INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 3

Pressurizer Simulator for Split-Range Control Across Asymmetric Actuators

Pages 53–80

Edited by John D. Hedengren

How to cite this chapter. "Pressurizer Simulator for Split-Range Control Across Asymmetric Actuators," 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. 3, pp. 53–80.

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

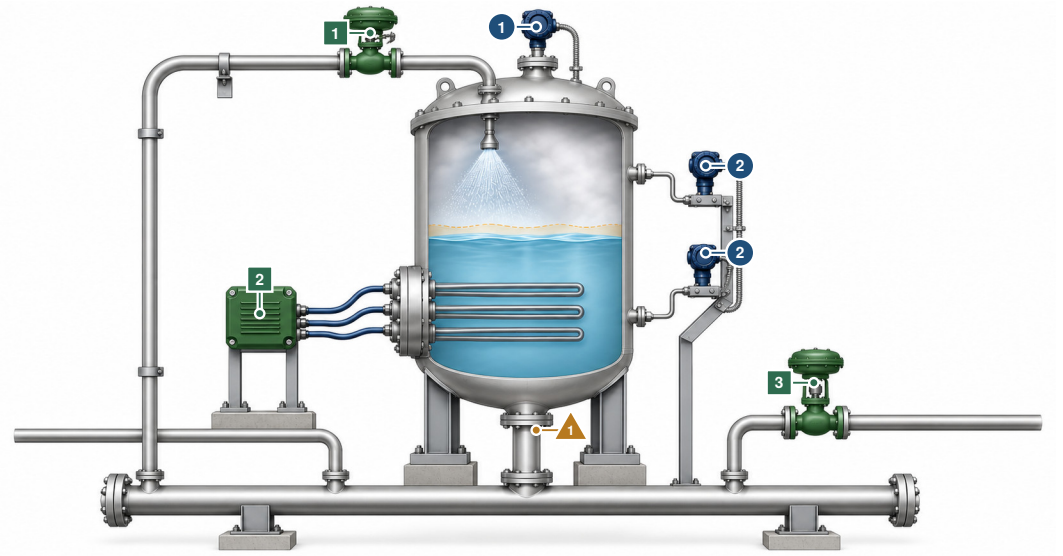
The app for this chapter
https://apmonitor.com/apps/pressurizer/

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

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

CHAPTER 3

Pressurizer Simulator for Split-Range Control Across Asymmetric Actuators



The controller moves 1 spray valve on the line from the cold leg
2 heater power, one modulated bank and four switched banks 3 charging valve into the coolant loop

The controller reads 1 pressure in the steam space 2 level, from taps in the steam and in the water

What disturbs it 1 a load change: coolant swelling up the surge line

The spray condenses steam at the top and the heaters in the water below make steam, two actuators that push the pressure opposite ways from one controller output. Illustration, not to scale.

IN PLAIN TERMS **One Controller for Heaters and a Spray**

The main idea. A nuclear reactor's cooling water is kept at high pressure by a tall tank with steam at the top. Heaters in the water make steam to raise the pressure, and a cold spray shrinks the steam to lower it.

The challenge. One controller runs both, and the spray works about nine times faster than the heaters. When the plant suddenly needs less steam, water swells into the tank faster than even a full spray can answer.

Why it matters. Knowing what a controller cannot fix shows where safety valves and simple safety rules must take over.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 3.2 and 3.3.8, the fact boxes of Sections 3.3 and 3.7, the try-it boxes of Section 3.7, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 3.7, Section 3.8, Section 3.9 and the open challenge

The control question Heaters make steam slowly and a spray condenses it about nine times faster, both on one controller output. How should the deadband, the valve trim and the loop structure be chosen, and which part of a load rejection can any controller win?

Learning objectives

- Compute the spray's condensing duty and the resting controller output from the model.
- Measure the two halves of a split range and explain their ratio.
- Choose a deadband and a valve trim from what fighting costs under noise.
- Derive the programmed level setpoint from the coolant expansion.
- Separate what tuning changes in a load rejection from what the plant fixes.

Prerequisites Sections 1.3, 1.4.5, 1.6.2 and 1.6.4, and Chapter 2. Section 1.6.3 for the predictive controller only.

The app https://apmonitor.com/apps/pressurizer/

See also Volume 4, Chapter 2 (a protective override), Volume 4, Chapter 3 (limiters on one valve through a selector) and Volume 4, Chapter 10 (three actuators of different speed and cost). Volume 11, Chapter 14 for a split range whose two ends cost money.

3.1 Introduction

Split-range control is usually drawn as one line. A controller output runs from zero to one hundred, one actuator is given the lower part and another the upper part, and a diagram of two triangles crossing in the middle is offered as the whole of the design. That picture makes the arrangement look like a wiring convenience. It is not. Two actuators on one output are two different plants sharing one controller, and the gain, the speed, the operating cost and the failure mode of each are properties of the equipment rather than of the split, so the single tuning has to be a compromise that the diagram gives no way to compute [1, 2]. The pressurizer of a pressurized water reactor is the sharpest available demonstration of that point. Its two actuators differ by a factor of 9.3 in authority, by a factor of five in response time, and by the direction in which each of them moves a second controlled variable, and every one of those numbers is measurable in the same experiment.

The class of problem is broader than the plant. A split range across asymmetric actuators with a deadband appears wherever a variable is pushed both ways by different equipment, such as heating and cooling on a jacketed reactor (Volume 8, Chapter 2), charge and discharge on a battery, acid and base on a neutralization tank, or fuel gas and a vent valve on a fuel header (Volume 11, Chapter 14). A second loop sharing the same vessel through an inventory coupling appears wherever a material balance and an intensive property are controlled on the same holdup. Both structures resist the standard multivariable treatment for the same reason, which is that the interaction lives in an equipment constraint rather than in a transfer function. The relative gain array identifies the pairing but says nothing about which actuator is cheap and which is scarce [3], and the industrial answers to the coupling are a programmed setpoint, a selector and an interlock rather than a decoupler [4]. Recent formalization of split range and of selector-based constraint switching has supplied the analysis those structures always lacked, and it treats them as control design rather than as plumbing [1, 2, 4].

Pressurizer models in the nuclear literature were built for accident analysis and for plant simulators rather than for controller design. The two-region treatment with mass and energy balances on a saturated liquid and a saturated vapor region dates from the 1960s, was extended with bubble rise and condensate drop terms to reproduce measured insurge and outsurge transients [5], and was carried into digital simulation with the spray, the heaters and the relief valves included [6]. Later work showed that thermal equilibrium between the regions is the assumption that fails first and that a non-equilibrium formulation is needed to follow the pressure through a fast surge [7], and comparisons of two-region and four-region models against plant data continue to be published [8]. However, none of that work is written for a student who has to choose a deadband, and a transfer function from heater power to pressure never discloses that the same output also commands a spray

valve nine times stronger. A block diagram of a split range shows the wiring and hides the equipment, which is exactly the information a course develops through worked examples rather than through the diagram itself [9, 10]. This chapter works with a browser pressurizer whose two actuators have to be measured one at a time before any tuning is possible, so that the asymmetry of a split range arrives as a measured number rather than as a caution. The simulator is a thirteen-state non-equilibrium model of the vessel, the reactor coolant system and the surge line, with a split-range mapping, an installed valve characteristic, a low selector, a heater interlock and an optional model predictive controller that plans the two actuators apart. Its scenarios measure the gain ratio, price the deadband, close the second loop and then meet a load rejection. The central result is that the peak of an insurge belongs to the plant rather than to the controller, which leaves the recovery, the level excursion and the equipment interlocks as the only discriminators between architectures.

Section 3.2 describes primary pressure control and its instruments. Section 3.3 gives every constant and every balance needed to rebuild the simulator, with its assumptions in Section 3.3.8. Section 3.4 measures the two halves of the split range and prices the deadband and the valve trim. Section 3.5 describes the three loops, the selector, the interlock and the predictive controller, and Section 3.6 closes the level loop on the same vessel. Section 3.7 follows a load rejection, an outsurge, two hidden failures and the graded challenge, and compares the split range with the predictive controller, and Section 3.8 sets the two loops in their plant hardware. Sections 3.9 to 3.11 state the limitations, discuss what generalizes and conclude, and the exercises follow.

3.2 Primary Pressure Control and Its Instruments

The pressurizer exists because water is nearly incompressible. A four-loop reactor coolant system holds roughly 306 $\mathrm{m^3}$ of liquid outside the pressurizer at about 310 °C, and a rigid circuit full of liquid has no way to accept the volume change that a temperature change forces on it. Attaching a partly steam-filled vessel to the hot leg solves both problems at once, because the bubble accepts volume and its saturation pressure sets the pressure of everything connected to it. Design pressure for the circuit modeled in this chapter is 15.51 MPa (2250 psia), at which saturation occurs at 344.95 °C, and the margin to departure from nucleate boiling in the core is demonstrated against that pressure rather than against a band.

Early plants held the pressure with heaters and a spray line under manual or simple on-off control, and the analysis that justified the arrangement came after the hardware. The first pressurizer models written for transient analysis treated the vessel as saturated liquid and saturated vapor in equilibrium, and adding bubble rise through the pool and condensate rain out of the steam space was what first reproduced the measured pressure during an insurge [5]. Digital simulation followed within a few years and put the spray, the proportional and backup heater banks and the relief valves into one model [6]. The equilibrium assumption survived until experiments showed that the pool and the bubble separate during a fast transient, and that forcing both onto the saturation line predicts the wrong pressure during an insurge and too rapid a recovery after an outsurge [7]. Non-equilibrium two-region formulations are now the standard control-oriented representation, and four-region models that resolve a subcooled layer are used where pool stratification matters [8, 11]. Spray effectiveness received its own treatment, because a droplet falling through a steam space does not reach saturation and the condensation it achieves depends on its size and fall height [12].

The control structure settled earlier than the modeling and has changed very little. Pressure is held by a single controller whose output is split, with the low half driving a modulating heater bank backed by contactor-switched banks and the high half driving spray valves supplied from the cold leg. Level is held by the chemical and volume control system on charging flow, against a setpoint programmed from reactor power rather than fixed, because the coolant inventory the pressurizer must absorb is a function of the average coolant temperature and therefore of load. Above the control band sit a power-operated relief valve at 16.10 MPa and spring-loaded code safety valves at 17.24 MPa, and reactor trips act at 16.55 and 13.10 MPa. They are the limiting and protection layers of Section 1.4.5.

The accident at Three Mile Island in 1979 began in this system, as the history segments of this chapter tell. The regulatory and design response that followed reshaped light water reactor safety practice well beyond the pressurizer itself [13], and it is the reason the distinction between what a control system can do and what a relief valve is for is worth teaching explicitly. Interest in the loop has returned with load following, because a plant asked to change power quickly puts the surge line to work several times a day rather than at startup [14, 15].

> ⧖ HISTORY **A Relief Valve That Did Not Close (1979)**
>
> On 28 March 1979 a pressurized water reactor at Three Mile Island was running at 97 percent power when its feedwater stopped. The coolant heated and expanded, the pressurizer level rose and the steam above it was compressed. Pressure reached 2255 psi, 100 psi above normal, and the power-operated relief valve opened as designed. It should have closed 13 seconds into the accident and instead stayed open for 2 hours and 22 minutes, while a control room light showed only that the signal to close it had been sent. As coolant escaped, the pressurizer level rose while the pressure fell. Operators trained to hold that level cut the high pressure injection from 1000 to under 100 gallons per minute, and much of the core was uncovered and severely damaged [16].

3.3 Process Model

Pressurizer PZR-101 is a 51 m^3 vessel of 3.575 m^2 cross-section connected to the hot leg of a four-loop reactor coolant system by a surge line. At full power it holds 30.6 m^3 of water (18 088 kg at 591 $\mathrm{kg/m}^3$) under a 20.4 m^3 steam bubble (2097 kg at 102.8 $\mathrm{kg/m}^3$), which is 60 % level. Heater bundle EH-101 sits in the bottom 17 % of the vessel and consists of a 300 kW modulating bank plus four 300 kW backup banks on contactors. Spray valves PCV-101 admit up to 28 kg/s of 292.75 °C cold-leg water into the steam space through a nozzle in the upper head, with a continuous bypass of 0.25 kg/s that keeps the line hot. Charging and letdown each pass 4.5 kg/s at the design condition. Pressure transmitter PT-101 and level transmitter LT-101 supply the two controlled variables. Every symbol used below, together with the instrument tags and the geometry the parameters refer to, is drawn in Figure 3.1.

> ✦ DID YOU KNOW? **A Pressurizer of the Same Size**
>
> The pressurizer of the Three Mile Island unit usually held 800 cubic feet of water under 700 cubic feet of steam [16], about 23 and 20 m^3. The vessel of this chapter holds 30.6 m^3 of water under a 20.4 m^3 bubble at full power.

Table 3.1 classifies the two controlled variables, the two manipulated variables, the disturbances and the thirteen integrated states. Working units are seconds, kilograms, cubic meters, megapascals, degrees Celsius and kilowatts throughout, and the pressure control band is 70 kPa either side of the setpoint.

3.3.1 Water and Steam Properties

Properties are least-squares polynomial fits to the IAPWS Industrial Formulation 1997 [17] evaluated on the saturation line at 12, 14, 16 and 18 MPa, with compressed-liquid points at 250 and 300 °C added for the liquid functions. The fitted quantities are $T_{sat}(P)$, $\rho_g^{sat}(P)$ and its inverse $P_{sat}(\rho_g)$, $h_{fg}(P)$, $u_g^{sat}(P)$, $h_g^{sat}(P)$, $\rho_f(T)$, $h_f(T)$ and $c_{p,f}(T)$. Residuals over the fitted range are 0.036 °C on T_{sat}, 0.34 % on saturated vapor density, 1.3 kJ/kg on latent heat, 0.010 MPa on $P_{sat}(\rho_g)$, 0.52 % on liquid density and 2.3 kJ/kg on liquid enthalpy. Arguments are clamped to 11–19 MPa and 200–366 °C, because the vapor-density fit turns over below about 11 MPa. At the design point the fits return $T_{sat} = 344.95\,^{\circ}\mathrm{C}$, $h_{fg} = 964.4\,\mathrm{kJ/kg}$ and $c_{p,f} = 7.697\,\mathrm{kJ/(kg\,K)}$.

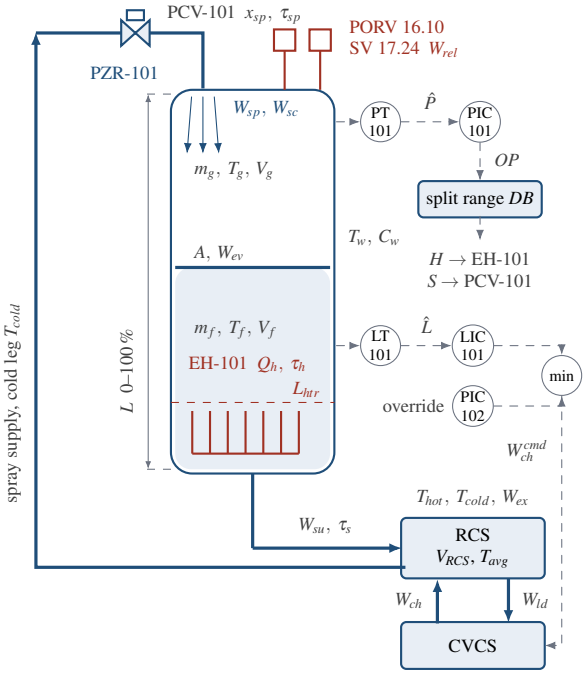


Figure 3.1 Pressurizer PZR-101 with the heater bundle, the spray line from the cold leg, the surge line to the reactor coolant system and the charging and letdown streams. Steam of mass m_g, temperature T_g and density ρ_g occupies V_g above water of mass m_f and temperature T_f occupying V_f, and the two exchange mass at W_{ev} across an interface of area A. Spray W_{sp} through PCV-101 condenses W_{sc} of steam directly, heaters EH-101 deliver Q_h into the pool and uncover below L_{htr}, and the surge line carries W_{su} set by charging, letdown, spray and the thermal expansion W_{ex} of the coolant at T_{avg}. PIC-101 splits one output OP into a heater demand H and a spray demand S across a deadband DB, and LIC-101 shares the charging valve with high-pressure override PIC-102 through a low selector.

3.3.2 Two Regions That Are Allowed to Disagree

The vessel carries four thermodynamic states rather than the usual two, because neither region is forced onto the saturation line. Mass conservation on the two regions is Equation (3.1).

$$\begin{aligned} \frac{\mathrm{d}m_g}{\mathrm{d}t} &= W_{ev} - W_{sc} - W_{rel} \\ \frac{\mathrm{d}m_f}{\mathrm{d}t} &= W_{su} + W_{sp} + W_{sc} - W_{ev} \end{aligned} \tag{3.1}$$

where m_g and m_f are the steam and liquid masses in kg, W_{ev} is the net interfacial evaporation, W_{sc} is the steam condensed by the spray, W_{rel} is the relief flow, W_{su} is the surge line flow positive into the vessel and W_{sp} is the spray flow, all in kg/s. The spray appears in both balances, once as liquid arriving and once as the steam it removes, which is the algebraic statement that a spray droplet is a mass transfer and a heat transfer at the same time. Pressure follows from the bubble alone, as shown in Equation (3.2).

$$P = P_{sat}(\rho_g) + z R \rho_g \left(T_g - T_{sat}(P_{sat}(\rho_g))\right) \tag{3.2}$$

where $\rho_g = m_g / V_g$ is the bubble density in kg/m^3, $R = 0.4615\,\mathrm{kJ/(kg\,K)}$ is the gas constant of steam and $z = 1.88$ is a real-gas multiplier on the ideal-gas $(\partial P/\partial T)_\rho$, because steam at 15 MPa departs strongly from ideal behavior. The second term is evaluated in kPa and scaled by 10^{-3} to match the MPa of the first. The second term in (3.2) is what an equilibrium model omits, and it is large.

Table 3.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
P (PT-101)	CV, measured	15.51 MPa	15.0–16.0 MPa SP
L (LT-101)	CV, measured	60 %	20–90 % SP
OP (PIC-101)	MV	40.75 %	0–100 %
W_{ch}	MV	4.5 kg/s	0–12 kg/s
W_{ld}	operator input	4.5 kg/s	0–16 kg/s
turbine load	DV, measured	100 %	0–110 %
T_{avg} bias	DV, unmeasured	0 K	−6 to 6 K
banks lost	DV, unmeasured	0	0–4
spray stuck open	DV, unmeasured	0 %	0–100 %
insulation	DV, unmeasured	1×	0.5–6×
m_g, T_g	states	2097, 344.5	kg, °C
m_f, T_f	states	18 088, 344.96	kg, °C
T_w	state	344.4 °C	40–380 °C
Q_h	state	141.8 kW	0–1500 kW
x_{sp}	state	0	0–1
T_{avg}	state	309.5 °C	—
P_{th}	state	100 %	0–120 %
W_{ch}, W_{ld} lags	states	4.5 kg/s	≥ 0
V_c	state	0 m^3	−8 to 8 m^3
W_{ex}	state	0 kg/s	±400 kg/s

Perturbing the bubble along the saturation line at the design density of 102.8 kg/m^3 gives a stiffness of 0.0936 MPa per kg/m^3, while compressing it adiabatically through the model's own energy balance gives 0.1949 MPa per kg/m^3, a factor of 2.08. An insurge therefore raises the pressure about twice as fast as a saturated-bubble model predicts, which is the mechanism the non-equilibrium pressurizer literature was written to capture [7, 8].

The interface transfers mass in whichever direction the pool departs from saturation, with two coefficients that differ by more than two orders of magnitude, as shown in Equation (3.3).

$$W_{ev} = -\frac{k\,(T_{sat} - T_f)}{h_{fg}}, \quad k = \begin{cases} k_c & T_f < T_{sat} \\ k_f & T_f > T_{sat} \end{cases} \tag{3.3}$$

where $k_c = 40\,\mathrm{kW/K}$ is the condensation conductance onto a subcooled pool and $k_f = 10000\,\mathrm{kW/K}$ is the bulk flashing conductance out of a superheated one. Condensation onto a quiescent pool is a surface process across 3.575 m^2, and with a pool heat capacity of 139 220 kJ/K its relaxation time $m_f c_{p,f}/k_c$ is 3480 s. Flashing is bulk nucleation through the whole inventory and relaxes in $m_f c_{p,f}/k_f = 13.9\,\mathrm{s}$. The 250:1 ratio is the reason a depressurization is self-limiting and a repressurization is not, and it is a property of the fluid rather than of the control system.

3.3.3 Spray, Heaters and the Vessel Wall

The spray condenses steam in proportion to the enthalpy deficit the droplets carry into the steam space, as shown in Equation (3.4).

$$W_{sc} = \eta_{sp}\, W_{sp}\, \frac{h_f(T_{sat}) - h_f(T_{cold})}{h_{fg}} \tag{3.4}$$

where $\eta_{sp} = 0.75$ is the fraction of that deficit actually used before the droplets reach the pool, which stands in for the droplet size distribution and fall history a spray model resolves explicitly [12]. At the design point $h_f(T_{sat}) = 1632.8\,\mathrm{kJ/kg}$ against $h_f(292.75\,^\circ\mathrm{C}) = 1299.4\,\mathrm{kJ/kg}$, a deficit of 333.5 kJ/kg

against a latent heat of 964.4 kJ/kg, so the full 28 kg/s of spray condenses 7.261 kg/s of steam. That is a condensing duty of 7.00 MW. The whole heater bank is 1.5 MW, and the first half of the asymmetry is already visible in those two numbers.

Spray flow follows the installed valve characteristic of Equation (3.5).

$$W_{sp} = W_{sp,\min} + \left(W_{sp,\max} - W_{sp,\min}\right) x_{sp}^{n} \qquad (3.5)$$
$$\tau_{sp} \frac{\mathrm{d}x_{sp}}{\mathrm{d}t} = \frac{S}{100} - x_{sp}$$

where x_{sp} is fractional valve travel, S is the spray demand in percent, n is the characteristic exponent with $n = 1$ linear and $n = 2.5$ approximating equal-percentage trim, and $\tau_{sp} = 4\,\mathrm{s}$ covers the valve stroke and the line transport. Travel is clipped to 0–1 after every step, and the derivative is zeroed at each bound so that a fixed-step integrator has no way to overshoot the stop and come back. Heater duty follows the demand through the bank logic of Equation (3.6).

$$Q_h^{cmd} = Q_p \min\left(\frac{H}{H_p}, 1\right) + Q_b \sum_{i=0}^{n_b-1} b_i \qquad (3.6)$$
$$\tau_h \frac{\mathrm{d}Q_h}{\mathrm{d}t} = Q_h^{cmd} - Q_h$$

where H is the heater demand in percent, $Q_p = 300\,\mathrm{kW}$ is the modulating bank rated full at $H_p = 20\,\%$, $Q_b = 300\,\mathrm{kW}$ is each of the $n_b = 4$ backup banks, b_i is the latched state of bank i and $\tau_h = 20\,\mathrm{s}$ is the sheath thermal lag. Bank i picks up at $H \geq 24 + 18i$ percent and drops out at $H < 24 + 18i - 6$, so the pick-up demands are 24, 42, 60 and 78 % and each bank carries 6 % of hysteresis against contactor chatter.

The vessel wall participates as a single lumped capacity, as shown in Equation (3.7).

$$C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} = hA_{wf}(T_f - T_w) + hA_{wg}(T_g - T_w) - \iota\, UA_{amb}(T_w - T_{amb}) \qquad (3.7)$$

where $C_w = 20000\,\mathrm{kJ/K}$ represents the inner few centimeters of the forging, hA_{wf} and hA_{wg} are the wall conductances to water and steam, ι is an insulation-condition multiplier and $UA_{amb} = 0.203\,\mathrm{kW/K}$ is the loss to containment. That loss is 59.8 kW at the design point.

3.3.4 Energy Balances on the Two Regions

The liquid energy balance collects the surge line, the spray, the condensate, the flashing, the heaters and the wall, as shown in Equation (3.8).

$$\begin{aligned} m_f c_{p,f} \frac{\mathrm{d}T_f}{\mathrm{d}t} = {} & \max(W_{su}, 0)\left[h_f(T_{hot}) - h_f\right] \\ & + W_{sp}\left[h_f(T_{cold}) - h_f\right] \\ & + W_{sc}(h_g - h_f) - W_{ev}\left[h_g^{sat} - h_f\right] \\ & + Q_h - hA_{wf}(T_f - T_w) \end{aligned} \qquad (3.8)$$

where $h_f = h_f(T_f)$ is the pool enthalpy and h_g is the enthalpy of the bubble including its superheat. An outsurge carries pool enthalpy away and therefore contributes nothing to (3.8), which is why only the positive part of W_{su} appears. The steam energy balance carries the compression work explicitly,

as shown in Equation (3.9).

$$\begin{aligned} Q_g &= W_{ev} h_g^{sat} - W_{sc} h_g - W_{rel} h_g - 1000\,P \frac{\mathrm{d}V_g}{\mathrm{d}t} \\ &\quad - hA_{wg}(T_g - T_w) - k_{ifg}(T_g - T_f) \\ c_{v,g} \frac{\mathrm{d}T_g}{\mathrm{d}t} &= \frac{Q_g - u_g \dot{m}_g}{m_g} - \frac{\mathrm{d}F}{\mathrm{d}\rho_g} \frac{\mathrm{d}\rho_g}{\mathrm{d}t} \end{aligned} \tag{3.9}$$

where $c_{v,g} = 3.30\,\mathrm{kJ/(kg\,K)}$, $k_{ifg} = 3.0\,\mathrm{kW/K}$ is a sensible conductance between steam and interface, and $u_g = F(\rho_g) + c_{v,g} T_g$ with $F(\rho) = u_g^{sat}(P_{sat}(\rho)) - c_{v,g} T_{sat}(P_{sat}(\rho))$ is the internal energy written so that $\dot{u}_g$ is explicit. The factor 1000 converts MPa m^3 to kJ. Both sides of (3.8) and (3.9) carry kW.

3.3.5 The Surge Line and the Reactor Coolant System

The reactor coolant system is rigid and liquid-full, so the mass it can hold is a function of its average temperature alone and everything it cannot hold arrives in the pressurizer. That closure is Equation (3.10).

$$\begin{aligned} W_{su} &= W_{ch} - W_{ld} - W_{sp} + W_{ex} \\ \tau_{loop} \frac{\mathrm{d}W_{ex}}{\mathrm{d}t} &= -V_{RCS} \frac{\mathrm{d}\rho_f}{\mathrm{d}T} \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} - W_{ex} \end{aligned} \tag{3.10}$$

where $V_{RCS} = 306\,\mathrm{m^3}$ and $\tau_{loop} = 8\,\mathrm{s}$ is the loop transport time. The lag matters. T_{avg} is one lumped number but the water it describes has to travel, so the expansion reaches the surge line spread over a loop transit rather than in the instant the turbine valve moves, and without it the whole insurge would land inside a single spray valve stroke. At 309.5 °C the coolant loses 2.601 kg/(m^3 K), so one kelvin of T_{avg} is 796 kg up the surge line. The primary loop itself is a single energy balance with a proportional rod controller on a programmed average temperature, the loop that Chapter 2 develops, as shown in Equation (3.11).

$$\begin{aligned} C \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} &= (P_{th} - \ell)\, q_{pct} - K_{sg}(T_{avg} - T_{prog}) \\ \tau_{rod} \frac{\mathrm{d}P_{th}}{\mathrm{d}t} &= \mathrm{clamp}[\ell + K_{rod}(T_{prog} - T_{avg}),\, 0,\, 115] - P_{th} \end{aligned} \tag{3.11}$$

where P_{th} is reactor thermal power in percent, ℓ is the turbine load in percent, $q_{pct} = 34\,110\,\mathrm{kW/\%}$, $C = 3.5 \times 10^6\,\mathrm{kJ/K}$, $K_{sg} = 90\,000\,\mathrm{kW/K}$ lumps steam generator self-regulation with the steam dump, $K_{rod} = 5\,\mathrm{\%/K}$, $\tau_{rod} = 25\,\mathrm{s}$ and $T_{prog} = 294 + 0.155\,\ell$ in °C is the programmed average temperature. Hot and cold leg temperatures are $T_{avg} \pm \frac{1}{2}\Delta T_{loop}\,\ell/100$ with $\Delta T_{loop} = 33.5\,\mathrm{K}$, which puts the cold leg at 292.75 °C at full load. The coolant is not perfectly rigid either, and its compressibility is carried as an equivalent volume held outside the vessel, as shown in Equation (3.12).

$$\tau_s \frac{\mathrm{d}V_c}{\mathrm{d}t} = \kappa V \left(P - P_{ref}\right) - V_c, \qquad V_f = \frac{m_f}{\rho_f(T_f)} - V_c \tag{3.12}$$

where $\kappa V = 0.40\,\mathrm{m^3/MPa}$ is 306 m^3 of 300 °C water at an isentropic compressibility of 1.3/GPa, $P_{ref} = 15.51\,\mathrm{MPa}$ and $\tau_s = 2.5\,\mathrm{s}$ is the surge line hydraulic response. Level is $L = 100\,V_f/V$, and $V_g = V - V_f$.

3.3.6 Relief, Protection and Instruments

The relief valves watch the true pressure rather than the transmitter, because they are mechanical. Two power-operated relief valves passing a combined 53 kg/s lift at 16.10 MPa and reseat at 15.86 MPa, and code safety valves passing 159 kg/s lift at 17.24 MPa and reseat at 16.90 MPa. Reactor trips

Table 3.2 Model Parameters. Every Value Is Adjustable in the Application. The Two Marked * Are Hidden From the User Interface.

Sym.	Meaning	Value	Unit	Sym.	Meaning	Value	Unit
V	vessel volume	51.0	m^3	Q_p	modulating bank	300	kW
A	interface area	3.575	m^2	Q_b	each backup bank	300	kW
L_{htr}	heater uncovery	17	%	n_b	backup banks	4	—
$V_{g,\min}$	smallest bubble*	0.5	m^3	H_p	demand for full Q_p	20	%
C_w	participating wall	20 000	kJ/K	τ_h	heater sheath lag	20	s
hA_{wf}	wall to water	100	kW/K	τ_{ch}, τ_{ld}	charging, letdown	8.0, 6.0	s
hA_{wg}	wall to steam	12	kW/K	V_{RCS}	coolant volume	306	m^3
UA_{amb}	loss to containment	0.203	kW/K	κV	coolant compliance	0.40	m^3/MPa
k_c	condensation	40	kW/K	τ_s	surge line response	2.5	s
k_f	flashing	10 000	kW/K	τ_{loop}	loop transport	8.0	s
k_{ifg}	steam to interface	3.0	kW/K	C	RCS heat capacity	3.5×10^6	kJ/K
T_{amb}	containment	50	°C	K_{sg}	SG self-regulation	90 000	kW/K
η_{sp}	spray efficiency	0.75	—	K_{rod}	rod control gain	5.0	%/K
$c_{v,g}$	steam c_v	3.30	kJ/(kg K)	τ_{rod}	rod and fuel	25	s
z	real-steam factor	1.88	—	ΔT_{loop}	hot–cold leg rise	33.5	K
$W_{sp,\max}$	spray, full travel	28.0	kg/s	P_{ref}	compliance ref.*	15.51	MPa
$W_{sp,\min}$	continuous bypass	0.25	kg/s	τ_{Pm}	pressure transmitter	1.0	s
τ_{sp}	spray valve and line	4.0	s	τ_{Lm}	level transmitter	4.0	s

are annunciated at 16.55 and 13.10 MPa and the run continues past them so that the transient stays readable. Both transmitters are first-order lags of 1.0 s on pressure and 4.0 s on level. Each is stepped by the discrete form $\hat{y}_k = \alpha \hat{y}_{k-1} + (1-\alpha) y_k$ with $\alpha = \exp(-\Delta t/\tau)$, which is the exact solution over a zero-order hold and is therefore stable at any period, and each accepts an optional colored-noise term scaled so that its stationary standard deviation equals the value set on the disturbance slider. A heater energized with the level below $L_{htr} = 17\,\%$ accumulates uncovered time, and 30 s of it destroys the bundle permanently. The interlock counts the heater command rather than the delivered duty, because a blocked heater is de-energized at once and what decays over the following two minutes is heat already stored in the sheath.

3.3.7 Parameters, Initialization and Numerics

Table 3.2 lists the constants that fix the behavior discussed in this chapter, and every scenario starts from a solved steady state rather than from a guess.

Initialization solves the steady state directly rather than iterating on the whole vessel. The requested level fixes V_f and V_g, the spray position fixes W_{sc} through (3.4), steady bubble mass requires $W_{ev} = W_{sc}$, and (3.3) then fixes the pool superheat exactly. Only the two-way coupling between the steam temperature and the wall needs a loop, and it converges in a handful of passes because both sit within a kelvin of the pool. The liquid balance (3.8) is finally solved for the heater duty rather than integrated, which is what makes the resting duty a computed quantity. At the design condition it returns 141.8 kW, of which 59.8 kW leaves through the insulation and the rest re-boils the 0.25 kg/s of continuous bypass spray. That duty is 9.45 % of heater demand and, through the split-range mapping of Section 3.4, a controller output of 40.75 %. The application verifies at startup that this state produces derivatives whose largest relative magnitude is below 10^{-9}.

The thirteen states are advanced by fourth-order Runge-Kutta on a fixed 0.05 s step, with the three controllers executing on a 0.5 s scan that matches a real pressurizer control system and the historian keeping ten records per simulated second. Every random draw comes from one generator seeded at 20260318, so a run with transmitter noise repeats exactly. Speed, seed and Reset behave as Section 1.6.1 describes.

3.3.8 Assumptions

Each simplification below is deliberate, and each fails somewhere.

One pool and one bubble. The whole vessel is two lumped regions with one interface. A real pressurizer stratifies, because the surge line discharges into the bottom of a pool that can be twenty kelvin colder at the top than at the nozzle, and computational studies of turbulent penetration during an in-surge resolve that structure [11]. The lumped treatment overstates how quickly the pool responds to a surge and produces no surge line thermal load at all.

The interface coefficients are judged, not correlated. Condensation at 40 kW/K and flashing at 10 000 kW/K are defensible orders of magnitude, and neither is a measurement. The 250:1 ratio between them is the physically important part and it is robust, while the individual values set how long the application takes to recover from a subcooling event.

Spray is one efficiency number. Equation (3.4) carries no droplet size distribution and no fall height, so the model cannot show spray becoming less effective as the bubble shrinks and the fall path shortens [12].

Properties are fits over a bounded range. Outside 11–19 MPa and 200–366 °C the arguments are clamped, so a depressurization that would take the loop two-phase cannot be followed.

The reactor coolant system is one node. A single average temperature, one lumped heat capacity and one self-regulation coefficient generate a realistic insurge on a load rejection and say nothing about steam generator level, secondary pressure or reactivity feedback beyond a proportional rod controller [14].

The relief valves are ideal. Fixed-capacity orifices with clean lift and reseat leave out the failure that matters most in practice, which is a relief valve that lifts and does not reseat [13].

Heater burnout is a timer. Thirty seconds uncovered and energized destroys the bundle. The real failure depends on sheath material, heat flux and how the water left.

Disturbances are idealized. A spray valve that sticks at exactly 45 % is not a failure a real valve has, but it isolates a continuous condensing load that the loop compensates perfectly. A real governor would round the corner off a load rejection, while a step gives a response that can be read directly from the trend.

> ⧖ HISTORY **Warnings Before the Accident (1977 to 1978)**
>
> The pattern had been seen before. In September 1977 a relief valve stuck open at a plant of the same design running at 9 percent power. Its pressurizer level rose while the pressure fell, operators interfered with the emergency injection, apparently relying on that rising level, and the block valve was closed about 20 minutes later. An engineer at another utility analyzed rising level with falling pressure more than a year before 1979, and relief valves had stuck open nine times at plants of that design [16]. At the accident unit itself a relief valve had stuck open a year earlier. The indicator light installed after that event showed that a closing signal had been sent, not whether the valve had closed [16].

3.4 The Split Range and Its Two Plants

3.4.1 The Mapping

One controller output is shared by two actuators across a deadband, as shown in Equation (3.13).

$$H = \frac{100\left[(50 - \frac{DB}{2}) - OP\right]}{50 - \frac{DB}{2}} \qquad S = \frac{100\left[OP - (50 + \frac{DB}{2})\right]}{50 - \frac{DB}{2}} \tag{3.13}$$

where OP is the controller output in percent, DB is the deadband width in percent of output and both H and S are clamped to 0–100 %, so between the two edges neither actuator is asked for anything.

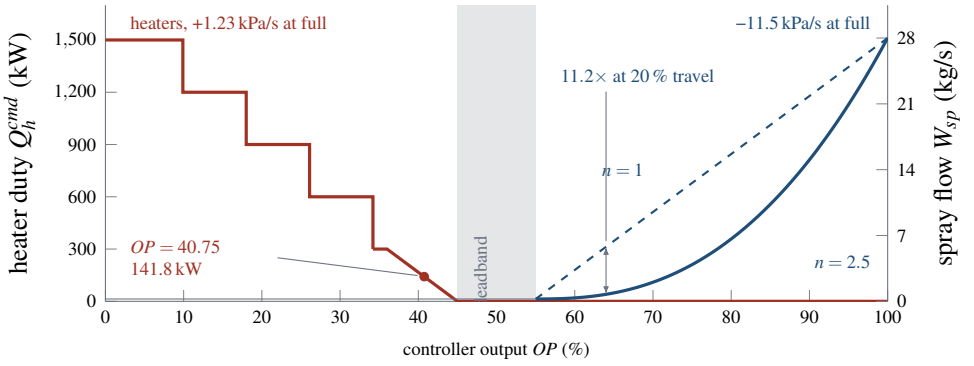


Figure 3.2 Split-range mapping from one controller output to two actuators of very different size. Heater duty is the staircase on the left axis, because the modulating bank saturates at 20 % demand ($OP = 36$) and the four backup banks pick up on contactors at 24, 42, 60 and 78 % demand. Spray flow is on the right axis, with the equal-percentage trim (solid curve) carrying 0.75 kg/s at 20 % travel against 5.80 kg/s for a linear valve (dashed line), and both trims reach the same 28 kg/s wide open. The design operating point sits at $OP = 40.75$, which is four points from the deadband edge rather than at the crossover.

At the default $DB = 10\,\%$ the heater band is $0 \le OP \le 45$ and the spray band is $55 \le OP \le 100$. Figure 3.2 draws the mapping through the actuators to the physical quantities they deliver.

3.4.2 Measuring the Two Halves

Scenario 1 places both actuators in manual at the design state and drives each one fully, and the measurement takes forty seconds. Driving the heater demand to 100 % in the first scan after Run raises the pressure at +1.2319 kPa/s averaged from $t = 10$ to 50 s. Driving the spray demand to 100 % from the same state lowers it at −11.5129 kPa/s. The ratio is **9.35**, from the two halves of one output. A demand set before Run acts half a second earlier, from $t = 0$, and gives a slightly smaller ratio over the same window.

Two mechanisms produce it, and they are separable. The duties differ, because 28 kg/s of water 52.2 K below saturation is the 7.00 MW condensing load of (3.4) against a heater bank of 1.5 MW. The paths differ as well. The spray removes steam directly from a bubble whose entire heat capacity is $m_g c_{v,g} = 6920\,\mathrm{kJ/K}$, while the heaters have to raise the pool above saturation before (3.3) produces any steam at all, and the pool and the participating wall carry 139 220 kJ/K and 20 000 kJ/K between them. Measured sixty seconds into a step to full heater demand, the delivered duty is 1431 kW and the flashing rate is 0.160 kg/s, which is 154 kW of steam production. Only 10.7 % of what the heaters put in is leaving as steam, and the rest is still going into sensible heat. A heater is a slow actuator because of the water in front of it and not because of its sheath.

The same asymmetry appears in the settling behavior, and it points the other way. Scenario 5 steps the pressure setpoint +100 kPa and then −200 kPa with the tuned loop in service. The upward move reaches a ±10 kPa band in 75.9 s with the heaters peaking at 1411 kW, overshoots the setpoint by 13.7 kPa and stays inside the band only from 207 s after the step. The downward move covers 154 kPa of its 200 kPa in the first forty seconds with the spray peaking at 17.5 kg/s, and then takes 126 s in total to reach the same band, because holding a lower pressure requires the pool to cool with the saturation temperature and eighteen tonnes of water do not cool quickly. Instantaneous authority and settling time are different questions, and a split range has to respect both.

3.4.3 Why the Answer Is Not a Smaller Gain

A single proportional gain has to serve a plant whose gain changes by 9.35 across the crossover. Reducing the gain until the spray half is comfortable gives up the heater half, which is already the

Table 3.3 Cost of the Deadband. Scenario 4 With 4 kPa of Transmitter Noise at the Shipped Seed, $K_c = 1500\,\%/\mathrm{MPa}$ and the Deadband Set Before the Run, Measured Over the 600 s Following the Setpoint Step at $t = 60\,\mathrm{s}$.

DB (%)	0	4	8	10	12	20	30
Fighting (s)	182.0	94.4	72.0	64.3	64.6	65.1	59.0
Heater energy (kWh)	47.01	43.03	41.64	41.34	41.11	40.67	39.83
Spray water (kg)	264	223	209	206	203	199	191

weak side. The installed valve characteristic solves the problem where it arises, at the actuator. With $n = 2.5$ in (3.5) the spray delivers 0.746 kg/s at 20 % travel against 5.800 kg/s for a linear valve, a factor of 11.2 less flow over the bypass at the position the controller normally occupies, and both valves deliver exactly 28.0 kg/s wide open where an insurge needs the capacity. The characteristic buys a lower gain in the region that matters without surrendering the capacity that a transient needs, which is the same argument that supports equal-percentage trim in any loop with a wide turndown [9]. The trim does not buy a faster setpoint response. With $n = 1$ the downward step of scenario 5 reaches its band sooner, in 85.6 s, because the linear valve has more gain near the seat. What the equal-percentage trim buys appears under noise. With 4 kPa of transmitter noise on the same 900 s run, the heaters and the spray are in service together for 234 s with $n = 1$ against 70 s with $n = 2.5$.

3.4.4 The Deadband Is Not Free

Scenario 4 runs the split range with the deadband at zero, a gain of 1500 %/MPa and 4 kPa of transmitter noise, and steps the setpoint +50 kPa at $t = 60\,\mathrm{s}$. The output has to pass through the crossover to move between actuators, and with noise on the measurement it does not pass through so much as loiter. Measured over 200 to 800 s, after the setpoint transient has cleared, the heaters and the spray are in service together for 139.5 s at $DB = 0$ and for 0.0 s at $DB = 10\,\%$. The noise is seeded, and over eight seeds the first figure ranges from 122 to 185 s while the second never exceeds 6.1 s. Over the whole 800 s run the plant consumes 57.61 kWh of heat and 340 kg of spray water at $DB = 0$ against 49.57 kWh and 256 kg at $DB = 10\,\%$, an increase of 16.2 % in heater energy and 33.0 % in spray water for no change in the pressure trace that an operator would notice.

Table 3.3 sweeps the width. The metrics improve steeply to about 8 % and then flatten, which makes the choice of ten points a design decision rather than a large number. Beyond 8 % every difference in the table is within the spread across noise seeds, which is 34 to 72 s of fighting over eight seeds between 8 and 30 %. The failure at the other end is not visible in this table, because a deadband wide enough to matter simply leaves the pressure sitting off setpoint until it drifts to an edge, and the run never asks it to.

The waste is the smaller part of the argument. Every spray cycle drops 292.75 °C water into a 344.95 °C pool and thermally cycles the nozzle and the surge line, and every heater cycle throws contactors rated for a finite number of operations. Neither is visible on a pressure trend, which is why the application counts the seconds directly.

3.5 Controller Implementation

3.5.1 The Shared Algorithm

The three loop controllers in the application are instances of the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement. In the notation of Equations (1.13) to (1.16), the running sum S_k of proportional and derivative increments is never clamped, v_k is the output the algorithm demands before the limit, and the integral I_k is the only term anti-windup touches, through the back-calculation tracking time T_t. The derivative filter on the pressure loop has $\tau_f = 4\,\mathrm{s}$. Manual operation and transfer are bumpless in both directions as Chapter 1 describes [18]. The scan period is $\Delta t_c = 0.5\,\mathrm{s}$, independent of the 0.05 s physics step.

3.5.2 The Three Loops

Loop action follows the convention of Section 1.3. PIC-101 is *direct* acting, so $a = -1$: a measurement above setpoint drives the output up into the spray half, and a rising output lowers the pressure. Its output limits are 0 and 100 % and its default tuning is $K_c = 800\,\%/\mathrm{MPa}$, $\tau_I = 240\,\mathrm{s}$ and $\tau_D = 12\,\mathrm{s}$. LIC-101 is *reverse* acting with $a = +1$, so a level below setpoint calls for more charging, and its output limits are 0 and 12 kg/s with $K_c = 1.0\,\mathrm{kg}/(\mathrm{s}\,\%)$ and $\tau_I = 400\,\mathrm{s}$. The high-pressure override PIC-102 is also reverse acting on the same valve, because charging water compresses the bubble, with $K_c = 90\,\mathrm{kg}/(\mathrm{s\,MPa})$, $\tau_I = 150\,\mathrm{s}$ and a setpoint of 15.72 MPa. PIC-101 and PIC-102 track at the shared default $T_t = \tau_I$, and LIC-101 uses $T_t = 30\,\mathrm{s}$. No output carries a rate limit, because the actuator lags of (3.5) and (3.6) already impose the physical rate. Pressure loop tuning is bounded to $0 \le K_c \le 2500$, $5 \le \tau_I \le 1200\,\mathrm{s}$, $0 \le \tau_D \le 120\,\mathrm{s}$ and $0 \le DB \le 40\,\%$.

The tracking time matters on this loop because the outsurge of scenario 7 takes the pressure a megapascal below setpoint, where the proportional and derivative sum S_k reaches −830 % against an output held at zero. Until 2026-10-08 the application tracked at $T_t = 30\,\mathrm{s}$ on PIC-101 and PIC-102. Back-calculation then drove I to +594 % so that v_k followed the clamp, and as the pressure recovered S_k shrank faster than I unwound, so v_k crossed the deadband with the pressure still 793 kPa low and the spray was commanded for 589 s while the pressure was more than 70 kPa below setpoint. That is integral windup produced by the anti-windup itself once the proportional term alone exceeds the output range eightfold. With $T_t = \tau_I$, I peaks at 71 % and the same condition lasts 12 s. The price is a slower unwind after a short saturation, visible as the 13.7 kPa overshoot of the upward step in scenario 5. The override suffered the same mechanism, and with $T_t = 30\,\mathrm{s}$ the selector handed it the charging valve at $t = 150.5\,\mathrm{s}$ in scenario 11 with the pressure 836 kPa below its setpoint.

Five architectures are selectable. MANUAL leaves all four stations to the operator. PRESSURE ONLY closes PIC-101 and leaves the level in manual. BOTH LOOPS adds LIC-101 against a fixed level setpoint with nothing between the two loops. BOTH PLUS OVERRIDES adds the three pieces of logic described in Section 3.6. MPC keeps that level side and replaces the split range with the predictive controller of Section 3.5.4.

3.5.3 The Selector and the Interlock

With the override architecture selected, LIC-101 and PIC-102 both run continuously and the charging valve takes whichever asks for less, as shown in Equation (3.14).

$$W_{ch}^{cmd} = \min[u_L,\, u_O]\,, \qquad I_{loser} \mathrel{+}= \frac{W_{ch}^{cmd} - u_{loser}}{T_{sel}}\Delta t_c \tag{3.14}$$

where u_L and u_O are the two unselected outputs and $T_{sel} = 20\,\mathrm{s}$. The loser is unwound toward the selected valve position by back-calculation and not by assigning its integral the selected output directly. The distinction is not cosmetic. Snapping the loser's integral forces it to reproduce a number the winner may itself be producing only because it is saturated, the two controllers then exchange roles on alternate scans, and the valve can sit shut while both loops call for it to open. That failure looks like a plant fault rather than a configuration error, and it is the classic way a selector scheme is got wrong [4]. Tracking the loser to the valve has its own consequence. The loser's output sits near the selected one, so its proportional term decides the hand-over, and a fast pressure rise passes the valve to PIC-102 before the pressure reaches its setpoint. With the override setpoint at its 16.20 MPa maximum, scenario 11 still hands it the valve for 23.5 s during the reload. This early proportional hand-over is retained.

The heater interlock is separate from both loops. When the measured level falls below $L_{htr} = 17\,\%$ in the override architecture, the heater demand is forced to zero regardless of what the pressure controller asks for. The loop that most wants to add heat when the pressure is low is exactly the loop that must not be allowed to when the level is low, and the two conditions arrive together during an outsurge.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.5.4 Model Predictive Control

The optional predictive controller plans the heater duty Q and the spray flow W as two inputs instead of one split output, within the general arrangement of Section 1.6.3, while charging stays with LIC-101, the programmed setpoint, the interlock and PIC-102. Every $T_s = 2\,\text{s}$, or four scans, it solves a quadratic program over $N = 90$ steps (180 s) and holds the first move until the next decision [19, 20]. Its model is the balance of Section 3.3 with the declared values of Table 3.2, not those of the run, linearized at 15.51 MPa, 60 % level, full load, charging and letdown at 4.5 kg/s and the balancing 141.82 kW (largest derivative 7.4×10^{-13}) and sampled with a zero-order hold, since per-input step-test fits cannot share the pool and bubble both inputs act through. It keeps the twelve vessel and transmitter states and omits the reactor coolant system, the contactors and the trim, which the planned kilowatts and kilograms per second are mapped through afterwards. The design point is on the flashing side of (3.3), so the model relaxes a subcooled pool through k_f. Over 300 s steps it follows a 150 kW heater step to 0.02 kPa root-mean-square and a 2.78 kg/s spray step to 1.4 % of the response, and one minute after a 3 kg/s charging cut it predicts −47.7 kPa against −53.5 kPa, but after a 5 kg/s rise 79.5 kPa against 300.8 kPa. A subcooled-side model predicts 295.8 and −177.5 kPa, so no single linearization serves both.

An unmeasured surge d_W (kg/s) entering as W_{ex} does and an unmeasured heat d_Q (kW) entering with the heaters augment the state as random walks [21]. A steady-state Kalman filter on the model sampled at Δt_c, with random-walk variances of $(2\,\text{kg/s})^2$ and $(20\,\text{kW})^2$ per scan against $(0.002\,\text{MPa})^2$ and $(0.1\,\%)^2$ on the transmitters, runs at every scan under every architecture and is told the requested duty and spray and the charging and letdown commands, so a lost bank or a stuck valve appears in d_Q or d_W. Over twelve move blocks of 2, 2, 2, 2, 4, 4, 6, 8, 12, 18, 30 and 90 s the plan minimizes Equation (3.15).

$$J = \sum_{k=1}^{N} \left(\frac{P_k - P_{\text{SP}}}{P_b} \right)^2 + \sum_{k=0}^{N-1} \left(c_h \frac{Q_k}{Q_{\max}} + c_s \frac{W_k}{W_{sp,\max}} \right) + \lambda \sum_{j=1}^{12} \left[\left(\frac{\Delta Q_j}{Q_{\max}} \right)^2 + \left(\frac{\Delta W_j}{W_{sp,\max}} \right)^2 \right] + J_s \tag{3.15}$$

where P_k is predicted with d_W, d_Q and the flow commands held, $P_b = 70\,\text{kPa}$, $Q_{\max} = 1500\,\text{kW}$, $c_h = 0.05$ and $c_s = 50$ are costs per step at full duty and full spray, ΔQ_j and ΔW_j are block moves measured first from the applied inputs, $\lambda = 10$, and $J_s = 1000\,(s_P/P_b)^2 + 10\,(s_L/5\,\%)^2$ prices the slacks of the soft bands. A step at full spray thus costs a thousand times one at full heat. The duty is hard within 0 to 1500 kW (zero while the heaters are blocked) and the spray within 0.25 to 28 kg/s, from the declared ratings, while the pressure is soft within 13.20 to 16.00 MPa, 100 kPa inside the trip and the PORV, and the level within 20 to 92 %, with limits read from the run. There is no terminal weight and no load preview. A dual active-set solver [22] solves the 24 decisions. The costs and λ are the largest within 10 % of the least IAE, 471.4 MPa s, among 90 settings on scenario 12 that spend no more spray and heat than the overrides. A plan that is not optimal, finite and in range is rejected, and PIC-101, which tracks the applied inputs through the inverse of (3.13), takes over without a bump until the next valid plan. The controller reads only the transmitters, setpoint, flow commands, load, interlock state and declared parameters, so its model shares the plant's equations but not its run, and a parameter the learner edits is a mismatch for the estimator.

3.6 Two Loops on One Vessel

3.6.1 The Interaction Is Not Symmetric

Scenario 6 closes both loops against fixed setpoints and asks for ten points more level. Filling the pressurizer *raises* the pressure, and the pressure controller answers with spray. Over the 440 s fill the

level reaches 69.55 %, the bubble is squeezed from 20.40 to 15.53 m^3, the heater duty falls below 5 kW within 70 s, the spray runs up to 6.91 kg/s, and 496 kg of the original 2097 kg of steam is condensed away. Returning the setpoint reverses it badly, because charging shuts but the only way out of the vessel is a letdown orifice at 4.5 kg/s, so the drain takes half as long again as the fill (654 s against 436 s to come within half a point) and the pressure falls to 15.2613 MPa with the heaters at their full 1500 kW. Moving a level ten points and putting it back costs 1969 kg of spray water and 196 kWh of heat from the step at $t = 60$ s to $t = 1000$ s.

Scenario 3 makes the opposite measurement. A step to 60 % spray demand delivers 7.99 kg/s into the steam space, condenses 1.95 kg/s of steam for 1.92 MW, and takes the pressure down at −3.14 kPa/s averaged over the 100 s step. It moves the level by −0.214 % over the same one hundred seconds. The spray draws its water from the cold leg, so the rigid liquid-full coolant system pulls the same mass back down the surge line and only the condensate stays. Charging is a level actuator with a large pressure side-effect, and spray is a pressure actuator the level can barely see. A relative gain array computed on those two gains identifies the pairing correctly [3], and it says nothing about the fact that both loops also share the surge line disturbance, which is where most of the difficulty actually lives.

Pause. Try Exercise 3.1 before reading on.

3.6.2 The Level Program and What It Does Not Cover

The industrial answer to the coupling is not a decoupler. It is a level setpoint programmed from turbine load, so that the target moves with the disturbance and the level controller has less to correct. The application programs 23 % at zero load to 60 % at full load, a slope of 0.370 % of level per percent of load. Converting the coolant expansion to pressurizer volume over the same span gives Equation (3.16).

$$\frac{\mathrm{d}L}{\mathrm{d}\mathrm{load}} = \frac{V_{RCS}\left[\rho_f(T_{avg,lo}) - \rho_f(T_{avg,hi})\right]}{\rho_{f,pz}\,V} \tag{3.16}$$

The coolant density is 738.38 kg/m^3 at 294 °C and 701.66 kg/m^3 at 309.5 °C, and the pressurizer water $\rho_{f,pz}$ is at 591.12 kg/m^3. With these values Equation (3.16) gives 11 238 kg, 19.01 m^3 and a slope of 0.3728 % per percent, which matches the programmed 0.370 % to 0.8 %. The setpoint is the plant's own thermal expansion written down rather than an operating preference.

The program is nonetheless not the whole level swing, and measuring it directly shows why. Settling the model at a series of loads with the pressure loop closed and charging matched to letdown, so that only the surge line moves the inventory, gives a natural level slope of 0.4931 % per percent between 100 and 95 % load. The excess over (3.16) is exact and has a name. Falling level grows the bubble at constant pressure, and the steam to fill it has to be boiled out of the pool. Between 100 and 95 % load the pool loses 743 kg, of which 614 kg is coolant expansion up the surge line and 129 kg becomes steam in the growing bubble, and the same decomposition holds to the kilogram at 70 % load with 4360 kg, 3602 kg and 758 kg. The consequence for the control system is measurable. At 70 % load the program asks for 48.90 % level while the process would settle at 45.54 % of its own accord, so the level controller still charges about 840 kg. The remaining roughly 176 kg of the pool-level gap is steam that condenses. That residue is small against the 11 % error a fixed setpoint carries, and it is not zero.

3.7 Results

The thirteen scenarios are ordered as a laboratory sequence. Scenarios 1 to 3 measure each actuator by hand and by a step test, scenarios 4 and 5 set the deadband and the valve trim, and scenario 6 closes the level loop. Scenario 7 is the load rejection, scenario 8 an outsurge, scenarios 9 and 10 two failures that hold every setpoint, scenario 11 adds the overrides, scenario 12 the predictive controller, and scenario 13 is the graded challenge. Each one loads an architecture, a tuning, a set of

disturbance positions and a schedule of timed events, and its briefing, completion rule and Reset work as Section 1.6.4 describes. Reset keeps the learner's own settings, which is how the comparisons in scenarios 1, 4, 5 and 11 are made. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt.

3.7.1 The Load Rejection

Scenario 7 is the transient the application exists for. The turbine sheds thirty percent of its load in a step at $t = 30$ s, and two transients follow in opposite directions from the single disturbance.

The reactor cannot follow instantly, so T_{avg} climbs 0.674 K while the program is falling, 306 m^3 of coolant expand, and the surge line peaks at 34.72 kg/s four seconds after the step. The pressure reaches 15.7462 MPa at $t = 39.6$ s, a rise of +236 kPa in ten seconds, and the spray demand goes to 100 %, although the valve reaches only 96 % of its travel by $t = 44.5$ s, after the peak. Even full spray could not stop it. Full spray condenses 7.26 kg/s while the surge is compressing the bubble several times faster than that, and the remaining 354 kPa of margin to the relief valve is the plant's rather than the controller's.

> **TRY IT IN THE APP** **7 · The Load Rejection: An Insurge You Cannot Stop**
>
> Load scenario 7 and press Run. Watch the pressure after the turbine sheds thirty percent of its load at 30 s. The spray demand goes to 100 %, yet the pressure still peaks at 15.7462 MPa at 39.6 s. Keep watching: the reversed surge then takes the pressure down to 14.5028 MPa while every heater bank comes in.

Then rod control and the steam dump win, the dump being the subject of Chapter 5. T_{avg} settles at the 70 % program of 304.85 °C, 4.65 K below where it started, the coolant contracts, and the surge reverses to −100.7 kg/s. The level falls from 60 to 45.37 % and the pressure to 14.5028 MPa, a full megapascal below program. Every heater bank comes in and the pressure is back inside its 70 kPa band at $t = 208$ s. LIC-101 meanwhile charges at its 12 kg/s maximum from $t = 78.5$ to 297.5 s and is still at 8.1 kg/s at 700 s, and that 292.75 °C water compresses the bubble. The pressure overshoots by 66.8 kPa at $t = 270$ s, the output crosses into the spray half at $t = 221$ s, and from $t = 328$ to 700 s the heaters stay below 5 kW while the spray runs at 1.5 to 4.6 kg/s, condensing what the charging compresses. Ten minutes after the rejection the pressure is 10 kPa high at 15.5202 MPa. When the unit is reloaded the sequence runs backwards and the level peaks at 70.48 %, ten and a half points above a setpoint that never moved, and the second outsurge takes the pressure to 14.50 MPa at $t = 911$ s. Half an hour into the run the loop is recovering on heat alone, as shown in Figure 3.4.

3.7.2 The Peak Is Not the Part Tuning Wins

Table 3.4 puts the three closed-loop architectures on the same disturbance at the same tuning, with the architecture chosen before the run. The peak pressure moves by 20 kPa across all three, which is 8.3 % of the 236 kPa rise, while the spray water differs by 1107 kg and the time with the level more than five points off setpoint differs by a factor of 26. Figure 3.3 draws Table 3.4 as bars.

Tuning moves the peak more than architecture does, and it moves it far less than intuition suggests. Sweeping the pressure loop over the plausible box $K_c \in \{400, 800, 1600, 2500\}$, $DB \in \{0, 10, 20\}$, $n \in \{1, 2.5\}$ and $\tau_D \in \{0, 12, 40\}$ with the level loop closed against a fixed setpoint gives a lowest peak of 15.6427 MPa and a highest of 15.8180 MPa, a spread of 175 kPa on peaks that are 133 to 308 kPa above program. The lowest peak comes from $K_c = 2500$ with no deadband, linear trim and $\tau_D = 40$ s, which is the one combination a designer would refuse, because that configuration then uses 11 529 % of output travel to $t = 1500$ s against 1290 % for the shipped tuning, spends 731 s of the run with both actuators in service against 247 s, and under the 4 kPa of transmitter noise of scenario 4 fights for all 600 s of the measurement window and moves the output 15 295 % against 4498 % for the shipped tuning, which never fights. Buying 104 kPa of peak with a factor of 8.9 in valve travel is not a trade a plant makes.

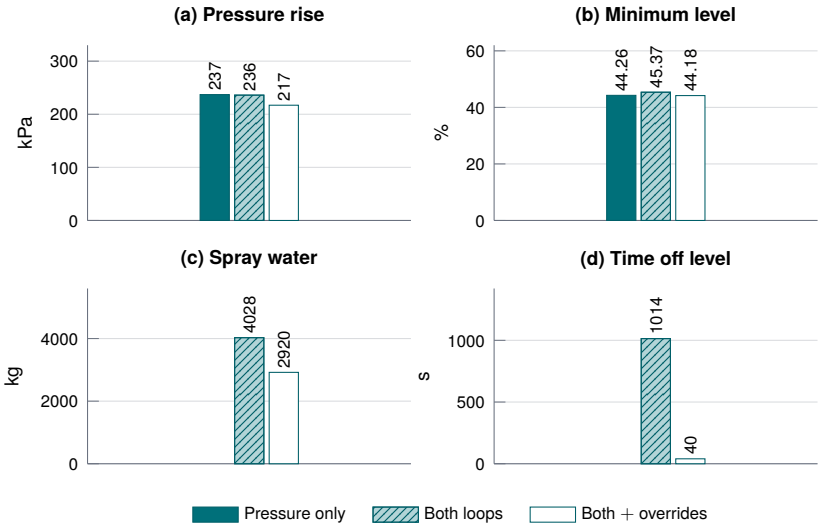


Figure 3.3 The −30 % load rejection of Table 3.4 under three architectures: (a) pressure rise, (b) minimum level, (c) spray water and (d) time with the level more than five points off setpoint. Solid: Pressure only. Hatched: Both loops. White: Both + overrides. Pressure only leaves the level in manual, so the table gives it no spray or level-time entry. Notice that the rise and the minimum level barely move, while the spray water and the time off level separate the architectures.

Table 3.4 The Same −30 % Load Rejection Under Three Architectures at the Shipped Tuning (K_c = 800 %/MPa, τ_I = 240 s, τ_D = 12 s, DB = 10 %, n = 2.5). Peaks Are Read Over the First 300 s and Totals at t = 1500 s.

	Pressure only	Both loops	Both + overrides
Peak P (MPa)	15.7466	15.7462	15.7269
Rise (kPa)	237	236	217
Minimum P (MPa)	14.4452	14.5028	14.4325
Minimum level (%)	44.26	45.37	44.18
Reload level peak (%)	—	70.48	63.04
P at 1500 s (MPa)	—	14.8501	15.2946
Spray water (kg)	—	4028	2920
Time off level (s)	—	1014	40

The same result appears in the graded challenge, where the disturbance is larger and the comparison is between two complete configurations rather than two gains. The challenge passes a run with no PORV lift at 16.10 MPa, the level always above 17 %, and the pressure within 100 kPa of 15.51 MPa over the last third of the run ($t \geq 1386$ s). Table 3.5 reports it. The shipped configuration and the suggested one peak within 26 kPa of each other, and neither lifts the relief valve. Everything after the peak separates them completely. The shipped configuration's worst error over the last third is twelve times the tolerance, although the level stays above 34.15 %. The suggested one is inside the tolerance by 73 kPa, with 75 kPa of PORV margin and the level above 33.2 %. No single change passes. Both loops with the overrides and K_c = 800 %/MPa, with nothing else changed, pass at 35 kPa. Tuning alone does not reach the target. The suggested split-range settings applied to the shipped two-loop architecture still end 320 kPa low, and 306 kPa low with the integral time halved to 120 s. The same settings with the override architecture end 1 kPa from program and pass even with the shipped level tuning, so the structure has to change before the gains do.

Table 3.5 Graded Challenge. A −55 % Load Rejection at $t = 60$ s, Read Over 2100 s, With the Last Third Judged From $t = 1386$ s as in the Application. The Shipped Configuration Has No Deadband, No Derivative, Linear Trim, a Gain Reduced to Stop Cycling and a Fixed Level Setpoint.

	Shipped	Suggested
Architecture	both loops	both + overrides
K_c (%/MPa), τ_I (s)	200, 500	800, 240
DB (%), n, τ_D (s)	0, 1, 0	10, 2.5, 12
Level K_c (kg/(s %)), τ_I (s)	2.5, 120	1.0, 400
Peak P (MPa)	16.0505	16.0253
PORV lifts	no	no
Minimum level (%)	34.15	33.20
Worst $\lvert P - P_{SP}\rvert$, last third (kPa)	1203	27
$\lvert P - P_{SP}\rvert$ at the end (kPa)	360	2
Spray water (kg)	2818	816
Heater energy (kWh)	519	524

✦ DID YOU KNOW? The Configuration That Passes Sprays Less

From the simulator: in the graded challenge the shipped configuration sprays 2818 kg of water and fails, while the suggested one sprays 816 kg and passes, for almost the same heater energy, 519 against 524 kWh (Table 3.5). The structure that holds the pressure also spends less water doing it.

The physical reason the peak resists tuning is in (3.2) and (3.4). The surge compresses the bubble at a rate set by the coolant expansion, and the only mechanism the controller has is condensation at a rate capped by the spray capacity and the cold leg enthalpy deficit. The one parameter that moves the peak at no cost is not a controller setting. Removing the coolant compressibility of (3.12) by setting κV to zero raises the rise from 236 kPa to 314 kPa, so the compliance of the reactor coolant itself takes 24.8 % off the peak. That is most of what the most aggressive tuning in the sweep buys, and it costs no valve travel at all.

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.7.3 Depressurization Is Self-Limiting

Scenario 8 opens a second letdown orifice with the level loop out of service, taking letdown from 4.5 to 14 kg/s against charging held at 4.5 kg/s. The level falls fifteen points from 60 to 45.04 % by the time the orifice shuts at $t = 430$ s, and the pressure falls only 305 kPa to 15.2050 MPa. The pool holds it up. As the bubble expands the saturation temperature falls, eighteen tonnes of water that were exactly at saturation become superheated, and (3.3) flashes them. Measured, the pool reaches 0.190 K above saturation at $t = 90$ s and flashes 1.95 kg/s, which is 1.90 MW and more than a quarter of what full spray can condense, and it is still flashing 1.84 kg/s when the orifice shuts. The heaters go to 1500 kW and stay there while the pressure falls anyway, which is the plainest available demonstration that heaters hold a steady state and do not answer transients.

3.7.4 Two Failures That Hold Every Setpoint

Scenario 9 removes two backup heater banks on a bus transfer and nothing happens. At 9.45 % demand the modulating bank alone is carrying the plant and no backup bank is in service, so the duty, the controller output, the pressure and every trend are exactly what they were. The failure becomes visible only when an outsurge asks for the heat, at which point the duty saturates at 900 kW instead of 1500 kW and the pressure sags 149 kPa to 15.3607 MPa. Scenario 10 sticks a spray valve 45 % open, which the equal-percentage trim makes 4.02 kg/s of cold water condensing 1.005 MW continuously. The loop holds the pressure to within 4 kPa of setpoint by running the heaters at 1500 kW with every

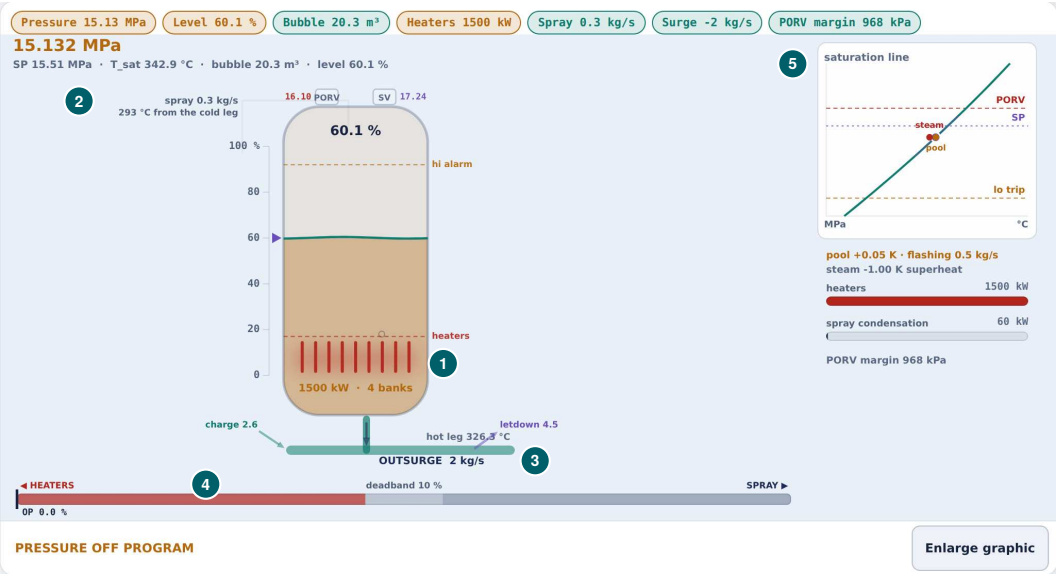


Figure 3.4 The load rejection of scenario 7 at $t = 1800$ s, during the heater-limited recovery from the outsurge that follows the reload. Pressure is at 15.132 MPa against a setpoint of 15.51 MPa and the banner reads PRESSURE OFF PROGRAM, the level is at 60.1 % against its fixed setpoint of 60.0 %, and the surge line is still running 2.1 kg/s outward. Every heater bank is in at 1500 kW while spray condensation supplies only 60 kW from the bypass, so the split-range bar under the vessel shows the output pinned at $OP = 0$ at the far end of the heater band with the whole deadband and the entire spray band unused. The saturation-line inset places the pool and the steam points below the setpoint line, and the readout under it reports 968 kPa of margin to the relief valve. The view is cut to the plant panel of the application. Callouts: (1) the heater bundle with every bank in, (2) the spray line from the cold leg, (3) the surge line running outward, (4) the split-range bar with the output pinned at zero, (5) the saturation-line inset.

bank in, around the clock. On the pressure trace everything is fine. What has happened is that the loop has spent all of its upward authority to stand still, and when a −15 % load change arrives at $t = 400$ s the pressure falls to 14.9779 MPa with nothing left to add. Two scorecard numbers would have caught it from the first minute, because the fighting counter reaches 369 s by $t = 400$ s and 869 s by $t = 900$ s, and the heater energy reaches 342 kWh over the run against the 35.5 kWh that 141.8 kW of standing duty would have consumed in the same 900 s. Neither is an alarm on a real plant and both are free to compute.

> **TRY IT IN THE APP** **10 · A Spray Valve That Will Not Shut**
>
> Load scenario 10 and press Run. After the spray valve sticks open at 30 s, the pressure stays within 4 kPa of setpoint. Watch the heater duty and the fighting counter instead: the heaters run at 1500 kW and the counter reaches 369 s by 400 s. Then the load change arrives and the pressure falls to 14.9779 MPa.

3.7.5 Comparison With the Predictive Controller

Table 3.6 runs scenario 7 (scenarios 7, 11 and 12 differ only in the architecture), scenario 8 and the challenge at its suggested settings with the same events, seed and instruments and the MPC as in Section 3.5.4.

Table 3.6 Split Range Against the Predictive Controller

Case	Structure	Peak P (MPa)	IAE (MPa s)	Off band (s)	Spray (kg)	Heat (kWh)	Fight (s)
Insurge[a]	Both loops	15.8653	589.8	891	4028	344.4	247
1500 s	Both + overrides	15.7394	441.1	949	2920	404.1	244
	MPC	15.8344	505.9	958	2835	371.0	133
Outsurge	Pressure only	15.5582	103.7	481	225	242.2	0
900 s	Both loops[b]	15.5681	26.7	90	1681	117.6	88
	MPC	15.5333	17.0	77	1670	111.8	77
Challenge	Both + overrides	16.0253	537.5	1077	816	523.8	24
2100 s	MPC	16.0609	535.5	1080	811	523.7	24

Totals and peaks are over the whole run. Off band is the time more than 70 kPa from setpoint, and Fight the time with more than 5 kW of heat and more than 0.375 kg/s of spray delivered together. [a]The peaks fall on the reload at 748 to 752 s, and the rejection's own peaks are 15.7462, 15.7269 and 15.7579 MPa. [b]Both + overrides is identical.

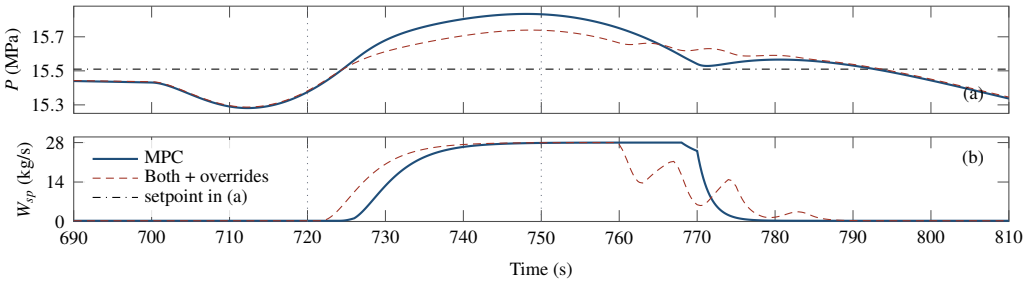


Figure 3.5 The reload of scenario 12 under the predictive controller and of scenario 11 under the split range with overrides, from the same initial state and the same load steps at 700 and 900 s. The dotted lines bound the decisive interval from 720 to 750 s. The predictive controller is the solid line, the split range with overrides the dashed line and the setpoint the dash-dotted line. (a) True pressure and setpoint. (b) Delivered spray flow.

On the outsurge, on the flashing side where the model was linearized, the MPC is ahead in every column. On the insurge it uses 2.9 % less spray and 8.2 % less heat than the overrides for a 14.7 % higher IAE. The loss is the reload's. The rejection peaks are 31 kPa apart near 39.5 s, the reload peaks 95 kPa apart at 748 s, as shown in Figure 3.5. The derivative of PIC-101 has the spray demand at 72 % at 722 s, with the pressure 74 kPa low, and full from 722.1 s, while the plans at 722 and 724 s keep the bypass, predicting 15.71 and 15.86 MPa at the horizon end, and the MPC's demand is full from 726.1 s. In diagnostic copies, not shipped, deciding every 0.5 s lowers the reload peak by 43 kPa, a subcooled-side model by 53 kPa, both by 96 kPa, and a zero spray cost by only 9 kPa. The spray cost sets the trade against pressure. With c_s at 0, 50 and 500, scenario 12 spends 3758, 2835 and 2457 kg of spray for an IAE of 265.3, 505.9 and 532.1 MPa s, while its reload peak stays within 18 kPa. Free spray buys a much lower IAE and leaves the reload peak nearly where it was, since the spray is already full on the reload under every price.

Neither structure ever commands heat and spray at once, with 0.0 s of the Fight condition on the commands in every case. Fight itself counts delivered overlap, and 116 of the MPC's 133 s and 225 of the split range's 244 s are heater duty decaying through the sheath lag after its command reached zero. The quadratic program adds a price on each actuator and pays it in pressure, and no setting in the sweep within the overrides' spray and heat reached their 441.1 MPa s.

On the challenge with its suggested settings the two are level on IAE, spray and heat and the MPC peaks 36 kPa higher. No plan was rejected in any scenario over 2100 s. The MPC solved 750, 450 and 1050 plans, and in Chrome a solve took a mean of 647 to 661 µs in three runs, at most 9.9 ms, about 0.03 % of one core, as wall clock on a 20-core desktop at a load average of 2 to 3.

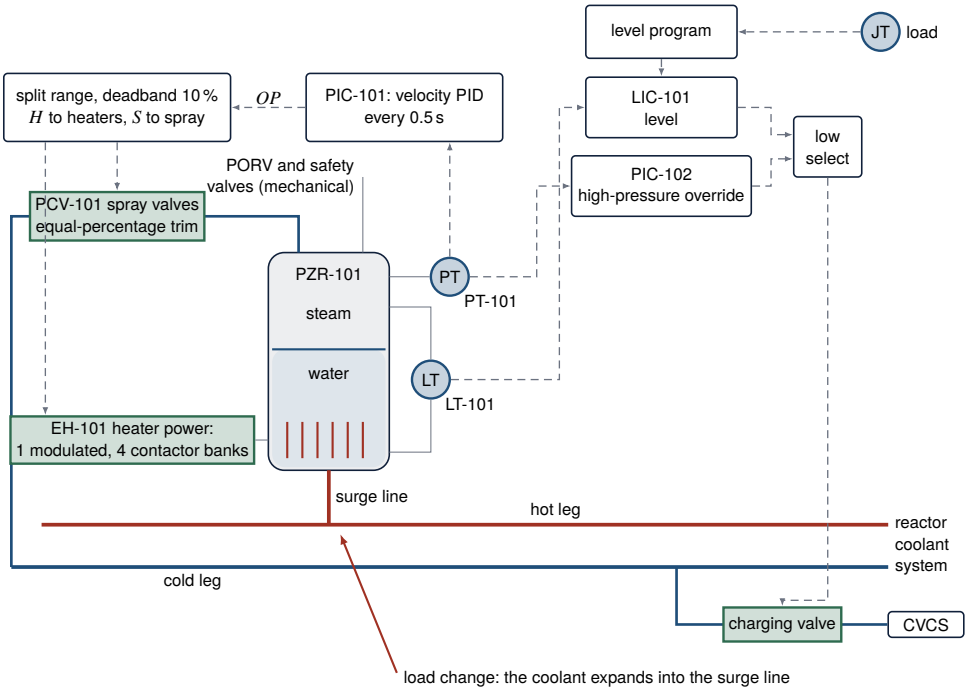


Figure 3.6 The pressure and level loops of Figure 3.1 with the hardware of a plant. A pressure transmitter (PT) on the steam space and a differential pressure level transmitter (LT) across the vessel feed the controllers, and a load measurement (JT) programs the level setpoint. PIC-101 runs every 0.5 s and its one output is split across the deadband between the heater power equipment and the spray valves, while LIC-101 and the override PIC-102 share the charging valve through the low selector. The relief valves act on their own. Thick solid lines carry hot-leg water (upper) and cold-leg water (lower and the spray line), and thin dashed arrows signals.

3.8 Technology in Practice

The tags of Figure 3.1 are, in a plant, a pressure transmitter and a level transmitter on the vessel, a heater bundle with its power equipment, spray valves on a line from the cold leg and the charging valve of the chemical and volume control system. Figure 3.6 draws the two loops with that hardware in place of the schematic's symbols.

What Measures It

The pressure of the whole primary circuit is read at the pressurizer, where the bubble sets it. Pressure transducers convert the force on a sensing element to a 4 to 20 mA signal or, in newer smart transmitters, to a digital one [23]. Level in a nuclear plant is most often measured as a differential pressure between taps at the top and the bottom of the vessel, against a filled reference leg [23]. The chapter carries both measurements as first-order lags, 1.0 s on PT-101 and 4.0 s on LT-101.

> **TECHNOLOGY** **Sensor: The Pressure Transmitter**
>
> A common pressure cell holds a taut diaphragm between two fixed plates, with a silicone fill fluid carrying the pressure in, so the deflection reads as a change of capacitance [24]. The transmitter sends a 4 to 20 mA signal, or a digital one from a smart transmitter, and the live zero at 4 mA reveals an open circuit [23]. Too much damping makes the transmitter report a slower process than the real one, and the instability that follows does not appear on the trend [24]. The chapter's PT-101 is a 1.0 s lag with optional noise.

What Moves It

Two final elements share the output of PIC-101. The heater bundle has one modulated bank and four banks switched by contactors, the two kinds of heater actuation found in plants, a regulated current and plain on and off [23]. The spray valves pass cold-leg water through a characterized trim, the equal-percentage form of Equation (3.5) with $n = 2.5$. Charging is the third element, shared by LIC-101 and PIC-102 through the low selector. The relief valves above them are mechanical and act on the true pressure, outside the control system, as the limiting and protection layers of Section 1.4.5.

> **TECHNOLOGY** **Final Element: Heaters and Spray on One Output**
>
> Valves that follow the output of one controller are called split-ranged. In the exclusive form both are shut at mid output and each opens toward its own end, so the fluid never passes through both at once [24]. An equal-percentage trim passes no more flow wide open than a linear one, but makes the installed flow closer to linear in stem position. A valve with stem friction moves in steps and ignores small reversals of its signal, and no tuning removes that [24]. Electric heaters are mostly switched on and off, though some are modulated by regulating their current [23]. PIC-101 drives both kinds of heater and the spray valve PCV-101.

Where the Controller Runs

Control loops of this kind run in field controllers with cycle times from milliseconds to minutes, under a supervisory layer that sends them setpoints [25]. The chapter's three controllers execute every 0.5 s above a 0.05 s physics step. In a function-block controller the split range is a block of its own: one PID block passes its output to a split-range block that drives two outputs, one for each final element [24]. Doing the split in the controller keeps each output a standard signal, and the deadband and the two bands of Equation (3.13) become settings of that block. The selector of Equation (3.14), the heater interlock and the level program of Equation (3.16) complete the chain.

> **IN PRACTICE** **Count the Thermal Cycles**
>
> A plant is designed for a set number of load cycles over its life. Under load following the pressurizer sees more thermal cycling, and the parts that suffer are the surge and spray lines, their nozzles and the vessel welds [26]. Plants that cycle therefore count the cycles on the surge line and monitor fatigue by measuring and recording wall temperatures against the design assumptions. Some change the pressurizer level and pressure control so that it makes fewer setpoint moves [26]. The fighting counter and the spray total on the chapter's scorecard are a first version of that count.

What Is Modernizing It

Three developments touch this loop. Smart transmitters diagnose themselves and report their own health [23], and on-line monitoring can show that a sensor has not drifted since its last calibration, so the calibration interval can be extended and the faults that maintenance itself introduces are avoided [23]. Continuous fatigue monitoring records wall temperatures and compares the observed load cycles with those the design assumed [26]. Plants that raise their power rating change the pressurizer pressure control to give finer control under reduced operating margins [23], which is where the deadband and trim choices of Section 3.4 return.

3.9 Limitations

The model is not a safety analysis and cannot support a claim about peak pressure, surge line loads or the licensing basis of any plant. Two lumped regions with judged interface coefficients cannot produce the stratified pool that a real in-surge creates, and the surge line thermal loads that follow are outside the model entirely [7, 11]. The spray is one efficiency number rather than a droplet field, so the loss of spray effectiveness as the bubble shrinks does not appear [12]. The reactor coolant system is a single node with one self-regulation coefficient standing in for both steam generator heat

transfer and the steam dump, which is enough to generate a realistic insurge and not enough to say anything about secondary pressure, steam generator level or reactivity feedback. The relief valves lift and reseat cleanly, so the failure that dominates the historical record is absent by design [13]. Charging, letdown and spray are first-order lags on commanded flow with no pump curves, no seal injection, no boron and no volume control tank, so the application can show what the continuous bypass spray costs in heater kilowatts but not why it exists. Reactor trips are annunciated rather than simulated. Even the repeatability is false, because a fixed integration step and a fixed scan give the same trajectory on every machine, and a plant control system is not that reproducible. What survived the simplification is what a controller has to be designed around, namely the 9.35 actuator asymmetry, the deadband, the installed characteristic, the inventory coupling, the selector and the interlock. Section 3.3.8 states where each assumption breaks down.

The predictive controller's model shares the plant's equations at the declared parameter values, so the comparison tests no structural model error, only the linearization side (79.5 kPa against 300.8 kPa one minute into an insurge) and the disturbance held over the horizon, and an edited plant parameter becomes a mismatch for the estimator. The held surge also puts the predicted pressure at the horizon end between 10.55 and 18.68 MPa on scenario 12 while the plant stays within 14.42 to 15.83 MPa.

OPEN CHALLENGE **Following Load Without Fatiguing the Surge Line**

Flexible operation aggravates cyclic loads, and a small change in operation can have an abrupt and large effect on fatigue. Thermal stratification in the surge line and actuations of the spray line are the examples given for the pressurizer [26]. These loads arise where hot and cold water meet in one pipe, which the chapter's two lumped regions cannot represent. Progress would be a pressure and level controller that knows the fatigue each spray cycle and each surge costs and trades it against pressure error, as the predictive controller of Section 3.5.4 already trades spray water against error.

3.10 Discussion

The result that generalizes beyond the pressurizer is the separation between what a control system can change and what it cannot. Table 3.4 shows a peak that moves 20 kPa across three architectures and a recovery that moves 445 kPa, so the peak belongs to the plant. Its value is set by the coolant expansion rate, by the bubble stiffness of 0.195 MPa per kg/m^3 under adiabatic compression, and by the 7.00 MW ceiling on spray condensation, and no arrangement of the same equipment beats those three numbers. The relief valves exist for that reason, and a control narrative promising to hold the pressure below the relief setpoint through every transient promises something the equipment does not support. The same reasoning applies wherever a fast disturbance outruns the strongest available actuator, which includes compressor surge (Volume 4, Chapter 2), batch reactor runaway and grid frequency after a large unit trip.

The three pieces of logic that do change the outcome are older and cheaper than multivariable control, and all three are auditable. Setpoint programming removes most of the interaction before any controller acts, because it makes the target a known function of the measured load. The low selector resolves a conflict over one valve with an explicit priority, and the interlock removes an actuator from the loop that would otherwise destroy it. Formal treatments of split range and of selector-based constraint switching now place these structures on the same footing as the optimization they approximate [1, 2, 4], which matters where a licensing basis is written around what the control system will do.

The model of Section 3.3 also supports methods beyond the implemented loops. The predictive controller of Section 3.5.4 plans the heaters and the spray with the relief setpoint as a soft constraint, and a third input on charging would replace the level loop and the selector with one three-input problem. State estimation on the same equations could infer the pool subcooling, which is not measured on a plant and which decides whether the next transient flashes or condenses. Coordinated

reactor and pressurizer control for load following is an active subject [14, 15], and the surge line closure of (3.10) couples the two problems.

3.11 Conclusion

This chapter has treated a pressurized water reactor pressurizer as an interactive laboratory for split-range control across actuators that are not the same size. The simulator integrates thirteen states, including separate steam and liquid temperatures so that the bubble superheats under compression at 0.195 MPa per kg/m^3 against 0.094 MPa per kg/m^3 on the saturation line, and so that the pool flashes at 1.95 kg/s on 0.190 K of superheat during an outsurge. Measured from the design state, full heaters move the pressure at +1.232 kPa/s and full spray at −11.513 kPa/s, a ratio of 9.35 that comes from a 7.00 MW condensing load against 1.5 MW of resistance heating and from a thermal path that holds 139 220 kJ/K of pool in front of the heaters and nothing in front of the spray. The deadband and the installed valve characteristic are the two design variables that make one gain serve both halves, because zero deadband costs 139.5 s of simultaneous operation in every ten minutes under 4 kPa of noise while ten points costs none, and an equal-percentage trim carries 11.2 times less flow above the bypass at 20 % travel with identical capacity wide open. On a −30 % load rejection the peak pressure moves only 20 kPa across three architectures at one tuning and 175 kPa across a wide tuning box, while the reactor coolant compressibility alone removes 24.8 % of it, so the peak is a plant property and the architectures separate on the recovery instead, at 15.2946 MPa against 14.8501 MPa after 1500 s and 2920 kg of spray water against 4028 kg. The anti-windup tracking time is a design variable on this loop, because a 30 s tracking time winds the pressure integral against a proportional term eight times the output range and sprays with the pressure still 793 kPa low, while tracking at the integral time keeps the spray shut until the pressure has recovered. Both pressure controllers retain this integral-time tracking. A predictive controller pricing the actuators spends 2.9 % less spray and 8.2 % less heat on the rejection and reload but peaks 95 kPa higher. Further work would add a stratified pool with a separate surge line node, a droplet-resolved spray model, a relief valve that can fail to reseat, and a steam generator secondary side so that the primary and secondary transients can be studied together. The lesson the app teaches by measurement is that two actuators on one output are two plants.

KEY POINTS

- Full spray moves the pressure 9.35 times faster than full heaters, and both hang on the two halves of one controller output.
- Zero deadband costs 139.5 s of simultaneous operation in ten minutes under 4 kPa of noise, while ten points cost none.
- On a −30 % load rejection the peak pressure moves only 20 kPa across three architectures, so the peak is a plant property.
- The architectures separate on the recovery, at 15.2946 MPa against 14.8501 MPa after 1500 s.
- Two actuators on one output are two plants.

Where else this applies. A split range across asymmetric actuators also sets heating against cooling on the jacketed reactor of Volume 8, Chapter 2 and fuel gas against a vent on the header of Volume 11, Chapter 14.

References

[1] A. Reyes-Lúa, C. Zotică, K. Forsman, and S. Skogestad, “Systematic design of split range controllers,” *IFAC-PapersOnLine*, vol. 52, no. 1, pp. 898–903, 2019.

[2] A. Reyes-Lúa and S. Skogestad, “Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy,” *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[3] 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.

[4] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[5] A. N. Nahavandi and S. Makkenchery, "An improved pressurizer model with bubble rise and condensate drop dynamics," *Nuclear Engineering and Design*, vol. 12, no. 2, pp. 135–147, 1970.

[6] R. C. Baron, "Digital model simulation of a nuclear pressurizer," *Nuclear Science and Engineering*, vol. 52, no. 3, pp. 283–291, 1973.

[7] S.-N. Kim and P. Griffith, "PWR pressurizer modeling," *Nuclear Engineering and Design*, vol. 102, no. 2, pp. 199–209, 1987.

[8] S. K. Moghanaki and M. Rahgoshay, "Simulation of two-region and four-region models for typical PWR pressurizer and benchmark obtained results using available results," *Annals of Nuclear Energy*, vol. 63, pp. 302–308, 2014.

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

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

[11] B. Tang and Y. Zhou, "Numerical investigation on turbulent penetration and thermal stratification for the in-surge case of the AP1000 pressurizer surge line," *Nuclear Engineering and Design*, vol. 378, p. 111176, 2021.

[12] J.-C. Jo, S.-K. Lee, and W.-K. Shin, "Heat transfer of a spray droplet in a PWR pressurizer," *KSME Journal*, vol. 5, no. 2, pp. 130–139, 1991.

[13] D. L. Phung, "Light water reactor safety before and after the three mile island accident," *Nuclear Science and Engineering*, vol. 90, no. 4, pp. 509–520, 1985.

[14] D. Bose, A. Hazra, S. Mukhopadhyay, and A. Gupta, "A co-ordinated control methodology for rapid load-following operation of a pressurized water reactor based small modular reactor," *Nuclear Engineering and Design*, vol. 367, p. 110748, 2020.

[15] H. K. Selim and N. M. El-Sahlamy, "Pressurizer system dynamic model for transient control in PWR," *Kerntechnik*, vol. 87, no. 6, pp. 672–682, 2022.

[16] President's Commission on the Accident at Three Mile Island, "Report of the President's commission on the accident at Three Mile Island. The need for change: The legacy of TMI," President's Commission on the Accident at Three Mile Island, Washington, DC, Tech. Rep. NP-24200, Oct. 1979.

[17] W. Wagner, J. R. Cooper, A. Dittmann, J. Kijima, H.-J. Kretzschmar, A. Kruse, R. Mareš, K. Oguchi, H. Sato, I. Stöcker, O. Šifner, Y. Takaishi, I. Tanishita, J. Trübenbach, and T. Willkommen, "The IAPWS industrial formulation 1997 for the thermodynamic properties of water and steam," *Journal of Engineering for Gas Turbines and Power*, vol. 122, no. 1, pp. 150–184, 2000.

[18] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *1989 American Control Conference*, 1989, pp. 1693–1698.

[19] 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.

[20] 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.

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

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

3.12 Exercises

3.1 Concept check · 2 min

The spray valve opens to pass 8 kg/s of cold-leg water into the steam space, with charging equal to letdown and the average coolant temperature steady. Use the mass balances of Equation (3.1) and the surge closure of Equation (3.10). What does the spray do to the total mass of water and steam in the pressurizer?

(a) It rises by 8 kg/s, because all of the spray water enters the vessel

(b) It rises by the condensation rate W_{sc}, because only condensed steam is new water

(c) It stays the same, because the surge line returns the spray and condensation only moves mass

(d) It falls by the condensation rate W_{sc}, because condensed steam leaves the bubble and the vessel

3.2 Concept check · 1 min

In the outsurge that follows a load rejection the pressure is far below its setpoint and the level is falling toward 17 %. With BOTH PLUS OVERRIDES, which element keeps the pressure controller from adding heat, and why?

(a) The low selector, because it gives the heaters to whichever loop asks for less heat

(b) The high-pressure override, because it takes the charging valve whenever the pressure is off setpoint

(c) The deadband, because an output pinned at zero lies outside both actuator bands

(d) The heater interlock, because a heater energized while uncovered is destroyed after 30 s

3.3 Concept check · 2 min

Which statement about the peak pressure of the 30 % load rejection in scenario 7 does the chapter's evidence support?

(a) It is fixed mostly by the plant, since full spray condenses far less steam than the surge pushes in

(b) It is set mostly by the pressure gain, since the spray demand saturates only when the gain is low

(c) It is set mostly by the architecture, since the overrides cut charging before the pressure peaks

(d) It is set mostly by the deadband, since a narrow band opens the spray valves much sooner

3.4 Theory and calculation · 6 min

At the design point the saturated liquid enthalpy is 1632.8 kJ/kg, cold-leg water at 292.75 °C carries 1299.4 kJ/kg, the latent heat is 964.4 kJ/kg, the spray efficiency is $\eta_{sp} = 0.75$ and the wall loses 59.8 kW to containment. The heater bank totals 1500 kW.

(a) Find the steam that full spray (28 kg/s) condenses, its condensing duty and the ratio of that duty to the full heater bank. *Checkpoint:* 7.26 kg/s, 7.00 MW and 4.67.

(b) A spray valve sticks at 45 % travel with the $n = 2.5$ trim. Find the spray flow and the steady heater duty that holds the pressure, if all spray water leaves down the surge line at the saturated liquid enthalpy. *Checkpoint:* 4.02 kg/s and 1400 kW, leaving about 100 kW in reserve.

(c) The measured ratio of the two pressure rates in scenario 1 is 9.35. Find the factor that the duty ratio of (a) leaves unexplained, and name the quantity in the model that produces it. *Checkpoint:* About 2.0.

3.5 Theory and calculation · 7 min

The programmed level of the app runs from 23 % at zero load to 60 % at full load. The coolant volume outside the pressurizer is 306 m^3, its density is 738.38 kg/m^3 at 294 °C and 701.66 kg/m^3 at 309.5 °C, the pressurizer water is at 591.12 kg/m^3 and steam at 102.8 kg/m^3, and the vessel holds 51 m^3.

(a) Evaluate Equation (3.16): the mass expelled from zero to full load, its volume in the pressurizer and the level slope. Compare the slope with the program's. *Checkpoint:* 11 236 kg, 19.01 m^3 and 0.373 % per percent against 0.370 %.

(b) The measured natural slope is 0.4931 % per percent. From 100 to 95 % load, find how much the bubble grows and the steam mass that has to be boiled out of the pool to fill it. *Checkpoint:* 1.257 m^3 and 129 kg.

(c) At 70 % load the program asks for 48.90 % while the plant settles at 45.54 % on its own. With the pressure held, find the pool water the gap represents and the water the level controller must charge to close it. *Checkpoint:* 1013 kg of pool water, of which about 840 kg is charged.

3.6 Theory and calculation · 6 min

The resting heater demand is $H = 9.4544\,\%$. Use the split-range mapping of Equation (3.13).

(a) Find the resting controller output with the deadband at 10 % and at zero. *Checkpoint:* 40.75 and 45.27 %.

(b) At $DB = 10\,\%$, find the outputs at which the first and the fourth backup banks pick up (heater demands of 24 and 78 %). *Checkpoint:* 34.2 and 9.9 %.

(c) Scenario 4 uses $K_c = 1500\,\%/\mathrm{MPa}$ and 4 kPa (one standard deviation) of transmitter noise. Find the output noise from the proportional term, and the distance from the resting output to the nearest edge of the spray band in units of that noise, at both deadbands. *Checkpoint:* 6 points, then 0.79 and 2.38 standard deviations.

3.7 App challenge · 7 min

Load scenario 1, Two actuators, one output, and they are not the same size (MANUAL), at the default seed. Set Heater demand to 100 % before Run, run to 50 s and export the CSV file. Press the reload button, set Spray demand to 100 % before Run, run to 50 s and export again. Both files have the time column Time (s).

(a) From the true pressure, column `p_MPa`, find each average rate between the rows at 10 and 50 s, and their ratio. Accept 0.005 kPa/s on the heater rate and 0.02 kPa/s on the spray rate. *Checkpoint:* +1.243 kPa/s, −11.450 kPa/s and 9.21.

(b) In the heater file, read `heater_kW` and `evap_kgs` at 50 s and find the share of the duty that leaves as steam, with a latent heat of 964.4 kJ/kg. *Checkpoint:* 1388.5 kW, 0.155 kg/s and about 10.7 %.

Explain why the ratio differs from the chapter's 9.35.

3.8 App challenge · 8 min

Load scenario 4, Split range with no deadband. It runs Pressure only at $K_c = 1500\,\%/\mathrm{MPa}$ with 4 kPa of transmitter noise, and steps the setpoint +50 kPa at 60 s. Use the default seed and speed 10, and run each case to 800 s. The exported file then has the time column Time (min), and the rows at 200 and 800 s read 3.33333 and 13.3333.

(a) Run the scenario as shipped. From column `metric_fighting_s`, find the seconds of fighting between 200 and 800 s, and read the Heater energy and Spray water tiles at 800 s. *Checkpoint:* 139.5 s, 57.6 kWh and 340 kg.

(b) Press Reset, set Split-range deadband to 10 % and run again. Read the same three values. *Checkpoint:* 0.0 s, 49.6 kWh and 256 kg.

Accept 0.5 s, 0.1 kWh and 1 kg. Find the increase in heater energy and in spray water that the zero deadband costs.

3.9 App challenge · 10 min

Load the challenge, Catch a fifty-five percent load rejection, at the default seed and speed 10. Each run covers the full 2100 s.

(a) Change only Control architecture to Both + overrides and run. Read the challenge result: the worst pressure error over the last third, whether the PORV lifted, and the lowest level. *Checkpoint:* 159 kPa, no lift, 33.2 %: not passed.

(b) Set Pressure loop gain to 800 %/MPa, press Reset and run again. Read the same three values and the Spray water tile. *Checkpoint:* 35 kPa, no lift, 33.2 %: passed, with 1000 kg of spray.

Accept 1 kPa on the errors. Compare both runs with the shipped configuration in Table 3.5, and say what each change did.

3.10 Application area · 4 min

A plant engineer prices the deadband from Table 3.3, which compares zero and ten points over 600 s at $K_c = 1500\,\%/\mathrm{MPa}$ with 4 kPa of noise.

(a) If the differences in heater energy and spray water between zero and ten points persisted around the clock, find the extra heater energy and the extra spray water in a year of 8760 h. *Checkpoint:* About 298 MW h and 3050 t.

(b) Name the consequence the chapter ranks above the energy, and say why no pressure trend shows it.

3.11 Application area · 4 min

A relief valve that lifts and does not reseat began the accident at Three Mile Island. The two PORVs pass 53 kg/s of steam, and the bubble holds 2097 kg at the design point.

(a) Find how long the PORVs would take to pass the bubble's whole mass if nothing replaced it. *Checkpoint:* About 40 s.

(b) Use Equation (3.3) with $k_f = 10\,000\,\mathrm{kW/K}$ and $h_{fg} = 964.4\,\mathrm{kJ/kg}$ to find the pool superheat at which flashing would replace the relief flow. *Checkpoint:* About 5.1 K.

(c) Name two limitations of the model that keep the app from following such an event to its end.

3.12 Review the agent · 5 min

An AI assistant reviews scenario 7 and writes: "The load rejection peaks at 15.7462 MPa, only 354 kPa below the PORV. The chapter's tuning sweep finds a lower peak of 15.6427 MPa with $K_c = 2500\,\%/\mathrm{MPa}$, no deadband, linear trim and $\tau_D = 40\,\mathrm{s}$. Adopt these settings as the standard tuning. The deadband and the equal-percentage trim only matter under transmitter noise, and a plant transmitter is quiet enough to ignore it." Find the flaw and say how you would check it.