the balance out of an error signal it measures backwards.

Because the level controller's output means the demand in one architecture and a trim in the other, and the scheduled gains change with the architecture, switching would otherwise bump the actuators. At every change of architecture, feedforward gain, scheduling or rod controller, the application puts the level, flow, power and rod controllers on the gains their next scan will use and makes each track the command the plant already has (less the feedforward for the level trim). Tracking sets s_k to the proportional term of the present error and I_k to the balance, so the next scan adds only the change in error since the switch. From a lined-out full-power state, no transition among the five loop architectures moves any actuator command by more than 0.023 % on its first scan, the same order as the loops' own motion between scans. During a transient the switch is as quiet. Switching from single to three element at 200 s of the single-element load change, at 85.4 % power with the indication 278 mm high, leaves the feedwater command at 86.83 % on the first scan, and three-element control then trims it to no less than 84.00 % over the next 10 s. A seed that cleared s_k instead would add the whole $K_c e_k$ on the first scan and drop the same command to 73.57 %.

6.5.4 Gain Scheduling

The coordinated architecture schedules two gains on measured states, as shown in Equation (6.21).

$$\begin{aligned} K_c^{\mathrm{LIC}} &= K_c \cdot \mathrm{sat}\left(\frac{0.72}{n}, 0.6, 2.2\right) \\ K_c^{\mathrm{RRC}} &= K_c \cdot \mathrm{sat}\left(\frac{f_{p,0}}{f_p}, 0.7, 3.0\right) \end{aligned} \tag{6.21}$$

where $\mathrm{sat}(\cdot, l, h)$ clamps to the bracket, n is the neutron density and f_p the pump drive fraction. The level term detunes the loop as power rises, on the argument that there are more voids to condense and the right-half-plane zero moves toward the origin. The recirculation term raises the controller gain as the pumps slow. Section 6.6.3 measures what each contributes, and the second term does not survive the measurement.

6.5.5 Protection

Five trips and two isolations are implemented and each is armed in at least one scenario. The Level 3 low-level scram at −450 mm and the Level 8 feedwater and turbine trip at +560 mm both fire in the feedwater pump trip scenario. The high-pressure scram at 7.72 MPa and the APRM high flux scram at 120 % both fire in the turbine valve scenario, as do the main steam isolation at 5.60 MPa and the relief valves at 7.86 MPa. The oscillation scram counts the seconds the decay ratio spends at or above unity, drains the count at 0.4 s per second below it, and fires when the count passes 15 s, which happens in the instability scenario and is the first trip to act there.

6.5.6 Model Predictive Control

The sixth architecture replaces RRC, LIC and FIC with one model predictive controller (MPC) that plans, every $T_s = 1$ s, the pump demand $f_{p,\mathrm{cmd}}$, the feedwater demand $W_{\mathrm{fw,cmd}}$ and the setpoint p^R handed to PIC, which still moves the valves at 20 Hz, with RWM behind the pumps. Its model is Section 6.3 at the declared values of Table 6.2 and the design disturbances, never the run's, so an edited parameter is model mismatch, and it reads only the loops' transmitters, commands and setpoints. Its 21 states are the plant's without T_{fw}, z_o and z_v, the three transmitter lags, the PIC integral at its live tuning, and unmeasured inputs d_ρ, d_E and d_M added to the reactivity, energy and mass balances, so the void feedback of (6.6) and the level of (6.15) are in every prediction. Each decision linearizes it about the estimate, with an exact hold and the rods frozen. Against ±5 % pump and feedwater steps and ±0.05 MPa setpoint steps from full power, 120 s each, one decision's linearization stays within 0.062 %, 1.98 mm and 0.000 35 MPa of the plant, close because the model is the plant's own balance.

Each plan minimizes Equation (6.22) over $N = 120$ steps in twelve move blocks, from 1 step to 75.

$$\mathscr{J} = \sum_{k=1}^{N} \left[q_Q e_{Q,k}^2 + q_L e_{L,k}^2 + q_p e_{p,k}^2 \right] + \sum_i \left(a_i \varepsilon_i^2 + b_i \varepsilon_i \right) + \lambda \sum_j \left[\Delta f_{p,j}^2 + \Delta W_{\mathrm{fw},j}^2 + \left(\Delta p_j^R / 0.01 \right)^2 \right] \tag{6.22}$$

The errors are power in %, indicated level in mm and pressure in MPa less their setpoints, with $q_Q = 1$, $q_L = 0.01$ and $q_p = 10^4$, so a percent, 10 mm and 0.01 MPa cost alike, and Δ is a move between blocks, priced at $\lambda = 100$. The bounds are hard: RRC's 28–100 %, FIC's feedwater limit, p^R inside the pressure band, and 2 %/s, 5 %/s and 0.01 MPa/s, never active in scenario 12 (largest moves 0.75 %, 0.27 %, 0.003 75 MPa), where λ alone smooths the moves. One slack ε_i each softens the ± 100 mm level band, the ± 0.10 MPa pressure band and the licensed 100 % power, at a_i of 1 per mm^2, 10^6 per MPa^2 and 1.2×10^7 per percent and b_i of 10 per mm, 10^4 per MPa and 1.2×10^7 per percent, judged every second step. A level or power bound already violated returns from the present value with a 30 s or 10 s time constant. Each decision is one receding-horizon quadratic program [17, 18] of 39 variables and 447 rows, solved by a dual active-set method [19] in a mean 3.7 ms in Chrome, 0.37 % of a core at real time. Section 1.6.3 sets out the shared form of the controller and its estimator.

A steady-state Kalman filter at the declared design point corrects the model every scan under every architecture from the three transmitters (σ of 0.002 MPa, 5 mm and 0.3 %), with d_ρ, d_E and d_M as random walks of variance $\rho_K (0.05\,\mathrm{pcm})^2$, $\rho_K (500\,\mathrm{kW})^2$ and $\rho_K (0.5\,\mathrm{kg/s})^2$ per scan and $\rho_K = 10$, which removes offset [20]. A grid pass on scenarios 12 and 8 chose λ and ρ_K among settings that held both bands without a trip or more travel than Coordinated. The Coordinated loops track the applied commands with no preload and take over, without a bump, whenever a plan is rejected as suboptimal, nonfinite or out of bounds, until the next valid plan. On that or a switch away p^R returns to the operator's setpoint at 0.01 MPa/s. A scram, an isolation or a Level 8 trip, whose logic reads no architecture, stands the MPC down for the run, so it solves and writes nothing after it, and a move of the pump station (runback, pump trip or operator) pins the pumps until Reset.

6.6 Results

Table 6.5 lists all thirteen scenarios in the order they are meant to be run, the mechanism each one shows and what completes it. Each one presets the architecture, all five sets of tuning constants, the disturbance sliders and whatever timed events it needs. Briefings, completion credit, Reset and the challenge rules are those of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance, including firing a disturbance button, voids the attempt, while the prescribed scenario events remain fair.

6.6.1 Single-Element Against Three-Element Control

Scenarios 5 and 6 are the same plant, the same recirculation ramp from 95 % to 70 % pump drive in ten steps over 60 s, and the same level gains of $K_c = 30\,\%/\mathrm{m}$ and $\tau_I = 90\,\mathrm{s}$. The only difference between the two runs is where the steam flow measurement goes. Both end at 77.2 % core flow, the single-element run at 83.9 % power with the steam flow 17.0 % below rated and the three-element run at 84.7 % power with the steam flow 15.4 % below. Table 6.6 reports the outcome.

Single element reaches +287.4 mm and spends 221.6 s outside the control band, against −63.0 mm and 0.0 s for three element, a factor of 4.6 on the same disturbance with the same gains. The mechanism is worth following. Power comes off, so steam flow comes off, so the level should rise, and it does, but slowly at first, because the falling power also means less steam below the surface, and less steam pulls the indication down. The single-element controller therefore sees almost nothing for the first thirty seconds and keeps feeding at full rate, and by the time the mass imbalance appears in the indication it is minutes of overfeeding old. The collapsed level confirms it: the vessel really took on 428.8 mm of water it did not need, against 166.4 mm for the three-element run. Figure 6.5 is that

Table 6.5 Scenarios and Completion Requirements

#	Scenario (kind)	Mechanism	Required to complete
1	Full power by hand *experiment*	four actuators, one internal state	In Manual, rods at 33 % with power above 101.5 % for 10 s, pumps at 90 % with power below 98 % for 10 s, turbine valves parked at 86 or 87.5 % until the dome pressure moves less than 0.2 kPa/s for 30 s, then feedwater at 105 % for 2 min.
2	Shrink: the level goes the wrong way *experiment*	shrink against the inventory	With the pressure regulator in service, the indication below −20 mm while the collapsed level is above zero within 10 s of the 110 % feedwater step at 30 s, then the run continued to 100 s through the return to 100 %.
3	Power without moving a rod *experiment*	flow as the power handle	Pumps at 80 % with no rod travel for 30 s between 150 and 220 s, then the run continued to 300 s through the return to 95 %.
4	The same power change with rods *experiment*	self-regulation and rod speed	The bank seen in transit from 35 to 42 % inserted, then total reactivity within 1 ¢ of zero for 20 s after 250 s.
5	Single-element feedwater through a load change *demonstration*	inverse response inside the loop	The indication above +200 mm after the 95–70 % pump ramp, and the full 400 s run.
6	Three-element feedwater, same load change *experiment*	the structural fix	The same ramp run to at least 300 s in three-element control at $K_{\text{ff}} = 1.0$ with less than 5 s outside ±100 mm.
7	Feedwater pump trip *demonstration*	swell while the inventory falls	The indication above +100 mm while the collapsed level is below zero within 10 s of the trip, the Level 3 scram, and 30 s after it.
8	Loss of feedwater heating *experiment*	reactivity through the condensate system	Power above 108.5 % for 60 s after the heater loss, then a switch to Manual and pumps lowered until power stays below 103 % for 10 s.
9	A pressure step is a reactivity step *demonstration*	pressure as reactivity, no equilibrium with the level loop in service	The parked-valve run until the isolation valves have closed and the reactor has scrammed, then Reset and the regulator-fails-closed event, ending in a Level 3 scram with the isolation valves open and no high-pressure trip.
10	Recirculation pump runback *experiment*	the map, and which way a runback moves	The indication below −150 mm between 40 and 80 s, then the decay ratio above 0.4 for 20 s after 200 s.
11	Into the exclusion region *demonstration*	the plant, not the loop, going unstable	The decay ratio above 0.7 for 20 s on the first plateau (150 to 220 s), then the run continued to 300 s through the second flow reduction.
12	Void feedback in the model *experiment*	three handles planned against the swell	The full 600 s under MPC, then Coordinated, Reset and the 600 s again: the load ramp of the challenge with every loop set up well, and a pressure setpoint step at 420 s.
13	Challenge: 70 to 100 % on recirculation flow *challenge*	structure against tuning	Over 600 s with the model parameters and disturbances as given, the indicated level within ±100 mm and the dome pressure within ±0.10 MPa throughout, the decay ratio below 0.60, and the power within 0.5 % of 100 % from 480 s.[a]

[a] As given, the challenge fails with a worst indicated level of 410 mm, four times the band, a worst pressure error of 0.259 MPa, a peak decay ratio of 0.47 and the power still 26.15 % from 100 % over the last 2 min. With the suggested configuration it passes with 67 mm (33 mm inside the band), 0.021 MPa, 0.45 and 0.07 %.

run at its worst point, with both surfaces drawn on the vessel view and the two readings 141 mm apart. Single element is not badly tuned here. It is badly informed, and the fact that would have solved it was sitting on a transmitter the loop was not connected to.

Table 6.6 Single-Element Against Three-Element Feedwater on an Identical 95–70 % Recirculation Ramp With Identical Level Gains, Read Over 400 s

	Single element	Three element
Worst indicated level (mm)	+287.4	−63.0
Time of that worst value (s)	221	130
Seconds outside ±100 mm	221.6	0.0
Indicated level at 400 s (mm)	+33.2	+2.7
Worst collapsed level (mm)	+428.8	+166.4
Feedwater demand range (%)	80.1–100.2	84.3–100.0

> **TRY IT IN THE APP** **6 · Three-Element Feedwater, Same Load Change**
>
> Load scenario 6, press Run and let the recirculation ramp from 95 % to 70 % play. Read the worst indicated level, −63.0 mm, and the time outside ±100 mm, none. Then load scenario 5, the same ramp on single-element control, and run all 400 s: +287.4 mm and 221.6 s outside the band.

What did not change is instructive. At 395 s the three-element run holds the indication at +3 mm while the collapsed level reads +147 mm, a standing gap of 144 mm, because at 85 % power the vessel needs more water to show the same level. The gap does not close, and it is not supposed to. Three-element control does not remove the non-minimum-phase behavior from the plant. It removes it from the feedback path.

6.6.2 Runback and the Exclusion Region

Running the recirculation drive back to 45 % takes the core flow from 100 % to 56.2 % and the power down the 35 % rod line from 100 % to 68.9 %, with no rod motion. The operating point moves further left than it moves down. The power-to-flow ratio rises from 1.00 to 1.23 and the decay ratio from 0.250 to 0.460, with a transient peak of 0.501. The level dips −247 mm at 14.2 s on the way, while the collapsed level stands above its starting value. The core-average void rises as the flow comes off, but the falling power cuts the steam passing through the separators, and the steam held below the surface falls with it. The feedwater that the three-element controller cuts to match the falling steam flow does not deepen the dip, because the vessel is gaining water at that instant. A runback is normally a protective action that prevents something worse, and this is its price: every runback moves the plant toward the corner of the map where it is least stable, and it gets there by removing power.

Under the MPC the runback pins the pumps at 45 % at 40.005 s, the scan it acts under three-element, and the indication rises to +114.9 mm, 5.7 s outside the band against three-element's 35.3 s. Under Coordinated the runback and the trip of both pumps also pin the pumps on their station until Reset, with the power loop and the rod controller standing aside. Fired at 400 s of the suggested challenge, the runback takes the power to 72.4 % (MPC 70.8 %). The trip of both pumps scrams Coordinated on low level at 412.5 s, while the MPC keeps the level within 241 mm.

Scenario 11 goes the rest of the way, from a high rod line with the ordinary boiling noise switched on. Table 6.7 reports the three plateaus.

At the first plateau the decay ratio is 0.765, the exclusion-region warning has already annunciated at 92 s, and nothing is out of limits. The void oscillation amplitude has grown by a factor of 1.37, from 0.281 % to 0.385 % peak to peak, but the power noise has grown only from 5.99 % to 6.41 % peak to peak and its standard deviation from 1.23 % to 1.24 %, so the growth is visible on the void and not yet on the flux. The second reduction is unambiguous. The decay ratio crosses unity 14 s after the flow reduction, the power oscillation grows from 5.4 % to 13.6 % peak to peak over the following 15 s, and the oscillation scram fires 29 s after the reduction. That is 15 s after the decay ratio first reached unity, because it stayed there and the trip logic counts 15 s of it. The low-level scram follows 4 s later, so the oscillation trip is the one that acts.

Table 6.7 Two Recirculation Reductions From a High Rod Line, With Boiling Noise at 1.5×. Values Over 20–60 s, 150–220 s, and From the Decay Ratio Reaching Unity to the Oscillation Trip (234–249 s).

	Initial	Plateau 1	Plateau 2
Pump demand (%)	62	35	12
Core flow (%)	70.9	49.8	38.3
Thermal power, mean (%)	99.2	80.1	68.4
Decay ratio	0.504	0.765	1.00–1.07
Power, peak to peak (%)	5.99	6.41	13.8
Power, std. dev. (%)	1.23	1.24	3.10
Void oscillation, peak to peak (%)	0.281	0.385	0.952

Table 6.8 The Challenge Ramp With One Change at a Time From the Configuration as Given. Every Row Is the Same 600 s Disturbance Sequence, and Off Is the Time the Indicated Level Spends Outside ±100 mm.

Change from as given	Worst L (mm)	Worst p error (MPa)	Power range (%)	Off (s)
none	410	0.2589	59–113	412
K_{ff} 0.6 to 1.0	384	0.2524	60–113	404
gain scheduling on	3097	1.5051	11–109	540
rod controller on	77	0.1661	67–100	0
six tunings restored	170	0.0216	70–100	199
tunings and K_{ff}	47	0.0217	70–100	0
all four (suggested)	67	0.0206	66–100	0
architecture to MPC	13	0.0459	70–100	0

What made the plant unstable was the flow and not the power. The power fell at every step. That is why the exclusion region occupies the upper-left corner of the map and why it is entered by reducing flow, and it is a control problem in which the plant rather than the controller is the thing that goes unstable [5, 12].

> **✦ DID YOU KNOW? Less Power, Less Stable**
>
> From the simulator: as the recirculation flow is cut twice from a high rod line, the mean power falls from 99.2 to 68.4 percent, yet the decay ratio climbs from 0.504 to between 1.00 and 1.07 and the power swing grows from 5.99 to 13.8 percent peak to peak (Table 6.7).

6.6.3 Graded Challenge: Structure Against Tuning

The challenge ramps the power setpoint from 70 % to 100 % in six steps between 60 s and 260 s and asks for the indicated level held inside ±100 mm, the dome pressure inside ±0.10 MPa and the decay ratio below 0.60 for ten minutes, with the power within 0.5 % of rated over the last two. The coordinated architecture is selected and configured badly on purpose: the feedforward gain is 0.6 instead of 1.0, the level loop runs at 130 %/m with a 20 s reset instead of 30 with 90, the power loop at 4.5 %/% with a 60 s reset instead of 1.2 with 15, the pressure loop at 22 %/MPa with a 30 s reset instead of 80 with 6, and both the gain scheduling and the rod controller are switched off. Table 6.8 changes one thing at a time from that starting point, each applied at $t = 0$ before the run.

The suggested configuration meets the target, with a worst indicated level of 67.0 mm, a worst pressure error of 0.0206 MPa, a peak decay ratio of 0.450 and 100.03 % power reached at 97.0 % core flow. The only rod motion in the run is the valve position controller walking the bank from 35.00 % to 33.73 %, a total travel of 6.07 %, which is what relieves the pumps and leaves the fast handle authority in both directions. Reading the table row by row separates the four fixes, and two of

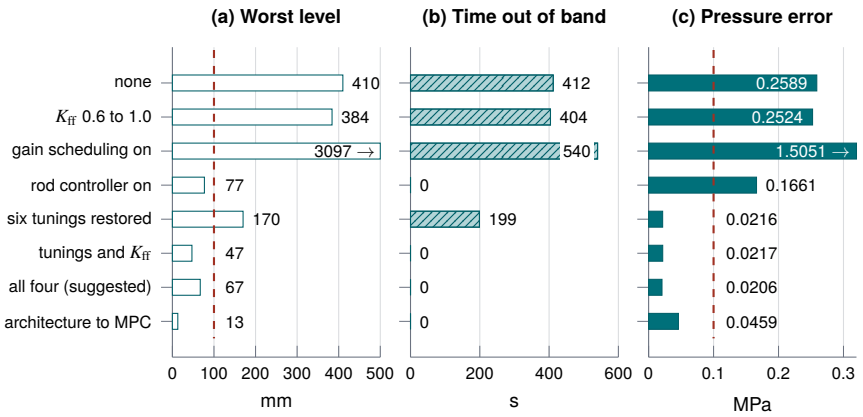


Figure 6.4 Table 6.8 drawn as bars: the challenge ramp with one change at a time from the configuration as given. (a) Worst indicated level, with the 100 mm band as the dashed line. (b) Time the indicated level spends outside ±100 mm. (c) Worst pressure error, with the 0.10 MPa band as the dashed line. The gain scheduling row runs off panels (a) and (c), so its bars end at the edge with the value printed. Notice that the rod controller alone does more for the level than the six restored tunings, and that only the last three rows stay inside both bands.

them do not behave the way the usual argument for them predicts. Figure 6.4 draws the three columns of the table as bars.

Restoring the six tuning constants does the pressure work on its own, taking the worst pressure error from 0.2589 MPa to 0.0216 MPa, and about half the level work, from 410 mm to 170 mm. Adding the feedforward on top finishes the level at 47 mm, which is better than the suggested configuration achieves. The feedforward alone with the original gains buys almost nothing, 384 mm against 410 mm, because a power loop oscillating between the pump stops generates a level disturbance far larger than the feedforward corrects. Turning the rod controller on alone is the single most effective change, 410 mm to 77 mm, because keeping the pumps away from their limits stops the excessive power gain from slamming them.

Turning the gain scheduling on alone is catastrophic. The plant scrams on low level at 103.7 s and trips on Level 8 at 145.9 s, and with the post-scram void collapse included the run records a worst level of 3097 mm, a worst pressure error of 1.5051 MPa and the power falling to 11 %. Equation (6.21) multiplies whatever gain it is given, and at the starting 46.7 % pump drive the recirculation factor is 2.03, so an already excessive K_c of 4.5 becomes 9.15. With the gains and the feedforward correct the scheduling makes no useful difference either, giving 102.6 mm against 46.5 mm with the rod controller off and 67.0 mm against 67.5 mm with it on. That agrees with Section 6.4: the process gain from pump speed to power changes by 6 % across half the pump range, not by the factor of two the scheduling law is scaled for.

Switching only the architecture to the MPC of Section 6.5.6 also meets the target, with a peak decay ratio of 0.450, at a pressure IAE of 9.0 against 2.0 MPa s for the suggested configuration, because the slow pressure loop as given still sits under it. The challenge is unchanged.

6.6.4 The Design Point and What the Application Verifies

Every scenario opens at the state in Table 6.9, which the Newton solve of Section 6.3 reaches to machine precision. Holding the design scenario for 60 s moves no reported quantity in the fifth decimal place.

Thirteen application-specific checks run headlessly on the shipped model, in addition to the framework's own structural, bumpless-transfer, actuator-bound, CSV-consistency and reproducibility

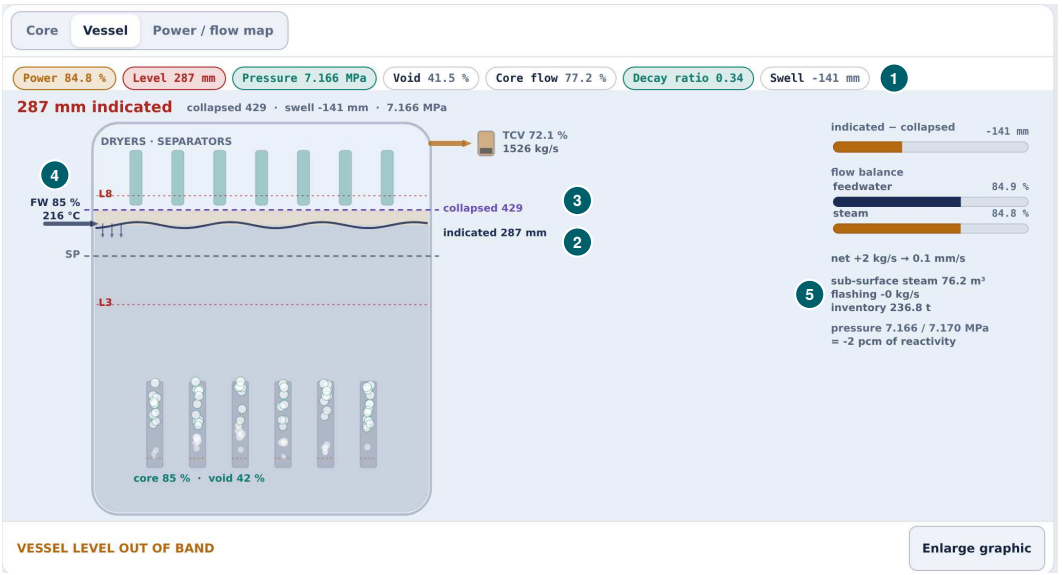


Figure 6.5 Single-element feedwater control in the Vessel view of the app at $t = 3.68$ min, at the worst point of the 95–70 % recirculation ramp of scenario 5, which is the failure the comparison of Section 6.6 rests on. Callouts: (1) the status tiles, (2) the indicated level, (3) the collapsed level, (4) the feedwater inlet with its demand and temperature, (5) the readouts of sub-surface steam and inventory. The vessel drawing carries both surfaces and they disagree by 141 mm: the level transmitter reads 287 mm against a setpoint of zero while the collapsed level, which is what the water inventory actually is, stands at 429 mm. The controller is acting on the smaller of the two numbers, so it is feeding a vessel that is already 236.8 t full, and the scene reports the indication out of band. The plant holds 84.8 % power on 77.2 % core flow at 7.166 MPa, and the void fraction has risen to 41.5 % as the flow came off. The turbine valves have closed to 72.1 % against a steam flow of 1526 kg/s, feedwater and steam are within 2 kg/s so the surface is nearly stationary at 0.1 mm/s, and the 7.166 against 7.170 MPa pressure error is annotated on the panel as the −2 pcm of reactivity it is worth.

Table 6.9 The Design Operating Point, Solved Rather Than Assumed

Quantity	Symbol	Value	Quantity	Symbol	Value
Core thermal power	Q_{th}	3330.0 MW	Core exit void fraction	$\alpha(x_e)$	65.23 %
Generator output	P_e	1100.1 MW	Core-average void fraction	$\bar{\alpha}$	40.32 %
Dome pressure	p	7.170 MPa	Non-boiling fraction	f_{nb}	10.57 %
Saturation temperature	T_{sat}	287.49 °C	Average fuel temperature	T_f	699.49 °C
Core flow	W_c	13 100 kg/s	Rod bank insertion	x_{rod}	35.0 %
Core mass flux	G	1506 kg/m^2/s	Recirculation drive	f_p	95.0 %
Steam flow	W_{st}	1800.2 kg/s	Turbine valve position	y_T	85.0 %
Feedwater enthalpy	h_{fw}	920.9 kJ/kg	Sub-surface steam volume	V_{sd}	80.71 m^3
Core inlet subcooling	Δh_{sub}	48.87 kJ/kg	Sub-surface steam as level	V_{sd}/A_s	2522 mm
Core inlet subcooling	—	9.97 °C	Vessel inventory	M_v	227.2 t
Core exit quality	x_e	13.742 %	Decay ratio	DR	0.250

tests. They verify the equilibrium, the design point, the sign of every actuator path, the pressure-to-reactivity coefficient, the direction of every closed loop, the shrink transient, the flow-to-power gain with the rod bank frozen, the single-element against three-element comparison, the location of the

Table 6.10 Coordinated Loops Against the MPC at Scenario 12's Settings, Same Seed, Events and Noise. IAE in % s, mm s and MPa s and Travel From the 10 Hz Historian

Case and measure	Coordinated	MPC
Scenario 12, ramp, 0–420 s		
Worst level, 60–420 s (mm)	49.3	27.6
IAE power, level, pressure	438.0, 10 124, 2.016	500.7, 5788, 0.985
Pump travel (%)	140.7	47.2
Within 0.5 % of 100 % from (s)	314.6	302.4
Scenario 12, pressure step, 420–600 s		
Level dip (mm)	−114.9	−17.7
IAE pressure	0.224	0.551
IAE power, level, pressure over 600 s		
Scenario 9, regulator out	143.6, 3163, 182.9	93.6, 1820, 177.9
Case 1, admission −5 %	72.9, 1585, 0.481	56.6, 1712, 0.579
Scenario 8, heater loss	194.2, 5885, 0.507	193.9, 4297, 0.765

exclusion region, the equilibrium of the three-element scenario before its ramp, that an architecture or AUTO switch on a running plant is bumpless, and that the challenge target is met by the suggested configuration and missed by the one given. Fourteen more check the MPC. All forty-seven checks pass on the shipped code, as do the fourteen checks that walk each scenario's completion rule.

6.6.5 Model Predictive Control Against the Coordinated Loops

Scenario 12 is the challenge ramp with every loop set up well, then the pressure step of Figure 6.3, and Table 6.10 sets Coordinated against the MPC. No run in it trips or rejects a plan.

On the ramp the MPC cuts the level IAE by 43 %, the pressure IAE by 51 % and pump travel by two thirds, but loses the power IAE, 500.7 against 438.0 % s, because λ prices each pump move. The 67 mm of Table 6.8 is that Coordinated run's opening rod move at 4.7 s, not the ramp. On the admission step and the heater loss it loses pressure IAE, moving p^R to help power and level. With 10 mm of level noise and boiling noise at 1× it loses power IAE, 724.5 against 621.6 % s, and moves the feedwater 274 against 165 %.

Every trip is the plant's own, so an instant moves only with the plant's path. With the regulator failing closed at 150 s of scenario 12 the MPC scrams on high pressure later, at 189.08 s against 179.81 s, because its plan has no pressure term once the valves are parked, and after each of the six trips measured under it the MPC solved no plan. Switched from three-element at $t = 0$ it peaks at 105.29 % power, against 104.39 % into Coordinated, below the 120 % flux scram, so only the power tile turns amber.

6.7 Technology in Practice

The tags of Figure 6.1 are, in a plant, pressure, level and flow transmitters, neutron detectors in the core, and the controllers of a plant control system. Figure 6.6 draws the four loops with that hardware in place of the schematic's symbols.

What Measures It

Pressure transmitters convert the pressure force to an analog signal such as 4 to 20 mA or, in newer smart transmitters, to a digital signal, and the most common level measurement in nuclear plants is a differential pressure between two taps on the vessel. Neutron detectors range from fission and ion chambers to self-powered detectors, and the readings of many in-core detectors are combined to calculate the overall reactor power [21]. Measurements in the vessel of a boiling water reactor that monitor adequate cooling of the core have their own standard, IEC 61343 [21]. The chapter carries the pressure transmitter as a 0.15 s lag, the level transmitter as a 1.2 s lag on Equation (6.15), the

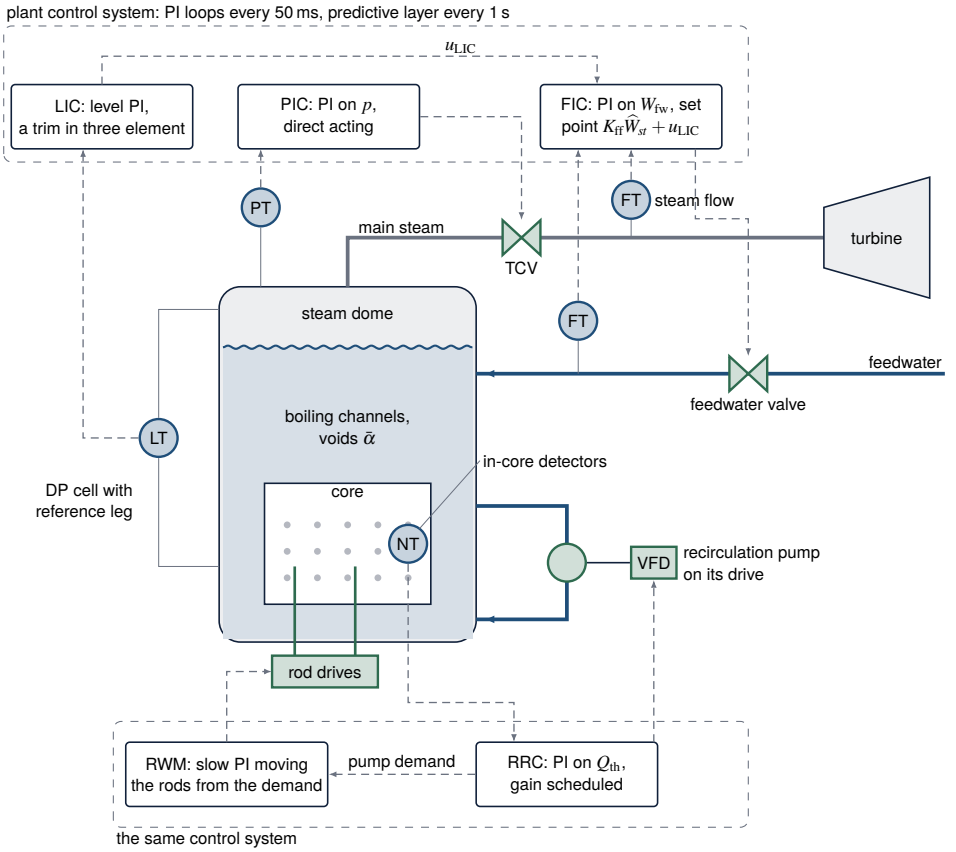


Figure 6.6 The four loops of the chapter with plant hardware in place of the symbols of Figure 6.1. Sensors are blue circles and final elements green: the turbine control valves (TCV) under PIC, the feedwater valve under FIC with steam flow fed forward and LIC trimming its set point, the recirculation pump on its variable-frequency drive (VFD) under RRC, and the rod drives under the slow valve position controller RWM, which reads the pump demand. Notice that the level transmitter's two taps straddle the boiling surface, so the reading moves with the steam below it.

power signal through a 0.6 s filter and the flow transmitters as 0.8 s lags. Figure 6.7 shows why the level reading is a head and not a height.

> **TECHNOLOGY Sensor: The Level Transmitter and Its Reference Leg**
>
> Most level measurements in nuclear plants are differential pressures between a tap low on the vessel and a tap high on it, with a filled reference leg on one side [21]. The pressure of a liquid column gives its height only if the density of the liquid is known and constant. The reference leg holds a constant head only while its own density stays constant, and a zero shift at calibration cancels that head. With the leg on the high port the output falls as the level rises, so the connection is chosen for a fail-safe reading [22]. An impulse line that clogs, traps gas or fouls makes the reading inaccurate or sluggish [21]. The chapter's LT converts with the densities of its ranging pressure, behind a 1.2 s lag.

What Moves It

The final elements are the turbine control valves, the feedwater supply, the two recirculation pumps and the control rod bank. In the chapter the valves stroke as a 0.25 s lag, the feedwater flow follows its demand through a 3.5 s lag and the pump drive through a 5.0 s lag, while the rods are rate limited at

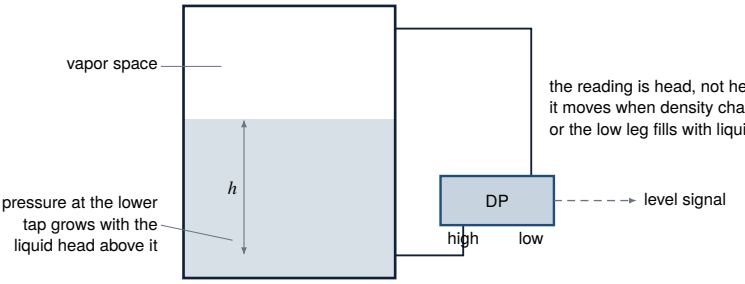


Figure 6.7 Level by differential pressure, from the book's instrument library. Notice that the transmitter reads the head between its two taps, which depends on the density of what lies between them. In this chapter that is a mixture of water and steam, which is why the indicated level of Equation (6.15) can fall while the vessel fills.

0.55 %/s. These are idealizations with no dead band and no slew limit on the pumps or the feedwater, as Section 6.8 states.

> **TECHNOLOGY** **Final Element: The Recirculation Pump on a Variable-Frequency Drive**
>
> Changing the flow through the core is the normal way to move a boiling water reactor's power between about 30 and 100 percent along its highest rod line, because more flow sweeps the voids out faster. The pump speed comes from a variable-frequency drive, an inverter that sets the motor's supply frequency and keeps the voltage in proportion. Such drives have replaced the older motor-generator sets, and one plant reported a fourfold rise in the reliability of its recirculation system after the change. Newer digital drives carry redundant power cells: a drive can lose two and still deliver full power, and a failed cell is swapped in eight cycles with a small loss of speed [23]. The chapter's drive follows its command through the 5.0 s lag τ_p.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, and nested loops in which the set point of one loop comes from the measurement of another [24]. The chapter's loops execute every 0.05 s over a 5 ms integration step, and the predictive controller plans every second and hands its pressure setpoint to PIC. The structure is a set of function blocks: a direct-acting PI on dome pressure, a level PI whose output is added to the steam flow feedforward to set a feedwater flow PI, a power PI on the pump demand, a slow valve position PI that moves the rods from that demand, and a gain scheduler on two of them.

> **IN PRACTICE** **Power Changes Ride on the Recirculation Flow**
>
> A boiling water reactor that operates flexibly can follow load on recirculation flow alone, along a curve of power against core flow at a fixed rod position. This suits the upper range, about 60 to 100 percent of rated power, where the power changes faster and with fewer changes in its distribution across the core. With no rods moving, the fuel sees smaller temperature swings, and the pressure and temperature in the main steam system stay constant. Changes beyond the recirculation range are made with the control rods [25]. The chapter's rod controller moves the bank only slowly, to walk the pumps back toward the middle of their range.

What Is Modernizing It

Three developments reach these loops. Variable-frequency drives have replaced motor-generator sets on recirculation pumps, and twenty-six were in operation in boiling water reactors in the United States by 2014 [23]. On-line monitoring mines the data the control system already records for changes in the relationships between measured parameters, and can extend the calibration period of a sensor that has

not drifted [21]. Operating plants are moving from analog to digital control through individual system upgrades [21], and control systems that adopt office computers, operating systems and network protocols have far less isolation from the outside world and a greater need to be secured [24].

6.8 Limitations

The model is a control-structure teaching model and not a safety analysis code. The couplings, the signs, the gains and the time constants are in the right places and of the right size, and the lessons transfer directly, but there is no hot channel, no critical power ratio, no regional neutronics and no emergency cooling, so the application cannot say anything about fuel damage, about margin to dryout, or about what happens more than a few minutes after a scram. The decay ratio of (6.16) is a calibrated correlation rather than a solved channel, so the onset boundary is meaningful and the limit-cycle amplitude is illustrative, and the drift-flux correlation is least reliable exactly in the low-flow corner the application most wants to demonstrate. The turbine has no bypass valves and no shaft dynamics, which removes the system a real plant relies on to absorb a load rejection. Disturbances enter as commanded steps and scheduled events instead of as the slow wandering a plant does, and the exact repeatability that a frozen step and a frozen scan buy is itself an unreality. Where each individual assumption breaks down is set out in Section 6.3.8.

The pump drive and the feedwater flow have no rate limits. They follow their commands through the first-order lags τ_p and τ_{fw} alone, where a real drive and a real feedwater pump set would also be slew limited. The scenarios step these demands by up to 50 points at once (the runback), so a limit of a few percent per second would stretch each step into a ramp of many seconds and change the fast transients this chapter measures.

The MPC's model leaves out the density-wave oscillator, so the decay ratio and the exclusion region are not in its plan. It also leaves out the noise, the feedwater temperature, the admission capacity, the isolation and relief valves, the trips, the rods' future moves and the regulator's anti-windup.

> **OPEN CHALLENGE** **Oscillations the Core Average Can Hide**
>
> Uprated boiling water reactors may run at up to 120 percent of their original licensed power with core flow as low as 80 percent of rated. Simulations of a turbine trip without scram from there show the core going unstable once both recirculation pumps trip. They also show that stability cannot be judged from the core average power alone: as the oscillation moves from the core-wide mode to higher modes, the average may fall while single channels swing harder. Good data at these conditions are scarce, so the predicted cladding temperatures carry a large uncertainty [26]. Progress would be monitors and models that resolve the regional modes, which the point model of this chapter cannot represent, checked against such data.

6.9 Discussion

Three of the results in Section 6.6 generalize past this plant. The first is that a structural change beat a tuning change by a factor of 4.6 on the same disturbance with identical gains, and that no gain on the single-element loop reaches the three-element result, because the limiting quantity is a right-half-plane zero rather than a margin [3]. The second is that a scheduling law scales whatever gain it is handed. Turning the scheduling of (6.21) on with the challenge's excessive power gain in place multiplied the worst level excursion by a factor of 7.6, through a low-level scram, and turning it on with correct gains bought nothing (67.0 mm against 67.5 mm with the rod controller in service, 102.6 mm against 46.5 mm without it), because the process gain it was written to compensate does not vary the way the derivation assumed. A scheduling rule is a claim about a gain, and the gain is worth measuring before the rule is written. The third is that a computed quantity can carry more diagnostic value than a measured one. The collapsed level of (6.14) is what an operator needs during any feedwater transient and no plant instrument measures it, and in Figure 6.2 it is the only signal that never has the wrong sign.

The engineering consequence is that the level indication and the trips reading it are correct instruments giving misleading information at exactly the moment the information matters. A reactor tripped on low water level 0.835 s after the turbine valves drift shut, while the vessel is gaining 276 kg, is neither an instrument failure nor a controller failure. It follows from a differential-pressure cell measuring the mixture it straddles, and the remedy in service is a collapsed-level calculation on the operator display built on (6.13) and (6.14) from signals a real plant already has.

Methods beyond the implemented architectures fit the same equations. Seventeen states with hard actuator limits, a right-half-plane zero and a stability requirement written on the operating point are a nonlinear model predictive control problem, and posing it that way turns the exclusion region of (6.16) into a constraint instead of a procedure and coordinates the feedwater, recirculation and valve moves as one problem rather than four [27]. The linear predictive controller of Section 6.5.6 is the first step, and it leaves the decay ratio out of its plan. The broader operational context is that reactors are increasingly asked to follow variable renewable generation rather than to hold base load [28], which moves them toward the reduced-flow part of the map where the recirculation gain, the level loop's zero and the decay ratio are all at their least favorable.

6.10 Conclusion

This chapter studied a boiling water reactor as an interactive laboratory for coupled loops with sign-reversing dynamics. The simulator integrates seventeen states, including two-group point kinetics with Doppler and void feedback, a drift-flux boiling channel with a subcooled-boiling correction, jet-pump recirculation against a quadratic loop resistance, a dome pressure on a lumped energy balance, and a vessel inventory split exactly into liquid, sub-surface steam and dome steam so that the indicated and the collapsed level are separate outputs. Measured on the running model, recirculation flow moves 0.659 % of rated power per percent of core flow with 63 % of the change in 3.3 s and no rod motion, a 0.10 MPa pressure change is worth 43.4 pcm, the plant returns 30.6 pcm of negative feedback per percent of power, and 28 °C of lost feedwater heating raises power 9.45 % with nothing on the reactor touched. A +10 % feedwater step drives the indication −40.9 mm for 6.10 s while the inventory gains 193 kg. Three-element control holds a 95–70 % recirculation ramp to −63.0 mm and no time outside the band, against +287.4 mm and 221.6 s for single element at identical gains. On the graded challenge, correcting the six tuning constants and the feedforward gain reduces the worst level excursion from 410 mm to 47 mm, while enabling the gain scheduling with the original gains in place scrams the plant on low level and raises it to 3097 mm. A model predictive controller cuts the pressure step's level dip from 114.9 mm to 17.7 mm, at 2.5 times the integrated pressure error. A fuller model would add a two-node axial power shape so that the regional oscillation mode appears, a hot channel with a critical power ratio, turbine bypass valves and a shaft, and xenon dynamics. The boiling water reactor is an accessible interactive plant on which to learn how loops are paired when the plant is already controlling itself.

KEY POINTS

- The steam bubbles in the core slow the chain reaction, so moving the valves, the feedwater, the pumps or the rods changes the bubbles, and each loop pushes on the other three.
- A +10 % feedwater step drives the indication −40.9 mm while the inventory gains 193 kg.
- Three-element control holds the ramp to −63.0 mm, against +287.4 mm for single element at identical gains.
- A scheduling rule is a claim about a gain, and the gain is worth measuring before the rule is written.
- The exclusion region is entered by reducing flow.

Where else this applies. Chapter 4 carries the same three feedwater structures on a drum, and Volume 6, Chapter 2 makes the same pairing decision before any loop is tuned.

References

[1] K. J. Åström and R. D. Bell, "Drum-boiler dynamics," *Automatica*, vol. 36, no. 3, pp. 363–378, 2000.

[2] K. Iinoya and R. J. Altpeter, "Inverse response in process control," *Industrial & Engineering Chemistry*, vol. 54, no. 7, pp. 39–43, 1962.

[3] R. H. Middleton, "Trade-offs in linear control system design," *Automatica*, vol. 27, no. 2, pp. 281–292, 1991.

[4] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[5] J. March-Leuba and J. M. Rey, "Coupled thermohydraulic-neutronic instabilities in boiling water nuclear reactors: a review of the state of the art," *Nuclear Engineering and Design*, vol. 145, no. 1-2, pp. 97–111, 1993.

[6] C. Lin and L. M. Grossman, "Optimal control of a boiling water reactor in load-following via multilevel methods," *Nuclear Science and Engineering*, vol. 92, no. 4, pp. 531–544, 1986.

[7] J. A. Mowrey, P. R. Boylan, and S. I. Abdel-Khalik, "Modeling of boiling water reactor feedwater control algorithms," *Nuclear Technology*, vol. 124, no. 3, pp. 223–242, 1998.

[8] N. Zuber and J. A. Findlay, "Average volumetric concentration in two-phase flow systems," *Journal of Heat Transfer*, vol. 87, no. 4, pp. 453–468, 1965.

[9] P. Saha and N. Zuber, "Point of net vapor generation and vapor void fraction in subcooled boiling," in *Proceedings of the 5th International Heat Transfer Conference*. Begell House, 1974, pp. 175–179.

[10] W. C. Lipinski, "Control aspects of the experimental boiling water reactor power plant," in *1956 IRE National Convention*. IEEE, 1956, pp. 84–89.

[11] L. E. Boing, E. A. Wimunc, and G. A. Whittington, "Design-development and operation of the experimental boiling-water reactor (EBWR) facility, 1955–1967," Argonne National Laboratory, Argonne, IL, Tech. Rep. ANL-91/13, Nov. 1990.

[12] F. Araya, K. Yoshida, M. Hirano, and Y. Yabushita, "Analysis of a neutron flux oscillation event at LaSalle-2," *Nuclear Technology*, vol. 93, no. 1, pp. 82–91, 1991.

[13] B. R. Upadhyaya and M. Kitamura, "Stability monitoring of boiling water reactors by time series analysis of neutron noise," *Nuclear Science and Engineering*, vol. 77, no. 4, pp. 480–492, 1981.

[14] J. R. Larson, "Comparison of TRAC-BF1 calculations with the LaSalle 2 instability event," Idaho National Engineering Laboratory, Idaho Falls, ID, Tech. Rep. EGG-NRE-10803, May 1993.

[15] W. Wulff, H. S. Cheng, and A. N. Mallen, "Causes of instability at LaSalle and consequences from postulated scram failure," in *International Workshop on Boiling Water Reactor Stability*, Upton, NY, Oct. 1990, brookhaven National Laboratory report BNL-NUREG-45273. [Online]. Available: https://www.osti.gov/biblio/6397094

[16] N. E. Todreas and M. S. Kazimi, *Nuclear Systems Volume I: Thermal Hydraulic Fundamentals*, 3rd ed. Boca Raton, FL, USA: CRC Press, 2021.

[17] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[18] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[19] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

[20] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[21] International Atomic Energy Agency, *Core Knowledge on Instrumentation and Control Systems in Nuclear Power Plants*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2011, no. NP-T-3.12.

[22] 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.

[23] T. Quinn, J. Mauck, R. Bockhorst, and K. Thomas, "Digital actuator technology," Idaho National Laboratory, Idaho Falls, ID, Tech. Rep. INL/EXT-14-33132, Sep. 2014.

[24] 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.

[25] International Atomic Energy Agency, *Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2018, no. NP-T-3.23.

[26] L.-Y. Cheng, J.-S. Baek, A. Cuadra, A. Aronson, D. Diamond, and P. Yarsky, "BWR anticipated transients without scram leading to instability," in *2013 ANS Winter Meeting and Nuclear Technology Expo*, Washington, DC, Nov. 2013, brookhaven National Laboratory report BNL-101129-2013-CP. [Online]. Available: https://www.osti.gov/biblio/1091775

[27] J. D. Hedengren, R. A. Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[28] A. Ho, B. W. Billings, J. D. Hedengren, and K. M. Powell, "Flexible operation of nuclear hybrid energy systems for load following and water desalination," *Renewable Energy Focus*, vol. 51, p. 100641, 2024.

6.11 Exercises

6.1 Concept check · 1 min

Which assignment of the four handles does the coordinated architecture of this chapter use?

(a) Pressure on the turbine valves, level on feedwater, power on the pumps, rods behind the pumps

(b) Power on the turbine valves, level on feedwater, pressure on the pumps, rods on core power

(c) Pressure on the pumps, level on the turbine valves, power on feedwater, rods on the level

(d) Power on the rods, level on feedwater, pressure on the turbine valves, pumps held fixed

6.2 Concept check · 2 min

With the pressure regulator in service and the level loop out, the feedwater demand steps from 100 to 110 %. The indicated level falls 40.9 mm in the first 3.45 s. What is happening in the vessel?

(a) The transmitter lags the surface by its 1.2 s time constant and still shows the old level

(b) The colder inlet water shortens the boiling length and condenses steam below the surface while the inventory rises

(c) The colder feedwater lowers the dome pressure, so the cell, ranged at the design pressure, reads low

(d) The vessel loses water at first, because the valve opens before the extra feedwater reaches the vessel

6.3 Concept check · 1 min

The model carries the equation $C_e\,\mathrm{d}p/\mathrm{d}t = Q_{\mathrm{th}} + W_{\mathrm{fw}}h_{\mathrm{fw}} - W_{\mathrm{out}}\left(h_f + h_{fg}\right)$. What is it?

(a) The mass balance on the vessel inventory that sets the level

(b) The point-kinetics balance on the neutron density

(c) The vessel energy balance that sets the dome pressure

(d) The drift-flux correlation for the channel void fraction

6.4 Theory and calculation · 6 min

The rod bank has a total worth $W_{\mathrm{rod}} = 0.030$ on the S-curve of Equation (6.6), and one dollar is $\beta = 0.0056$. The bank starts 35 % inserted.

(a) Compute the differential worth at 35 % in pcm per percent of travel and in cents per percent. *Checkpoint:* 47.6 pcm and 8.51 ¢ per percent

(b) Compute the integral worth of a move from 35 to 42 % inserted, in pcm and in cents, and compare it with seven times the answer to (a). *Checkpoint:* 366 pcm, 65.4 ¢, against 333 pcm

(c) The plant returns 30.6 pcm of negative feedback per percent of power. Estimate the settled power change after the move of (b) and compare it with the 11.46 % the chapter measures. *Checkpoint:* about 12.0 % lower

6.5 Theory and calculation · 6 min

At the design point the dome is at $p_0 = 7.17\,\mathrm{MPa}$, the turbine valves are at $y_{T,0} = 0.85$ and pass $W_{s,0} = 1800.2\,\mathrm{kg/s}$, and the energy storage is $C_e = 1.40 \times 10^7\,\mathrm{kJ/MPa}$.

(a) Compute the nozzle coefficient K_t of Equation (6.12). *Checkpoint:* 295.4 kg/s per unit valve per MPa

(b) Compute the saturated steam enthalpy $h_f + h_{fg}$ at p_0 from the fits of Equation (6.1). *Checkpoint:* 2770.7 kJ/kg

(c) Compute the initial rate of change of dome pressure per percent of valve opening, with power and feedwater held, and the initial rate after the 4 % park of scenario 9. Compare the second with the measured peak of 17.7 kPa/s and say which term of the balance makes up the difference. *Checkpoint:* −4.19 kPa/s per percent, −16.8 kPa/s for 4 %

6.6 Theory and calculation · 6 min

The decay ratio of Equation (6.16) has $\mathrm{DR}_0 = 0.25$, $a_{\mathrm{DR}} = 1.20$ and $b_{\mathrm{DR}} = 0.77$.

(a) After the runback of scenario 10 the plant settles at 68.9 % power on 56.2 % flow. Compute the decay ratio. *Checkpoint:* 0.459

(b) Compute the damping ratio ζ of Equation (6.17) at that decay ratio, and the decay ratio at which ζ reaches zero. *Checkpoint:* $\zeta = 0.123$, zero at $\mathrm{DR} = 1$

(c) With both pumps stopped the plant settles at 50.4 % power on 34.6 % flow. Compute the decay ratio there and say which of the two terms of the exponent contributes more. *Checkpoint:* 0.716

6.7 App challenge · 6 min

Load scenario 2, “Shrink: the level goes the wrong way”, as shipped (default seed). Press Run at the default speed of 2 on the Vessel view, pause between 40 and 50 s and export the CSV, whose time column is in seconds.

(a) Find the minimum of `L_ind_m` after the feedwater step at 30 s, its time, and `L_col_m` in the same row. Give the levels in mm to within 0.5 mm. *Checkpoint:* −40.9 mm at 33.4 s, collapsed +11.8 mm

(b) Read `V_steam_sub_m3` in that row and before the step, and express the change as a fraction of the design 80.71 m^3. Give the change to within 0.05 m^3. *Checkpoint:* −1.68 m^3, 2.1 %

(c) The step adds 180.0 kg/s. From the mass balance alone, at $\rho_f = 736.5\,\mathrm{kg/m}^3$ over $A_s = 32\,\mathrm{m}^2$, what would the level have risen by at the time of (a)? *Checkpoint:* 7.64 mm/s, about 26 mm

6.8 App challenge · 8 min

Load scenario 5, “Single-element feedwater through a load change”, as shipped (default seed), and run all 400 s at speed 10. Read the Off level tile and export the CSV, whose time column is in minutes. Repeat with scenario 6, “Three-element feedwater, same load change”.

(a) For each run give the Off level tile, the largest deviation of `L_ind_m` from zero in mm to within 1 mm, and its time in minutes to within 0.02 min. *Checkpoint:* single: 222 s, +287.4 mm at 3.69 min. three: 0 s, −63.0 mm at 2.17 min

(b) Give the largest `L_col_m` of each run in mm and say what the difference between the two runs means for the water put into the vessel. *Checkpoint:* 428.8 mm against 166.4 mm

(c) The level gains are identical. Explain in two sentences why single element loses by a factor of about 4.6. *Checkpoint:* the error it acts on reads backwards for the first tens of seconds

6.9 App challenge · 8 min

Load the challenge, “Challenge: 70 to 100 % on recirculation flow”, at the default seed, and change one setting at a time from the case as given, each before Run. Run each attempt the full 600 s at speed 10 and read the verdict line (worst level, worst pressure, peak decay ratio).

(a) Set Rods behind the pumps to On (Advanced) and leave everything else as given. Give the verdict values and say which target fails. *Checkpoint:* 77 mm, 0.166 MPa, 0.45; the pressure band fails

(b) Reload the challenge, set only Gain scheduling to On, and run. Name the first trip and its approximate time from the marks on the trends. *Checkpoint:* Level 3 low-level scram near 104 s

(c) Explain from Equation (6.21) why switching the scheduling on was harmful here. *Checkpoint:* $4.5 \times 2.03 = 9.15$

6.10 Application area · 3 min

A plant engineer who has run fired drum boilers asks why a boiling water reactor puts its turbine valves on pressure and moves power with the recirculation pumps. Answer with the two measured coefficients the chapter gives.

6.11 Application area · 4 min

After a trip of both recirculation pumps the plant settles on natural circulation at about half power and a third of rated flow. (i) Why is this the dangerous corner of the operating map, and how is it left? (ii) Name two limitations of this chapter’s model that bound what it can say about the event.

6.12 Review the agent · 5 min

An AI assistant reviewing the challenge writes: “Jet-pump head goes as the square of pump drive, so at 46.7 % drive the gain from pump speed to power is less than half its value at 95 %. Switch Gain scheduling on first. It scales the power-loop gain by $f_{p,0}/f_p$ and restores the loop’s performance at low drive.” Find the flaw and say how you would check it on the app before trusting the rule.