INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 4

Three-Element Drum Level Simulator for an Inverse Response That Moves With Load

Pages 81–111

Edited by John D. Hedengren

How to cite this chapter. "Three-Element Drum Level Simulator for an Inverse Response That Moves With Load," 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. 4, pp. 81–111.

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

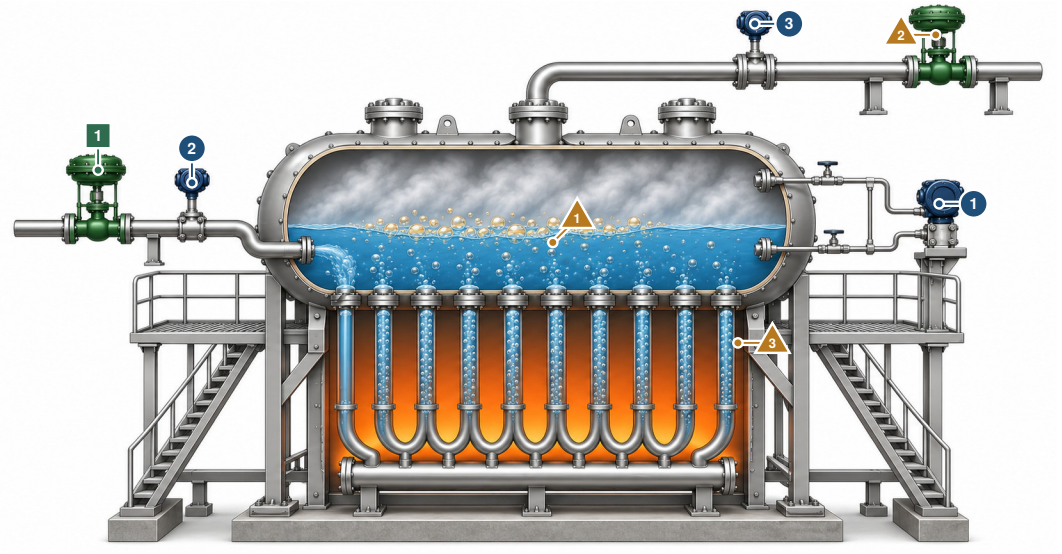
The app for this chapter
https://apmonitor.com/apps/steam-generator/

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

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

CHAPTER 4

Three-Element Drum Level Simulator for an Inverse Response That Moves With Load



The controller moves 1 feedwater valve on the inlet line

The controller reads 1 drum level, from a pressure difference between a tap in the steam and a tap in the water
2 feedwater flow 3 steam flow to the turbine

What disturbs it 1 steam bubbles under the surface, which swell and collapse
2 turbine governor, opening with the steam demand 3 heat from the firing under the tubes

Notice the bubbles below the water line: they hold up the surface, so a load change or a burst of cold feedwater first moves the measured level the wrong way. Illustration, not to scale.

IN PLAIN TERMS Keeping Water in a Boiling Drum

The main idea. A boiler drum must keep its water level in a narrow band while steam leaves and new water comes in. The gauge sees the water and also the steam bubbles hidden under the surface.

The challenge. Cold water added to the drum first collapses bubbles, so the level drops before it rises. A controller that watches only the level is fooled, and it is fooled for longer when the boiler runs at low output.

Why it matters. Too little water harms the boiler tubes and too much sends water to the turbine, so either limit stops the plant.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 4.1 and 4.2, the fact boxes of Sections 4.3 and 4.6, the try-it boxes of Sections 4.4 and 4.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	25 min	The first page, Section 4.6, Section 4.8, Section 4.9 and the open challenges

The control question A drum level is the water inventory plus the steam held below its surface, and the two terms answer a feedwater or load change in opposite directions. What does the wrong-way response that this creates cost a level loop as the load falls, and what does each added measurement of the single-element, two-element and three-element ladder buy against it?

Learning objectives

- Compute the steam that a kilogram of subcooled feedwater condenses, and the void contribution to the level at a given load, from the balances and the property fits.
- Distinguish the depth of the inverse response, which scales with the step, from its duration, which moves with load.
- Explain why a steam-flow feedforward deepens the dip of a load rejection while it removes the overshoot, and why it is exact only when stated in flow.
- Show on a pump trip why an inner feedwater flow loop covers faults that no level tuning reaches.
- Evaluate the level error a differential-pressure cell ranged at the wrong pressure leaves behind a loop that holds its indication on setpoint.

Prerequisites Sections 1.4.4 (shrink, swell and the moving zero), 1.6.2 (velocity form) and 1.6.4 (how an exercise names a run). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/steam-generator/

See also Volume 3, Chapter 2 for matching loop structure to where a disturbance enters, Volume 3, Chapter 3 for the inner-loop speed a cascade needs and Volume 3, Chapter 9 for a right-half-plane zero that moves with load. Volume 1, Chapter 6 for a vessel whose level swells while it loses water, and Volume 1, Chapter 2 for a level loop whose dynamics move with the operating point.

4.1 Introduction

The water level in a steam drum is among the most heavily instrumented measurements in a power plant and among the most frequent reasons a unit comes off line. Neither direction is tolerant. A level low enough to uncover the steam-generating tubes overheats them within seconds, and a level high enough to carry water over with the steam delivers liquid to a machine turning at 3000 rpm. Protection is therefore hard-wired on both sides, the band between the two trips is a fraction of a meter, and the loop that has to stay inside it reads a measurement that moves the wrong way for the first minute of every disturbance. Level control in the steam generator of a pressurized water reactor is documented as a source of unplanned shutdowns, and the difficulty concentrates at low power rather than at full power [1, 2]. The same loop on fossil-fired and heat-recovery boilers carries the same reputation in the practitioner literature [3, 4].

> ⧖ HISTORY **Emergency Feedwater That Could Not Reach the Steam Generators (1979)**
>
> At 36 seconds after four in the morning on 28 March 1979, the pumps feeding the two steam generators of Unit 2 at Three Mile Island tripped while the reactor ran at 97 percent power, and the turbine was shut down automatically. Three emergency feedwater pumps started, but a valve was closed on each of the two emergency lines. The two control room lights that showed it were missed, and one was covered by a maintenance tag. About a minute and 45 seconds into the event the steam generators boiled dry. Eight minutes in, the closed valves were found and opened, and water rushed in. The commission found that this loss of emergency feedwater had no significant effect on the outcome, but it added to the confusion [5].

What the transmitter reports is not the water inventory. It is the height of a water surface, and beneath that surface is a population of steam bubbles that occupies real volume. The indicated level is therefore the sum of two states, a water volume and a sub-surface steam volume, and the second responds to the feedwater valve with the opposite sign to the first. Feedwater arrives subcooled, absorbs sensible heat on its way to saturation, and takes that heat out of the bubbles already in the drum. Each bubble that collapses gives back about forty times its own mass in volume, so admitting

water lowers the indication before it raises it. The opposite disturbance is more dangerous. Opening the turbine governor takes steam away faster than the firing rate can replace the energy, the drum pressure falls, saturated steam density falls with it, and every bubble under the water line expands. The level rises while the drum loses mass, which is the transient in which a controller reading only level closes the feedwater valve on an emptying boiler.

A linearized model of that behavior carries a numerator root in the right half plane, and the classical consequences follow. The root cannot be canceled, it cannot be moved by tuning, and it caps the achievable closed-loop bandwidth at a value no gain can raise [6], as Section 1.4.4 shows for a two-term drum model. The process control literature has treated inverse response as a distinct design problem since the early 1960s [7], and tuning rules exist for the specific combination of an integrator, dead time and a right-half-plane zero [8]. The drum adds a property those treatments usually hold fixed. Natural circulation is driven by the density head the heated risers develop, so the circulating flow falls with the firing rate and steam spends longer below the water surface before it separates. The zero therefore migrates toward the origin as the boiler is turned down, so the achievable bandwidth is lowest at the load a cycled plant spends its nights at and highest at the load it was commissioned on.

That class is wider than boilers. Any plant whose measurement is a sum of two states that answer the manipulated variable with opposite signs and different speeds has the same shape, and vessels with a phase change in them are full of examples, including flash drums, evaporator bodies, reboiler sumps and refrigeration receivers. The flash vessel of Volume 1, Chapter 6 swells on a rise in steam demand while it empties, for the same reason. The broader class is any nonminimum-phase plant whose zero location is a function of the operating point rather than a constant of the hardware. The fixed version of the problem as it appears in a textbook does not require the loop to be retuned, gain scheduled or transferred to a different architecture as the plant moves, and this one does.

Simulators of drum boilers are usually built to answer a modeling question rather than a control-architecture question [9–11]. However, a transfer function whose numerator root was chosen by the author of a textbook problem presents inverse response as a property of algebra, and a student who observes that a smaller move produces a smaller dip draws the wrong conclusion from it, because the depth of the dip belongs to the step while the interval the response spends traveling backward belongs to the plant. That presentation also cannot show the outcome the practitioner literature reports, which is that most drum level events are decided by which measurements were wired in and by whether the transmitter is reading correctly, long before anybody chooses a gain [3].

This chapter works with an interactive natural-circulation steam generator in which shrink and swell emerge from a mass and energy balance rather than from a numerator coefficient. It measures what a load-dependent right-half-plane zero costs and what each rung of the single-element, two-element and three-element ladder buys. The drum model is a teaching-scale reduction of published work [11, 12], and three-element control has been standard practice for decades [4, 13]. On this plant the three architectures share a single actuator, a single sample period and a single set of limits, so the comparison between them is a matter of wiring and of nothing else. The differential-pressure cell's density error is computed from the physics rather than injected as a bias. The scenario sequence separates what the plant decides from what the engineer decides, and it ends on two faults no gain reaches.

The principles the plant carries are inverse response and the distinction between the depth of a wrong-way excursion and its duration, feedforward and cascade in the form in which they are actually wired, the time-scale separation a cascade requires, actuator limits and anti-windup, and the rule that a loop holding a biased measurement holds the bias [14]. The physical fundamentals are three conservation statements on one vessel, a mass balance on the water, a mass balance on the steam held below the surface and an energy balance that carries the pressure, coupled through saturated properties and separated in time by a circulation that is fast against the drum and slow against the valve. The industrial warrant is direct, because drum level is among the most frequent causes of unplanned outages on both fossil and nuclear units, and the fault modes that cause those outages are structural rather than arithmetic [1, 3].

Section 4.2 traces the architecture ladder and the drum models behind it. Section 4.3 states the balances, the property fits, the constants and the integration scheme, so that the plant can be rebuilt from this chapter. Section 4.4 measures the inverse response in open loop and its dependence on load. Section 4.5 describes the three architectures, the PID algorithm they share, the limits of a feedforward delivered as a valve position and the predictive controller. Section 4.6 compares the architectures on load changes, a pump trip, a swell, a miscalibrated transmitter, a noisy level and the graded challenge, and sets the predictive controller against three-element control. Sections 4.7 and 4.9 discuss what the results mean for plant practice and where the model stops, Section 4.8 sets the loop in its plant hardware, and Section 4.10 sums up. The exercises follow in Section 4.11.

4.2 Background

Feedwater regulation is older than feedback control as a discipline. The first devices were mechanical, a float in the drum or a thermostatic expansion tube across the water line, acting directly on the feed valve. Both are proportional controllers whose gain is set by geometry and whose offset nobody measures, and both respond to the indicated level, which means both respond to shrink and swell in the wrong direction. Practitioner accounts of boiler control describe the sequence from those regulators through single-element, two-element and three-element schemes as the industry's response to exactly that defect [3, 4].

The quantitative case arrived with the first dynamic analyses of boilers. A 1958 study derived the drum water level response from mass and energy balances and identified the sub-surface steam volume as the cause of the wrong-way motion [9], and the effect of rapid steam take-off on natural circulation and drum water level was studied in its own right once turbines were expected to pick up load quickly [15]. A reduced nonlinear model of a drum boiler and turbine unit intended for control design appeared in 1972 [10] and reached its standard form as a fourth-order model whose drum, downcomer and riser balances reproduce the measured level response including the sub-surface steam volume [11]. Later work treated the void fraction distribution and the circulation loop as the mechanism that sets both the size and the duration of the excursion [12], and reviews written for the control audience place shrink and swell alongside pressure and temperature as the three couplings a unit controller has to respect [16].

The architecture settled before the modeling did. Single-element control uses drum level alone. Two-element control adds a steam flow measurement as a feedforward, on the argument that a kilogram leaving has to be replaced by a kilogram entering and that the steam flow signal moves the instant the governor does. Three-element control adds a feedwater flow measurement and an inner flow controller, so that the level controller writes a flow setpoint rather than a valve position. By the time the standard process control references were written the three-element scheme was described as the normal arrangement for any drum that swings load, with single-element control retained for low loads at which the feedwater flow measurement stops being trustworthy [4, 13]. The nuclear industry reached the same structure by a different route and has continued to publish on it, because a U-tube steam generator has stronger shrink and swell than a fossil drum and its worst case is at low power where the plant is least willing to accept an excursion [1, 2]. A boiling water reactor carries the same three-element structure on its vessel level, where the voids are also the moderator (Chapter 6).

Two developments have kept the loop current, and neither is a tuning problem. The first is instrumentation. A differential-pressure cell infers a height by dividing a pressure difference by a density difference, and the densities it uses are the ones it was ranged with rather than the ones in the drum, so its error moves with operating pressure and is shared by the trips that read the same transmitter [3, 17]. The second is duty. Thermal plants commissioned as base-load machines are now cycled to accommodate variable renewable generation, which moves them into the low-load, fast-ramp region where the level loop is hardest and where the architecture chosen at design time may no longer be the right one [18].

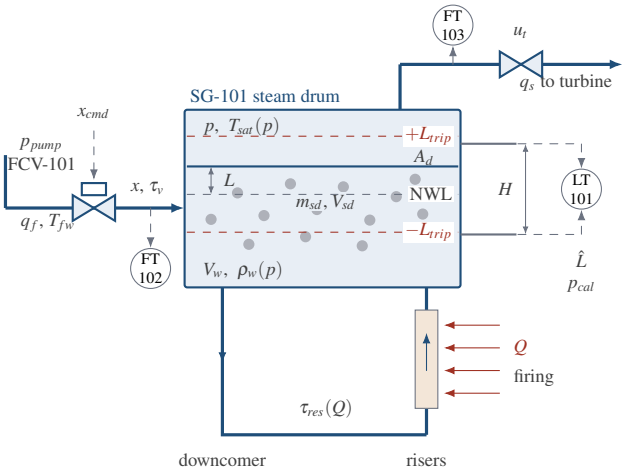


Figure 4.1 Natural-circulation steam generator SG-101 with its feedwater valve, turbine governor and drum level measurement. Water volume V_w occupies the drum below a surface of area A_d, and steam mass m_{sd} occupying volume V_{sd} is carried below that surface by a circulation whose residence time τ_{res} falls with the firing rate Q. Feedwater q_f at T_{fw} enters through FCV-101 at stem position x against pump pressure p_{pump}, and steam q_s leaves through the governor at opening u_t. LT-101 is a differential-pressure cell across taps H apart, ranged at pressure p_{cal} and reporting $\hat{L}$. Both level trips act on $\hat{L}$.

> **HISTORY** **From Three Elements to n Elements (1979)**
>
> A paper prepared in 1979 for a nuclear society meeting observed that plants were still being shut down by loss of steam generator water level control. It named the difficulties: a process with several interacting variables, a level that answers a steam flow change the wrong way, a delayed answer to feed flow, flows sensitive to valve position and uncertain measurements. Many were worse at low power, so the control schemes of the day ran automatically only above 15 percent power and left the range below to manual control. The paper proposed state feedback from a model of the steam generator, a natural extension of the classical three-element controller to n elements, one for each state of the model [19].

4.3 Steam Generator Model

Steam generator SG-101 is a natural-circulation drum boiler rated at 25 kg/s of saturated steam at 40 bar, which is 250.3 °C of saturation temperature and 53.0 MW of absorbed heat. Water leaves the drum down an unheated downcomer, boils in heated risers and returns to the drum as a two-phase mixture. The steam separates and leaves through the turbine governor. Feedwater at 160 °C replaces it through equal-percentage valve FCV-101, which sits 70 % open at the design condition against a feed pump holding 55 bar. Level transmitter LT-101 is a differential-pressure cell straddling two drum taps 1.2 m apart and reports a 600 mm span with normal water level at its center. The arrangement, the instrument tags and each symbol used below appear in Figure 4.1. Both level trips act at ±250 mm and the deviation alarms at ±150 mm, so the whole usable band is half a meter.

Table 4.1 sorts the signals into controlled, manipulated, disturbance and state. Five states are integrated. The working units are the ones used in boiler practice, mass in kg, flow in kg/s, volume in m^3, pressure in bar, enthalpy in kJ/kg, heat in kW and time in sec. Level is carried in meters, so the display factor of 1000 is the only unit conversion anywhere in the application. Every controller works in percent of transmitter span and percent of rated steam flow, because that is what a distributed control system faceplate shows and because it makes one gain number mean the same thing in all three architectures.

Table 4.1 What the Level Loop Interacts With. The Nominal Column Is the Full-Load Design Condition.

Symbol	Role	Nominal	Range
$\hat{L}$ (LT-101)	CV, measured	0 mm	±150 mm SP
q_f (FT-102)	inner CV, measured	100 %	0–160 % SP
x_{cmd} (FCV-101)	MV	70 %	0–100 %
steam demand	DV, measured	100 %	20–120 %
T_{fw}	DV, measured	160 °C	100–185 °C
p_{pump}	DV, measured	55 bar	40–70 bar
firing bias	DV, unmeasured	0 %	±25 %
LT bias	DV, unmeasured	0 mm	±40 mm
LT noise σ	DV, unmeasured	0 mm	0–12 mm
FT noise σ	DV, unmeasured	0 %	0–6 %
V_w	state	10 m^3	$\geq$ 0.2 m^3
m_{sd}	state	300 kg	$\geq$ 0
p	state	40 bar	5–90 bar
Q	state	53030 kW	0.05–1.4 Q_0
x (stem)	state	0.70	0–1

✦ DID YOU KNOW? Feedwater at Plant Scale

The two main feedwater pumps of Three Mile Island Unit 2 could supply its two steam generators with up to 15,500 gallons of water a minute [5], nearly a cubic meter every second. The boiler of this chapter is rated at 25 kg/s of steam.

4.3.1 Indicated Level, Water Volume and Sub-Surface Steam

The indicated level is algebraic rather than a state, and it is the relation every other result in this chapter depends on, as shown in Equation (4.1).

$$L = \frac{V_w - V_{w,ref}}{A_d} + k_{void}\,\frac{V_{sd} - V_{sd,ref}}{A_d}, \qquad V_{sd} = \frac{m_{sd}}{\rho_s(p)} \tag{4.1}$$

where L is the true level in m from normal water level, V_w is the drum water volume, A_d is the water-surface area, V_{sd} is the volume of steam held below the surface, k_{void} is the fraction of that volume that appears as indicated level, and the two reference volumes are the design values. Writing both terms as deviations makes L zero at the design condition by construction rather than by an offset. Setting k_{void} to zero deletes the second term, and with it the inverse response, which is the cheapest available test of what that term is responsible for. Equation (4.1) is the nonlinear counterpart of Equation (1.7). Its first term is the integrating water term of that equation, and its second term is the void term. The void relaxation time τ_{void} of Equation (4.3) below plays the part of τ_v there, since this chapter uses τ_v for the valve stroking time.

The water inventory is a pure integrator of the flow imbalance, corrected for the mass that has just left the liquid to become vapor, as shown in Equation (4.2).

$$\rho_w(p)\,\frac{\mathrm{d}V_w}{\mathrm{d}t} = q_f - q_s - \frac{\mathrm{d}m_{sd}}{\mathrm{d}t} \tag{4.2}$$

where q_f is feedwater flow and q_s is steam flow, both in kg/s. The bubble correction is worth about 2.5 % at 40 bar and it is what closes the mass balance. With $A_d = 20\,m^2$ and $\rho_w = 798.6\,kg/m^3$, a total loss of feedwater at full load drops the water term 1.57 mm/s.

The sub-surface steam is the state that carries shrink, and it relaxes toward the holdup the circulation would settle at, as shown in Equations (4.3) and (4.4).

$$\frac{\mathrm{d}m_{sd}}{\mathrm{d}t} = \frac{m_{sd}^{ss} - m_{sd}}{\tau_{void}} \tag{4.3}$$

$$m_{sd}^{ss} = \tau_{res}(Q)\left[\frac{Q}{h_{fg}(p)} - \frac{q_f\left(h_w(p) - h_{fw}\right)}{h_{fg}(p)}\right] \tag{4.4}$$

where τ_{void} is the void relaxation time, Q is the heat absorbed in kW, h_{fg} is the latent heat, h_w is the saturated water enthalpy and h_{fw} is the feedwater enthalpy, all in kJ/kg. The first term inside the bracket is steam made in the risers and the second is steam destroyed in the drum by the incoming subcooled water. The second term is shrink, written as a flow. Both m_{sd} and m_{sd}^{ss} are floored at zero.

Circulation is what makes the second term a function of load, as shown in Equation (4.5).

$$\tau_{res}(Q) = \frac{\tau_{res,0}}{\sqrt{0.25 + 0.75\,Q/Q_0}} \tag{4.5}$$

where $\tau_{res,0}$ is the riser residence time at full firing and Q_0 is the design heat input. Natural circulation is driven by the density head the risers develop, and that head falls with the firing rate, so the residence time runs from 12.0 sec at full firing to 13.6 sec at 70 %, 16.2 sec at 40 % and 18.1 sec at a quarter. The constant 0.25 floors the residence time at twice its full-load value, 24 sec, and it is a physical statement rather than a numerical convenience, because a drum boiler on standby still circulates and an unfloored (4.5) would predict a stagnant boiler at zero firing.

Pressure carries the energy balance and is where swell originates, as shown in Equation (4.6).

$$C_e\,\frac{\mathrm{d}p}{\mathrm{d}t} = Q - q_s\,h_{fg}(p) - q_f\left[h_w(p) - h_{fw}\right] \tag{4.6}$$

where C_e is the metal and water sensible heat per bar of saturation pressure, in kJ/bar. Saturated steam density is nearly proportional to pressure, so a pressure drop expands every bubble already under the water line through the $V_{sd} = m_{sd}/\rho_s(p)$ of Equation (4.1). Equations (4.1) and (4.6) taken together are swell. The feedwater term appears in both (4.6) and (4.4), which is why a feedwater step pushes the pressure down and the level down at the same time.

4.3.2 Firing Rate, Valves and Instruments

The fuel and air path is not part of the control problem. It is a load feedforward with a proportional pressure trim behind a first-order lag, as shown in Equations (4.7) and (4.8).

$$\tau_f\,\frac{\mathrm{d}Q}{\mathrm{d}t} = Q_{dem} - Q \tag{4.7}$$

$$Q_{dem} = Q_0\,u_t\left(1 + \tfrac{b_f}{100}\right) + K_m\,(p_0 - p) \tag{4.8}$$

where u_t is the governor opening, b_f is a firing-rate bias in percent, K_m is the boiler master pressure gain in kW/bar and Q_{dem} of (4.8) is clamped to $0.05\,Q_0 \le Q_{dem} \le 1.4\,Q_0$. The firing path is present because pressure is what carries a load change into the void volume. The governor follows the operator-set steam demand through a rate limit of $\dot{u}_t \le 0.5$ %/s by default.

Both flows are algebraic, as shown in Equations (4.9) and (4.10).

$$q_f = q_{f,\max}\,\frac{R^x - 1}{R - 1}\sqrt{\frac{\max(0,\ p_{pump} - p)}{\Delta p_0}} \tag{4.9}$$

$$q_s = K_t\,u_t\,p \tag{4.10}$$

where R is the equal-percentage rangeability, x is the stem position in the interval 0 to 1, $q_{f,\max}$ is the flow the valve passes wide open at design differential Δp_0, and K_t is the choked-nozzle coefficient in kg/(s bar). The characteristic in (4.9) is exactly zero at the seat and exactly one wide open. Its derivative is $R^x \ln R/(R-1)$, so at $R = 12$ the valve has 2.62 times the gain at 70 % open that it has at 31.2 % open, which is the opening that passes a quarter of rated flow. The square root in (4.9) is what makes a pump trip damaging, because a discharge pressure falling from 55 bar to 43 bar cuts the valve differential from 15 bar to 3 bar and the same stem position then passes $\sqrt{3/15} = 45\,\%$ of the flow it did. The stem itself strokes as a first-order lag of time constant τ_v toward the commanded position, with the derivative blocked at both stops so that the fixed-step integrator cannot pass through a limit and return.

Saturated properties come from quadratics fitted at 20, 40 and 60 bar, listed in Equation (4.11) with p in bar and T in °C.

$$\begin{aligned} \rho_s &= 0.4956 + 0.45663\,p + 8.306\times10^{-4}\,p^2 \\ \rho_w &= 923.09 - 3.7743\,p + 0.016556\,p^2 \\ h_w &= 677.3 + 12.8775\,p - 0.065625\,p^2 \\ h_{fg} &= 2098.1 - 11.32\,p + 0.04225\,p^2 \\ T_{sat} &= 161.745 + 2.849\,p - 0.0158625\,p^2 \\ h_{fw} &= 4.19\,T + 3.4\times10^{-4}\,T^2 \end{aligned} \tag{4.11}$$

The fits reproduce 30 and 50 bar to better than 0.5 %, and the pressure used for a property lookup is clamped to the interval 8 to 75 bar rather than extrapolated. At the design pressure they give $\rho_s = 20.09\,\mathrm{kg/m^3}$, $\rho_w = 798.6\,\mathrm{kg/m^3}$, $h_w = 1087.4\,\mathrm{kJ/kg}$ and $h_{fg} = 1712.9\,\mathrm{kJ/kg}$, against $h_{fw} = 679.1\,\mathrm{kJ/kg}$ for 160 °C feedwater.

The level transmitter is the last equation and the most consequential one. A differential-pressure cell infers a height by dividing a pressure difference by a density difference, and it uses the difference it was ranged with, as shown in Equation (4.12).

$$\hat{L} = \frac{\left(\frac{H}{2} + L\right)\left[\rho_w(p) - \rho_s(p)\right] + H\left[\rho_s(p) - \rho_s(p_{cal})\right]}{\rho_w(p_{cal}) - \rho_s(p_{cal})} - \frac{H}{2} \tag{4.12}$$

where H is the tap separation and p_{cal} is the pressure the cell was ranged at. Equation (4.12) returns L exactly when $p = p_{cal}$, which is what allows it to report only the real error. The controller then sees $\hat{L}$ plus an operator-settable bias plus band-limited noise, passed through a first-order transmitter lag τ_{LT}. Both flow measurements pass through their own lag τ_{FT}, and the steam flow measurement carries the flow transmitter noise because it reaches the valve through the feedforward.

4.3.3 The Design Point and the Numerics

Table 4.2 holds the constants. Everything else is algebraic in them and is recomputed once per controller execution rather than cached, because any parameter may be changed at any time. The derived design constants are $Q_0 = q_{s,0}(h_{fg,0} + \Delta h_{sub,0}) = 53030\,\mathrm{kW}$, $m_{sd,0} = \tau_{res,0}\,q_{s,0} = 300\,\mathrm{kg}$, $V_{sd,ref} = m_{sd,0}/\rho_s(p_0) = 14.93\,\mathrm{m^3}$, $q_{f,\max} = 58.6\,\mathrm{kg/s}$, $\Delta p_0 = 15\,\mathrm{bar}$ and $K_t = q_{s,0}/p_0 = 0.625\,\mathrm{kg/(s\,bar)}$, where $\Delta h_{sub,0} = h_w(p_0) - h_{fw} = 408.3\,\mathrm{kJ/kg}$ is the design subcooling.

Every scenario opens at a genuine equilibrium. The steam flow is computed from (4.10) at the scenario's demand, the feedwater is set equal to it, the firing is set to the value that makes (4.6) stationary, the sub-surface steam is set to (4.4), the stem is set to the opening that passes the required flow through (4.9), and the water volume is then chosen so that the *indicated* level sits on the operator's setpoint by inverting (4.12). A scenario therefore starts on setpoint even when its transmitter is miscalibrated. The application checks the result on itself, and the largest scaled state derivative at the design condition is below 10^{-12}.

Table 4.2 Constants of the Plant. Every One Is Operator Adjustable Except $V_{w,ref}$, Which Is Fixed in the Source.

Symbol	Meaning	Value	Unit	Visibility
$q_{s,0}$	rated steam flow	25	kg/s	adjustable
p_0	design drum pressure	40	bar	adjustable
C_e	pressure energy capacity	195000	kJ/bar	adjustable
$T_{fw,0}$	design feedwater temp.	160	°C	adjustable
A_d	drum surface area	20	m^2	adjustable
L_{span}	transmitter span	600	mm	adjustable
k_{void}	void-to-level coupling	0.526	—	adjustable
$\tau_{res,0}$	residence time, full fire	12	sec	adjustable
τ_{void}	void relaxation time	20	sec	adjustable
x_0	valve at design flow	70	%	adjustable
R	valve rangeability	12	—	adjustable
τ_v	valve stroking time	4	sec	adjustable
$p_{pump,0}$	design pump pressure	55	bar	adjustable
$\dot{u}_{t,\max}$	turbine ramp rate	0.5	%/s	adjustable
τ_f	firing response time	30	sec	adjustable
K_m	boiler master gain	2000	kW/bar	adjustable
τ_{LT}	level transmitter lag	3	sec	adjustable
τ_{FT}	flow transmitter lag	2	sec	adjustable
H	level tap separation	1.2	m	adjustable
p_{cal}	transmitter range pressure	40	bar	adjustable
L_{alarm}	deviation alarm	150	mm	adjustable
L_{trip}	low-low and high-high trip	250	mm	adjustable
$V_{w,ref}$	water volume at NWL	10	m^3	hidden

The engine of Section 1.6.1 integrates the plant by fourth-order Runge-Kutta on a fixed 0.05 sec step, which is eighty steps per valve stroke and is therefore sized against the fastest mode in the model. Both controllers execute on a 0.5 sec period, which is a distributed control system scan, with the first execution 0.5 sec after the one that initializes the controllers at $t = 0$, and the historian logs ten samples per simulated second. Noise is drawn from a generator seeded at 20260829, so scenario 13 replays identically on every machine. The speeds run from $0.5\times$ to $10\times$.

4.3.4 Assumptions

The model is detailed where a controller has to face the consequence and thin everywhere else. Each item below names the point at which the simplification stops being safe.

Five states and lumped saturated properties. This is a teaching-scale reduction of the published drum-boiler balances [11]. It reproduces shrink, swell, their dependence on load and the right-half-plane zero they create, and it says nothing quantitative about any particular boiler.

The whole circulation loop is one number. Equation (4.5) replaces the riser and downcomer system with a single residence time, so there is no distributed void profile up the riser, no slip between the phases and no downcomer subcooling [12]. There is therefore no circulation instability, which matters because a real natural-circulation boiler at very low load can circulate unstably and low load is where this application is otherwise most interesting.

Two volume effects are neglected. The thermal expansion of the drum water during a pressure transient is not tracked as a separate contribution, which understates swell by roughly one part in forty of the water-side term. The steam space above the water line is not a state, and at about 10 m^3 of steam space a pressure ramp of 0.05 bar/s is a 0.26 kg/s error against flows of 25 kg/s, so it costs about 1 % during the fastest transient and nothing at steady state.

The fuel side is not modeled as a control system. Equations (4.7) and (4.8) carry no air-fuel cross limiting, no firing rate limit and no pressure controller, all of which a real boiler master carries [16]. The firing bias is the only way to make the fuel side misbehave.

The turbine is a choked nozzle with a ramp limit. There is no steam line volume, no drum-to-throttle pressure drop and no superheater, so steam pressure control cannot be studied here and neither can its interaction with level control.

The instrument chain is idealized apart from the density error. Equation (4.12) is exact, but reference-leg boiling, tap plugging and the standing waves of a real drum surface are absent, and each of those is a documented reason a drum level indication is wrong [3, 17].

The valve has no stiction, backlash or positioner. Feedwater valve stiction is a common cause of real level cycling and is exactly the fault an inner flow loop half conceals, which makes its absence worth stating.

A trip does not stop the plant. Both level trips latch and annunciate without stopping the run, which no real protection system would do. The transient after a trip is the part worth reading, so it is kept on the trend. The only protective action is the feedwater interlock of Section 4.5.

The trajectory repeats better than a plant does. Fixing the step, the scan period and the seed produces one history on every machine, which a real distributed control system does not.

4.4 Inverse Response and Its Dependence on Load

The experiment that establishes the plant's central property takes no controller at all. In MANUAL at full load, with feedwater and steam balanced at 25.0 kg/s, FCV-101 is opened from 70 % to 78 % at 60 sec (the scripted step of scenario 1) and the feedwater rises to 31.8 kg/s. More water is entering than leaving and the drum is gaining mass, yet the level falls to −3.95 mm 19 sec after the step and does not return to where it started until 38 sec. The level is reported decomposed into its two terms, and at the bottom of the dip the water contribution is +7.1 mm against a void contribution of −11.1 mm. The mass balance is behaving correctly and the void term is larger.

> **TRY IT IN THE APP** **1 · Drive the Valve by Hand (Manual)**
>
> Load scenario 1 and press Run. At 60 sec the script opens the feedwater valve from 70 to 78 %, so more water enters than leaves. Watch the drum level: it falls first, to −3.95 mm 19 sec after the step, and does not return to where it started until 38 sec, because the entering water collapses steam bubbles.

The arithmetic is entirely in the enthalpy of Equation (4.4). Saturated water at 40 bar carries 1087.4 kJ/kg and 160 °C feedwater carries 679.1 kJ/kg, so every kilogram entering the drum absorbs 408.3 kJ before it can boil. Latent heat is 1712.9 kJ/kg, so each kilogram of feedwater condenses 0.238 kg of steam on its way to saturation. That steam occupied space at 20.09 kg/m^3 and the water replacing it occupies space at 798.6 kg/m^3, so a bubble that collapses returns 39.8 times its own mass in volume. The void loss exceeds the mass gain for the first half minute and is smaller than it afterward.

Pause. Try Exercise 4.1 before reading on.

Two properties of the response behave differently, and Table 4.3 separates them. Depth is a gain. Increasing the feedwater by 10, 25, 50 and 100 % at full load troughs at −1.47, −3.70, −7.46 and −15.21 mm, a factor of ten across a factor of ten in the step. The crossing time is a dynamic property. All four come back through zero between 38 and 39 sec. Confusing the two leads to the attempt to fix an inverse response by moving less, which buys a smaller excursion and no speed at all. The rows are model runs rather than interface procedures, because each starts from the exact equilibrium opening, the step positions fall between the half-point increments of the manual valve station and the scripted step of scenario 1 is removed. Table 4.3 reports the drum level rather than the transmitter output, because the pressure falls during the step and the differential-pressure cell of Equation (4.12)

Table 4.3 Open-Loop Feedwater Step in MANUAL, With the Valve Moved From Its Equilibrium Opening to the Position That Passes the Stated Increase in Feedwater. Trough Depth Scales With the Step and the Crossing Time Does Not. Times Are Measured From the Step.

	Trough (mm)	Time of trough (sec)	Crossing (sec)
Step size, at 100 % load			
+10 % feedwater	−1.47	19	38
+25 % feedwater	−3.70	19	38
+50 % feedwater	−7.46	20	38
+100 % feedwater	−15.21	20	39
Load, at a +25 % feedwater step			
100 % load	−3.70	19	38
70 % load	−4.05	23	47
40 % load	−3.73	26	58
25 % load	−3.04	29	65

therefore returns a trough about 7 % shallower at full load and 4 % shallower at a quarter, and a crossing about 2 sec earlier.

Pause. Try Exercise 4.2 before reading on.

Changing the load does the opposite. The trough stays within 0.7 mm of −3.7 mm across the whole range while the crossing time runs from 38 sec at full load to 65 sec at a quarter. Equation (4.5) is the mechanism. Slower circulation holds steam under the water line longer per kilogram being made, so the sub-surface steam volume falls only from 14.93 m^3 at full load to 8.05 m^3 at 40 % while the steam production falls to two fifths, and the void contribution to the indicated level runs from 0 mm at full load to −181 mm at 40 % and −244 mm at a quarter. The zero has moved toward the origin, and the closed-loop bandwidth a feedback controller can carry has come down with it [6], through the bound that follows Equation (1.7).

No numerical zero location appears anywhere in this chapter, and the omission is deliberate. The coefficients of (4.4) and (4.5) are functions of the operating point, so a single value would be a linearization about one load presented as a property of the boiler. The crossing time in Table 4.3 is what a step test on the plant returns, and it moves for the same reason the zero does. The output voltage of a boost converter carries a right-half-plane zero that also moves with load, and Volume 3, Chapter 9 measures what an inner current loop returns against it.

The second term of (4.1) can be tested by deletion. Zeroing k_{void} and rerunning the closed-loop load rejection of Section 4.6 gives no undershoot at all, +15.9 mm of overshoot and 6 % of valve travel, against −29.8 mm, +37.7 mm and 36 % with the term in place. The drum becomes an ordinary integrating level and every difficulty in the application disappears with it.

Swell is the same mechanism with the sign reversed. Steam density in (4.1) is a function of pressure, so anything that moves the pressure moves the indicated level with no change in inventory at all. A load pickup is the dangerous direction, because the firing rate lags the demand by $\tau_f = 30$ sec in (4.7) while the governor opens at up to 5 %/s, so the pressure falls, every bubble expands and the level climbs while the drum empties. That effect was quantified for rapid steam take-off on a real natural-circulation boiler decades ago [15].

4.5 Controller Design

Every architecture in the application drives the same feedwater valve through the same first-order stem, so the manipulated variable in the trends and in the comma-separated export is always FCV-101 and what changes is only what decides it. The three forms are given in Equation (4.13) and drawn in

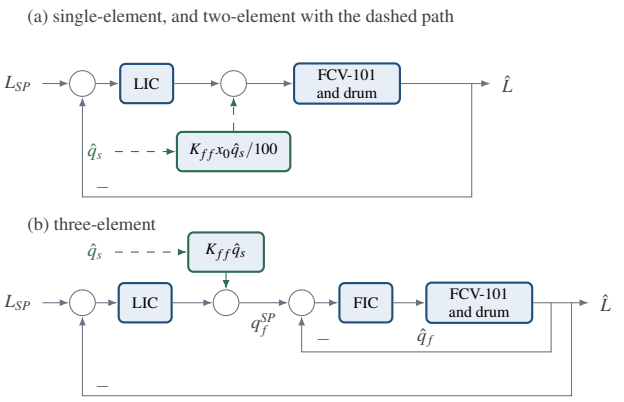


Figure 4.2 The three drum level architectures as implemented. All of them drive the same valve. Single-element control is the solid path in (a). Two-element control adds the dashed steam-flow feedforward, which has to be delivered as a valve position because there is no flow loop to receive it. Three-element control in (b) makes the level controller write a feedwater flow setpoint q_f^{SP}, states the feedforward in flow, and closes an inner loop on the measured feedwater flow $\hat{q}_f$. The output of the block on the right is the commanded valve position x_{cmd} of Figure 4.1.

Figure 4.2.

$$\begin{aligned} \text{single: } x &= \tfrac{1}{100}\,\mathrm{LIC}\big(L_{SP},\hat{L}\big) \\ \text{two: } x &= \tfrac{1}{100}\left[K_{ff}x_0\frac{\hat{q}_s}{100}+\mathrm{LIC}\big(L_{SP},\hat{L}\big)\right] \\ \text{three: } q_f^{SP} &= K_{ff}\hat{q}_s+\mathrm{LIC}\big(L_{SP},\hat{L}\big), \\ x &= \tfrac{1}{100}\,\mathrm{FIC}\big(q_f^{SP},\hat{q}_f\big) \end{aligned} \tag{4.13}$$

where $\hat{q}_s$ and $\hat{q}_f$ are the lagged steam and feedwater flow measurements in percent of rated steam flow, K_{ff} is the feedforward gain and x_0 is the design valve opening in percent. The feedforward is delivered as the output bias of the level controller rather than as a separate summing junction, so the integral term of the LIC carries exactly the difference between what the feedforward asked for and what the plant needed, and the application reports those two numbers side by side. A fourth structure, the model predictive controller of Section 4.5.3, writes the same flow setpoint q_f^{SP} to the same inner loop.

4.5.1 Control Algorithm

LIC-101 and FIC-102 run one algorithm between them, the velocity-form PID law of Section 1.6.2 [20], written here in the loop's own symbols. The filter and the three increments are Equations (4.14) to (4.17), and the recursion that assembles them is Equation (4.18). The running sum s_k is the S_k of Equation (1.14), and the demanded output c_k is the v_k of Equation (1.15).

$$\hat{L}_k^f = \alpha\,\hat{L}_{k-1}^f + (1-\alpha)\,\hat{L}_k, \qquad \alpha = \mathrm{e}^{-\Delta t_c/\tau_{df}} \tag{4.14}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1}) \tag{4.15}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{4.16}$$

$$\Delta u_k^D = -\frac{K_c\tau_D\left(\hat{L}_k^f - 2\hat{L}_{k-1}^f + \hat{L}_{k-2}^f\right)}{\Delta t_c} \tag{4.17}$$

$$
\begin{aligned}
s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\
c_k &= u_{bias,k} + s_k + I_{k-1} \\
u_k &= \min(u_{max},\ \max(u_{min},\ c_k)) \\
I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k)
\end{aligned}
\tag{4.18}
$$

where $e_k = L_{SP,k} - \hat{L}_k$ is the error in percent of transmitter span, $u_{bias,k}$ is the feedforward term of (4.13), $\hat{L}^f$ is the measurement filtered with a default τ_{df} of 2 sec for the derivative alone, Δt_c is the controller period, s_k accumulates the proportional and derivative increments without limit, c_k is the output the algorithm demands and T_t is the anti-windup tracking time. Equation (4.15) is written for a fixed gain. The app forms the increment as the change in $K_c e_k$, so a gain change acts on the whole error at the next execution. Both loops are reverse acting in the ISA sense, action $+1$ in the convention of Section 1.3, because more valve raises both the level and the feedwater flow, so a measurement above setpoint is met with less valve. The inverse response is a dynamic effect and does not change the sign of either process gain, which is the reason it is mistaken for one. The filter of (4.14) is advanced on every execution in MANUAL as well, so no derivative transient survives a transfer into AUTO. Windup protection is the back-calculation in the last line of (4.18), with T_t equal to each loop's own integral time (300 to 700 sec for the level controller in the shipped scenarios, and 8 or 120 sec for the flow controller). Only the integral is corrected, and the running sum s_k is never truncated. The LIC output is clamped to 0–100 % in single-element and two-element control, where it means a valve position, and to 0–160 % in three-element control, where it means a flow setpoint and a drum that has just swelled needs the headroom. The FIC output is clamped to 0–100 %. Neither controller carries an output rate limit, because the physical rate limit is the 4 sec valve stroke. Both execute every 0.5 sec, independently of the 0.05 sec integration step.

MANUAL inverts the same recursion for the integral, $I_k = u_{man} - u_{bias,k} - s_k$, as in Section 1.6.2. It places each controller output on the value it would have to deliver, the operator's valve position for the FIC and for the single-element and two-element LIC, and the measured feedwater flow for the three-element LIC. Changing architecture at run time is treated the same way, because the outer controller's output means a valve position in two of the three schemes and a flow setpoint in the third. For one execution after the switch both controllers track the valve the plant already has, which is the only way to avoid commanding a valve position of 25 % when the number in the integral was a flow setpoint of 25 %.

One protective action is implemented. After a high-high level trip, an interlock drives the feedwater valve shut, because water carried into a turbine is unrecoverable, and both controllers track the plant as in MANUAL while it holds, so neither integral winds up against it. It releases when the measurement is back inside the 150 mm deviation alarm rather than latching, and the release condition was chosen by measurement. Latched permanently, the interlock emptied the drum after every high trip, so a run that reached the high trip always finished at the low one and the transient that caused the original event became unreadable underneath a second, larger one.

4.5.2 Limits of a Feedforward Delivered as a Valve Position

The feedforward in three-element control is exact and needs no calibration, because a kilogram of steam leaving requires a kilogram of feedwater entering and that statement contains no valve, no pump and no operating point. $K_{ff} = 1$ is therefore the right answer rather than a tuned one. Two-element control has no flow loop to hand a flow setpoint to, so the same signal has to arrive as a valve position, and (4.13) scales the design opening x_0 by the steam flow fraction. That is a straight line through two points while the equal-percentage characteristic of (4.9) is a curve through the same two points, so the calculation is exact at one load and wrong everywhere else, as listed in Table 4.4.

The largest error in Table 4.4 is 14.5 points of valve travel at 40 % flow, and the sign reverses above the design point. Nothing in two-element control can remove it, because the scheme contains no measurement of what the valve delivered. Three-element control removes it by stating the feedforward

Table 4.4 Two-Element Feedforward Against the Valve It Is Driving. The Error Is Carried by the Level Controller's Integral Term.

Feedwater flow (% of rated)	Feedforward asks for (%)	Valve needs (%)	Error (points)
25	17.5	31.2	13.7
40	28.0	42.5	14.5
60	42.0	53.9	11.9
85	59.5	64.7	5.2
100	70.0	70.0	0
115	80.5	74.7	−5.8

in the units of the thing it is about, and the inner flow loop then absorbs the valve entirely, so the level controller faces a feedwater flow that follows its setpoint rather than an equal-percentage stem whose gain changes by a factor of 2.62 across the range.

4.5.3 Model Predictive Control

The model predictive controller (MPC) takes the place of LIC-101 in the three-element wiring of Equation (4.13) and writes q_f^{SP} to an unchanged FIC-102. It follows the receding-horizon scheme of Section 1.6.3 [21, 22], re-solving every $T_s = 2\,\mathrm{sec}$, or fourth execution. The voids of Equation (4.3) follow the measured steam flow, so a model can forecast a swell before the transmitter reports it.

The controller's model is a four-state reduction of Equations (4.1) to (4.8) with no pressure state, written from the parameters of Table 4.2 as they stood at the last reset, as shown in Equation (4.19).

$$\begin{aligned} \frac{\mathrm{d}\ell_w}{\mathrm{d}t} &= c_w\left(q_f - \hat{q}_s - \hat{d}\right), \qquad \tau_{void}\frac{\mathrm{d}b}{\mathrm{d}t} = b^* - b \\ \tau_{fl}\frac{\mathrm{d}q_f}{\mathrm{d}t} &= v - q_f, \qquad \tau_{th}\frac{\mathrm{d}z_f}{\mathrm{d}t} = q_f - z_f \\ b^* &= k_v\left[\tau_{res}(z)\,(z - \varepsilon q_f) - \tau_{res,0}\,q_{s,0}\right] \end{aligned} \tag{4.19}$$

Here ℓ_w is the level contributed by the water mass, counted from its estimate $\hat{L}_w$, b is the level contributed by the sub-surface steam net of the liquid it displaces, q_f is the delivered feedwater flow, z_f is that flow as the firing has replaced its sensible heat, v is the flow setpoint, $\hat{q}_s$ is the measured steam flow and $\hat{d}$ is an estimated unexplained outflow, all flows in kg/s. The heat input is counted as the steam flow it supports, $z = Q/h_{fg} = z_s + \varepsilon z_f$ with $\tau_{gov}\dot{z}_s = \hat{q}_s - z_s$, where $\varepsilon = \Delta h_{sub,0}/h_{fg} = 0.238$ and $\tau_{res}(z)$ is Equation (4.5) with $Q/Q_0 = z/[q_{s,0}(1+\varepsilon)]$. At 40 bar, $c_w = 1/(\rho_w A_d) = 6.261 \times 10^{-5}\,\mathrm{m/kg}$ and $k_v = (k_{void}/\rho_s - 1/\rho_w)/A_d = 1.2465 \times 10^{-3}\,\mathrm{m/kg}$, so b^* is k_v times the departure of Equation (4.4) from the design holdup, −172.3 mm at 40 % load against the plant's −180.9 mm, because the model holds the properties at design pressure. The predicted indication is $y = \hat{L}_w + \ell_w + b$.

The lags were identified by step tests on the simulator. The closed inner loop, stepped by ±10 % at the shipped flow tuning, fits a dead time of 1.15 to 1.25 sec and a lag of 2.05 to 2.43 sec at full load with residuals of 2.0 to 2.2 %, folded into $\tau_{fl} = 3.4\,\mathrm{sec}$ and scaled elsewhere as $\tau_{fl} = 3.4\,\mathrm{sec} \cdot R^{x_0 - x_q}(0.6/K_{c,F})(\tau_{I,F}/8\,\mathrm{sec})$, with x_q the opening that passes q_f. That gives 6.7 sec at 40 % load, where the fits sum to 6.9 to 7.4 sec, and 510 sec for the detuned loop of scenario 9 at full load. The firing follows a steam-flow change with $\tau_{gov} = 12.6\,\mathrm{sec}$ (fits of 11.99 to 13.25 sec) and replaces the sensible heat of extra feedwater through the pressure trim with $\tau_{th} = 104\,\mathrm{sec}$ (fits of 101.9 to 106.0 sec), with residuals of 11 to 16 % from the second-order shape of the trim. Run open loop beside the plant for 300 sec, the model's indication departs from the transmitter by 7.4 to 9.2 % of the excursion in root-mean-square after a valve step and by 15 to 20 % after a governor ramp. At each decision the model is linearized about the estimated heat input and the measured feed flow and discretized with an exact zero-order hold. The measured steam flow, held over the horizon, is the only

preview and enters through z_s as a known term, and it is the signal the three-element feedforward reads.

A steady-state Kalman filter estimates L_w, b, $\hat{d}$ and the transmitter output at every 0.5 sec execution under every structure, from the transmitter reading with an assumed noise of 2 mm and process noise per execution of 0.2 mm, 3 mm, 0.02 kg/s and 0.1 mm respectively, which gives the gain $[0.158,\ 0.914,\ -7.58\,\mathrm{kg/(s\,m)},\ 0.425]$. A persistent level error is absorbed into the random walk $\hat{d}$, the disturbance model that removes steady offset [23].

The plan chooses setpoints $v_1,\dots,v_{12}$ held over blocks of $1,1,1,1,2,2,3,4,6,9,15$ and 15 steps of a horizon of $N = 60$ steps, or 120 sec, longer than the 65 sec crossing time of Table 4.3 at a quarter load, to minimize Equation (4.20).

$$J = \sum_{k=1}^{N} e_k^2 + P e_N^2 + \lambda \sum_{j=1}^{12} \left(v_j - v_{j-1}\right)^2 + \rho_b s^2 + \sigma_b s \tag{4.20}$$

Here $e_k = y_k - L_{SP}$ in mm, v_0 is the setpoint applied now, $P = 10$, $\lambda = 300\,\mathrm{mm^2\,s^2/kg^2}$, and $\rho_b = 100$ per $\mathrm{mm^2}$ and $\sigma_b = 100$ per mm price the single slack $s \geq 0$ of the soft band $|e_k| \leq 150\,\mathrm{mm} + s$ at the deviation alarm. The only hard constraint is $0 \leq v_j \leq 1.6\,q_{s,0}$, the range of the three-element LIC. The valve keeps the clamp of FIC-102, neither PI loop carries a rate limit, and the trips are not constraints. Decisions are posed as deviations from v_0, so the small ridge that keeps the program strictly convex cannot pull the feed toward zero, and the program goes to the Goldfarb and Idnani dual active-set algorithm [24]. Only λ was tuned, on scenario 14, by the rule the predictive controllers of these apps share: among values from 0 to 3000 that reach no trip and use no more stem travel than three-element control, 152.6 %, keep the largest whose IAE at 600 sec lies within 10 % of the best. The best is 13801 mm s at $\lambda = 1$, and the rule keeps $\lambda = 300$ at 14844 mm s.

Any plan whose solver status is other than optimal, or whose first move is non-finite or outside the setpoint range, is discarded. LIC-101 tracks the applied setpoint at every execution, so it then takes over bumplessly and hands back once a valid plan returns. The first plan is solved at $t = 0$, and no plan was discarded in any run reported here. The controller reads only the transmitter, the measured flows, the setpoint and its copy of the parameters, which keeps the five-state plant, the instrument chain and the four-state model as three distinct objects.

4.6 Results

The app runs on the framework of Section 1.6.4. Manual mode positions the feedwater valve directly, automatic mode closes the architecture selected in Section 4.5, and every run exports as a comma-separated file. Fifteen scenarios are arranged as a laboratory sequence and are listed in Table 4.5. Each one presets the mode, the architecture, the tuning, the disturbances, any parameter it overrides and a list of timed events, and each figure quoted below came off the running application at the settings named next to it, changed before the run is started unless a Reset is named. Reset returns the plant to $t = 0$ and keeps every setting the learner changed, and reloading the scenario restores its own. Integrated absolute error keeps accruing after a transient ends, so every value reported here is the application's own running integral read 710 sec after the event that starts the transient, which makes the numbers comparable across controllers and meaningless in isolation. Valve travel accumulates in the same way and is read at the same instant. Two cases name their own window instead. The pump trip is still recovering at 710 sec, so its figures are read over the 900 sec following the trip, and the noise scenario applies no event at all and is read over its full fifteen minutes. The comparison with the predictive controller in Section 4.6.7 reads every case at the clock on which its lesson is credited instead. Peak deviation needs no reading instant. Every deviation quoted is the drum level rather than the transmitter output, except where the transmitter is the subject.

Each scenario also carries a completion rule, given in the last column of Table 4.5. A case is credited when the learner performs the listed actions and the simulation reaches the listed state, and

Table 4.5 Scenarios and Completion Requirements

#	Scenario (kind)	What it exercises	Required to complete
1	Drive the valve by hand *experiment*	inverse response with no controller	In MANUAL, the scripted valve step from 70 to 78 % at 60 sec followed to 120 sec. After Reset, a smaller step to between 70.5 and 76.5 % made before 60 sec, and after the scripted step feedwater matched to steam within ±0.3 kg/s for 30 sec.
2	Inverse response at 40 % load *experiment*	the zero moving with load	In MANUAL at 40 % load, the scripted valve step from 42.5 to 57.5 % at 60 sec followed to 150 sec, then the steam demand moved at least 10 points from 40 % and held there for 60 sec.
3	Single-element, 100 to 85 % *experiment*	one measurement that moves the wrong way	Single-element control at $K_c = 1.6$ held from the demand step at 90 sec to 600 sec, then after Reset a run at $K_c \geq 2.5$ held over the same window.
4	Same change at 40 % load *demonstration*	the same disturbance, harder plant	Single-element runs at 40 % load with $K_c = 1.6$, 1.0 and at least 2.5, each held from the step to 25 % at 90 sec to 600 sec.
5	Two-element feedforward *experiment*	a feedforward that fixes half the problem	Two-element control at $K_{ff} = 1$ with $K_c = 1.6$ and then $K_c \geq 3.5$, each held from the demand step at 90 sec to 600 sec.
6	Two-element at 40 % load *experiment*	a feedforward carrying a valve model	Two-element runs at 40 % load with valve rangeability 12 and then 2, each held from 90 sec to 600 sec.
7	Three-element control *experiment*	cascade plus an exact feedforward	Three-element control at $K_c = 3.5$ and then $K_c \geq 6$, each held from the demand step at 150 sec to 600 sec.
8	Three-element at 40 % load *experiment*	gain that no longer suits the load	Three-element runs at 40 % load with $K_c = 3.5$ and then $K_c \leq 2.5$, each held from 90 sec to 600 sec.
9	Inner flow loop detuned *demonstration*	the time-scale separation a cascade needs	Three-element runs with the flow loop at $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, at 0.2 and 30 sec, and at 0.6 and 8 sec, each held from 90 sec to 600 sec.
10	Feedwater pump trip *experiment*	a loop closed around the actuator	The pump trip at 90 sec ridden to 600 sec on two-element, three-element and single-element control in turn, each architecture in place from 90 sec.
11	Fast pickup, 25 to 120 % *demonstration*	swell, and the low-level excursion it leaves	Single-element control at $K_c = 1.0$ run until a level trip is reached, then runs held from the pickup at 60 sec to 600 sec at $K_c \geq 2.5$, on three-element control at $K_c = 3.5$, and on three-element control at $K_c = 3.5$ with the turbine ramp cut to at most 1 %/s.
12	DP cell ranged at 28 bar *demonstration*	an instrument fault no tuning reaches	The firing upset at 90 sec ridden to 600 sec with the cell ranged at 28 bar, then with it ranged at 40 bar.
13	Noisy level with derivative action *experiment*	actuator travel bought for nothing	Three idle runs held from the start to 900 sec, with $\tau_D = 20$ sec unfiltered, with $\tau_D = 0$, and with $\tau_D = 20$ sec behind a 5 sec derivative filter.
14	Three-element against MPC through a swell *experiment*	a model that forecasts the voids	At 40 % load with the turbine picking up to 100 % at 60 sec and 2 %/s, a run on the MPC held from 60 sec to 600 sec, then after Reset a run on three-element control selected before the pickup and held over the same window.
15	Graded challenge, three upsets *challenge*	choosing an architecture under time	Neither level trip reached in 900 sec through the pickup at 60 sec, the heater trip at 220 sec and the rejection at 380 sec, with the drum level inside ±30 mm of setpoint from 600 to 900 sec.[a]

[a] As shipped (single-element, $K_c = 1.0$, flow loop at 0.06 and 120 sec, governor ramp 2 %/s) the drum reaches the high-high trip at 559.5 sec and the worst level error over the last five minutes is 269.9 mm, 239.9 mm outside the band. The suggested settings (three-element, $K_c = 2.0$, $\tau_I = 500$ sec, $K_{ff} = 1$, flow loop at 0.6 and 8 sec) reach neither trip and hold the level within 14.8 mm, a margin of 15.2 mm, with the same result whether they are applied before the run or carried through Reset.

the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

4.6.1 Single-Element Control and the Limit on Gain

The architecture lessons require their settings to be held through the post-event comparison and are credited at $t = 600$ sec. Their Learn cards report that clock reading. Longer comparisons below state their horizon explicitly. In particular, Table 4.6 uses equal time after each event, accounting for scenario 7's later step.

Scenario 3 puts the level controller in service with drum level as its only measurement and backs the turbine from 100 % to 85 %, so the drum begins to gain water at 3.75 kg/s. The level falls to −29.8 mm at 168 sec, which is swell with the sign reversed. Less steam leaving raises the pressure, higher pressure compresses the bubbles, and the void contribution has to fall by 38 mm between 100 % and 85 % load. The controller reads a low level and opens the feedwater valve, which condenses more steam and holds the level down further, and by the time the mass balance exceeds the void term the drum has been overfilled. The level peaks at +37.7 mm at 286 sec, which is 67 mm peak to peak for a fifteen-point load change on the easiest part of the range.

Gain does not rescue it. Sweeping K_c through 1.0, 1.6, 2.0, 2.5 and 3.0 read at $t = 800$ sec gives integrated errors of 16074, 11486, 9938, 10414 and 18307 mm s for 17, 36, 60, 124 and 299 % of valve travel. The error stops improving at 2.0 and the loop oscillates at 3.0, on a disturbance that needs a five-point valve move. The number that barely moves is the instructive one. The trough runs only from −29.5 to −32.4 mm across the whole sweep, because the trough is the void collapse and feedback cannot reach it. What the gain buys is a faster recovery afterward and what it costs is overshoot, from +34.8 mm at $K_c = 1.0$ to +61.3 mm at $K_c = 3.0$. That is the right-half-plane zero setting a ceiling and gain converting the wrong-way response into an oscillation rather than into performance [13].

4.6.2 The Same Load Change at Reduced Load

Scenario 4 is the identical controller and the identical 3.75 kg/s at 40 % load rather than at full load. Peak deviation goes from −29.8 to −66.3 mm and integrated error from 11486 to 25341 mm s, a factor of 2.2. Two effects stack. The void swing per point of load is larger, 63 mm instead of 38, and the wrong-way part of the response now lasts 58 sec instead of 38 (Table 4.3), so the loop has to be slower at exactly the moment the disturbance has grown. Raising the gain makes it worse. At $K_c = 2.5$ the trough deepens to −83.8 mm and the overshoot reaches +123.3 mm with 162 % of valve travel, which is a full oscillation rather than a control response.

The obvious version of that comparison is wrong, and the application reports the reversal rather than hiding it. Measured at the same *relative* step instead, a fifteen percent load reduction from 40 % down to 34 %, the low-load case is slightly *better* than the full-load one, at −23.2 and +18.7 mm for 9107 mm s and 20 % of travel, against −29.8 and +37.7 mm for 11486 mm s and 36 %. That run is made on scenario 4 by setting the demand to 34 % just after its scripted step at 90 sec, which the 0.5 %/s governor ramp makes identical to a step to 34 %. The equal-percentage valve's gain falls faster with load than the bandwidth ceiling comes down, so a proportional disturbance at the bottom of the range is not harder than the same proportional disturbance at the top. The statement about the absolute change is the operationally relevant one. A house load, a soot blower, an auxiliary steam demand and a mill trip are fixed quantities of steam, every one of them is a larger fraction of a turned-down boiler than of a full one, and the same 3.75 kg/s arrives 2.2 times worse at 40 % load.

4.6.3 Comparison of the Three Architectures

Table 4.6 carries the three architectures on the identical load rejection at both loads. Two-element control makes the dip at full load *deeper*, from −29.8 to −50.2 mm, and removes the overshoot completely, and that combination is the most useful result in the application. The feedforward makes feedwater equal steam, which is exactly right for the water inventory of (4.2) and does nothing at all about the voids of (4.3). The bubble volume genuinely has to shrink by 38 mm between the two loads, so holding the mass constant means the level genuinely has to fall by 38 mm. What the feedforward

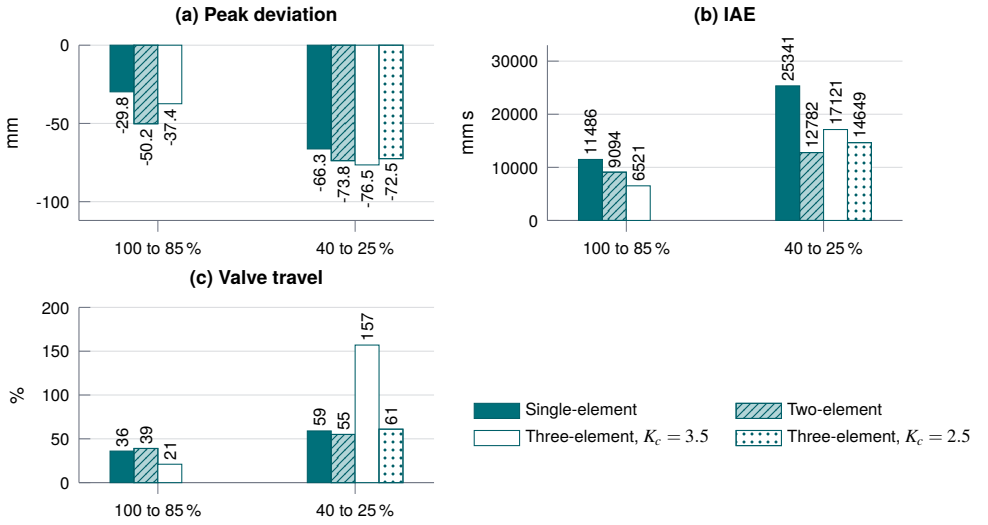


Figure 4.3 The fifteen-point load rejection of Table 4.6, from 100 to 85 % load and from 40 to 25 % load, under the three architectures: (a) peak deviation, (b) integrated absolute error read 710 sec after the load change and (c) valve travel. Solid: single-element. Hatched: two-element. White: three-element at $K_c = 3.5$. Dotted: three-element at $K_c = 2.5$, run at the lower load only. Notice that the ladder ranks cleanly at full load, while at the lower load two-element control posts the smallest error and the full-load gain of three-element control costs 157 % of travel.

Table 4.6 The Architecture Ladder on the Identical Fifteen-Point Load Rejection. Every Integral and Travel Is Read 710 sec After the Load Change. The Low-Load Runs Are the Same 3.75 kg/s of Steam.

	K_c	Peak (mm)	IAE (mm s)	Travel (%)
100 % to 85 % load				
Single (sc. 3)	1.6	−29.8	11486	36
Two (sc. 5)	1.6	−50.2	9094	39
Three (sc. 7)	3.5	−37.4	6521	21
40 % to 25 % load				
Single (sc. 4)	1.6	−66.3	25341	59
Two (sc. 6)	1.6	−73.8	12782	55
Three (sc. 8)	3.5	−76.5	17121	157
Three (sc. 8)	2.5	−72.5	14649	61

removed is the part of the excursion the controller was adding on its own. Figure 4.3 draws Table 4.6 as bars.

🕮 **Pause.** Try Exercise 4.3 before reading on.

The two-element feedforward gain is more interesting than one for one suggests. For an independent comparison, the learner reloads scenario 5, sets the gain before Run and reads every result at $t = 800$ sec. The initialized integral balances the selected feedforward. Gains 0, 0.25, 0.5, 0.75 and 1.0 give 11486, 8328, 5601, 5776 and 9094 mm s, with the same values when carried through Reset. Zero feedforward reproduces single-element under this startup procedure. A gain changed during a run can also disturb the existing integral balance, so it is a separate experiment. A half-strength feedforward splits the difference between the mass error and the void error at full load. At 40 % load the same 800-second comparison gives 18536 mm s at half strength against 12782 mm s at unity. The

preferred gain depends on load, which is a limitation of delivering feedforward as valve position. Three-element control instead states the feedforward in flow.

Three-element control is also the only architecture that tolerates gain. At $K_c = 6$ it still works, at 5809 mm s for 58 % of travel, where two-element control at $K_c = 3.5$ has already lost stability at 31635 mm s and 629 % of travel. What it does not do is remove the physics. The remaining 37.4 mm in Table 4.6 is the void collapse itself, and no architecture can prevent a void volume that has to shrink from shrinking.

The lower block of Table 4.6 does not rank the same way, and this chapter reports that rather than selecting the load at which the expected ordering holds. On the load rejection at 40 % load it is two-element control that posts the lowest integrated error at 12782 mm s, against 17121 mm s for three-element control at the full-load gain of 3.5 and 14649 mm s with the gain backed to 2.5. Backing the gain off is worth doing anyway, because it cuts the travel from 157 to 61 % for a 14 % improvement in error, and it is the practical case for gain scheduling this loop on load. What the load rejection does not separate, the pump trip does.

4.6.4 Inner Flow Loop Tuning and the Pump Trip

Scenario 9 breaks the cascade without changing the wiring. The following comparisons are read at $t = 800$ sec, 710 seconds after its load step. Detuning the flow controller to $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, which is slower than the level loop above it, raises the integrated error to 20683 mm s and leaves the level 45.2 mm high at 569 sec. Valve travel falls to 11 %, the lowest figure anywhere in the application, and the low figure is the misleading part. A flow loop that takes two minutes to reach setpoint is a slow valve positioner rather than a flow loop, so the level controller is facing the equal-percentage valve directly again. Nearly all of the recovery arrives with the first step away from the sluggish setting, at 7305 mm s for 0.2 and 30 sec and 6521 mm s for 0.6 and 8 sec. An inner loop does not have to be fast in absolute terms. It has to be several times faster than the loop above it. A fermenter oxygen cascade makes the same inner-loop speed a measured quantity (Volume 3, Chapter 3).

Scenario 10 decides the argument, and it involves no load change. A feedwater pump trips and the discharge pressure falls from 55 bar to 43 bar, so the differential across FCV-101 collapses from 15 bar to about 3 bar and (4.9) passes 45 % of the flow it did. On two-element control the feedwater falls from 25.0 to 10.4 kg/s within ten seconds with the valve still sitting at 69 % open, and nothing in the scheme measures it. The level has to fall far enough for the level controller to respond, and by then the level is falling at more than a millimeter a second. The bottom is −149 mm at 340 sec with the valve on its stop, the integrated error over the 900 sec following the trip is 86511 mm s, and the level is still 62 mm low thirteen minutes later. Single-element control is worse again at −164 mm and 99200 mm s. Three-element control does not encounter it as a level problem at all. FT-102 reports the collapse on the next scan, the inner controller drives the valve from 70 to 95 % inside ten seconds, and the level reaches −7.6 mm before recovering, for 1459 mm s over the same window and about the same valve travel. That is a factor of 59 on the integrated error for a fault that no level tuning addresses.

TRY IT IN THE APP **10 · Feedwater Pump Trip, on Two-Element Control**

Load scenario 10 and press Run. When the pump trips, the feedwater falls from 25.0 to 10.4 kg/s within ten seconds while the valve stays near 69 % open, and the level bottoms at −149 mm at 340 sec. Reload, select three-element control before the trip at 90 sec and run again: the level now reaches only −7.6 mm.

The property the inner loop supplies is coverage rather than speed. Everything that can go wrong between the controller output and the water entering the drum, a pump, a strainer, a stuck stem, worn trim or a suction pressure swing, is inside the inner loop and outside every other scheme.

4.6.5 Transmitter Range Error, Swell and Derivative Action

Scenario 12 ranges the differential-pressure cell at 28 bar and runs the drum at 40 bar. Equation (4.12) then reads 19.7 mm below the true level from the first sample. Denser water makes more pressure per millimeter, so a cell that assumes the water is denser than it is reports fewer millimeters than are there. The controller holds the indication on setpoint exactly, so the water sits 19.7 mm above normal water level and that much high-side trip margin has been spent by a calibration sheet rather than by an engineer. The firing-rate upset moves it in the direction that matters. Drum pressure climbs from 40.0 to 42.4 bar, the error grows to −23.6 mm, and at 600 sec the true level is +26.1 mm while the transmitter reports +2.6 mm. A correctly ranged cell subjected to the same 2.4 bar rise produces −3.5 mm of error, so the residual sensitivity is about 1.5 mm/bar even when nothing is wrong, which is small until a unit moves its pressure across a ten bar range in normal operation. Two consequences follow and neither is a control problem. Feedback cannot see the error, because a loop holding a biased measurement holds the bias and reports nothing, and the trips read the same transmitter [17].

Scenario 11 produces a trip. The boiler idles at 25 % load and the turbine picks up to 120 % at 5 %/s, which is faster than a large machine is permitted to move. Steam leaves the instant the governor opens, the firing rate needs half a minute to catch up, the pressure falls, and every bubble under the surface expands. The level rises to +225 mm at 136 sec, which the transmitter reports as +228.5 mm and is 21.5 mm short of the high-high trip, while feedwater falls short of steam by up to 31.1 kg/s. The controller reads a high level and closes the valve. The voids then stop growing, the mass balance is left with the deficit, and the level falls through zero, through the low alarm and into the low-low trip at 421 sec, bottoming at −254 mm. The excursion that threatens the boiler is therefore the second one and not the swell that caused it.

Two of the three available remedies work and the obvious one does not. Raising the gain to 2.5 avoids the low trip at −159.6 mm and costs 310 % of valve travel by 770 sec, and three-element control at $K_c = 3.5$ holds the falling half to −53.9 mm but reaches the high-high trip at 124 sec on the way to +265 mm, so each remedy trades one trip against the other. Slowing the governor ramp does not help at all. Reducing it from 5 to 2, 1, 0.5 and 0.25 %/s on the shipped tuning brings the indicated peak down from +228.5 mm through 221.6, 198.5 and 138.3 to +86.0 mm, and every one of those runs still reaches the low-low trip, at 432, 450, 494 and 627 sec respectively, because the mass deficit that produces the second excursion is set by the size of the load change and by how long the controller holds the valve shut, and not by how quickly the load arrives. A ramp rate limit protects the high side of the drum and leaves the low side to the controller. The two together divide the work. Three-element control at $K_c = 3.5$ with the ramp cut to 1 %/s reaches neither trip, with the indication peaking at +226 mm and the level bottoming at −29.6 mm.

Scenario 13 applies no disturbance at all and is one of the most expensive runs in the application. A drum surface at full boil is not flat, so the transmitter carries about 6 mm of noise, and derivative action is switched on at $\tau_D = 20$ sec with the filter removed. Fifteen minutes at steady load cost 2091 % of valve travel, twenty full strokes of a stem that is holding a constant level, and finish at 1121 mm s of integrated error against 1018 mm s with the derivative switched off at 208 % of travel. Three additional seeds give 2024 % to 2121 % of travel with derivative against 203 % to 210 % with PI, and derivative adds 227 mm s to 304 mm s of integrated error. Derivative action is wrong on this

Table 4.7 Three-Element Control and the Predictive Controller on the Same Cases, Read at the Lesson Clock

Case	Structure	Peak (mm)	IAE (mm s)	Stem (%)	Command (%)	Limit (%)	HH (mm)	LL (mm)	Last 5 min (mm)	Trip
7, load 100 to 85 %	three	−37.4	5451	20.9	21.3	0	238	212	14.3	none
	MPC	−22.2	3101	14.5	31.0	0	246	228	8.6	none
8, load 40 to 25 %	three	−76.5	14595	136.4	138.4	0	198	174	51.8	none
	MPC	−44.9	6455	42.2	50.5	0	242	205	7.6	none
9, detuned flow loop	three	+45.2	13501	8.1	8.2	0	205	218	45.2	none
	MPC	−29.5	6257	11.8	15.0	0	232	220	17.6	none
10, pump trip	three	−7.6	1083	31.1	36.8	0.6	246	242	1.6	none
	MPC	−5.7	554	32.6	34.4	10.7	246	244	0.5	none
11, 25 to 120 % at 5 %/s	three	+264.7	44448	184.7	191.1	15.7	−19	195	53.9	HH
	MPC	+217.0	25600	108.5	130.5	14.0	30	232	16.8	none
14, 40 to 100 % at 2 %/s	three	+167.6	27202	152.6	154.7	8.3	80	199	49.9	none
	MPC	+122.1	14844	102.7	116.7	5.7	127	238	11.4	none
15, suggested settings	three	−154.2	24868	69.7	74.4	0	214	97	14.8	none
	MPC	−83.4	14540	67.3	161.8	0	229	168	14.7	none
15, as shipped	single	+285.7	66180	211.1	211.7	13.8	−45	72	269.9	HH
	MPC	+106.1	41337	28.3	30.4	44.2	141	160	106.1	none

Peak and IAE use the true level, IAE being the scorecard value. Stem is the travel of the valve stem x and Command that of the valve command x_{cmd}, the scorecard's manipulated-variable travel. Limit is the share of samples with the valve command at 0 or 100 % or the flow setpoint at 0 or 160 %, whereas the scorecard's saturation counts the valve command alone and reads 0 on every MPC row but the pump trip. HH and LL are the smallest margins of the transmitter reading to the trips, and Last 5 min is the largest true-level error over the final 300 sec. Each MPC run made one plan per 2 sec with none rejected.

loop for a second reason beyond the noise. It acts on the direction the measurement is moving, and on a process with an inverse response the measurement moves the wrong way for the first minute, so an accurate derivative drives the valve away from where the mass balance needs it.

4.6.6 Graded Challenge

The challenge, scenario 15, applies a load increase to 115 % at 60 sec, a feedwater heater trip at 220 sec and a fifty-point load rejection to 65 % at 380 sec, starting on single-element control at $K_c = 1.0$ with the flow loop detuned to $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, and with the governor ramp at 2 %/s. The target is neither trip reached and the level inside ±30 mm for the last five minutes. As shipped the boiler trips, swinging to +285.7 mm and reaching the high-high trip at 559.5 sec after having been 181.5 mm low, for 211 % of travel. Two-element control at $K_c = 1.6$ trips the other way, reaching the low-low trip at 486.6 sec and −254.4 mm, retaining the shipped $\tau_I = 700$ sec after changing the architecture and gain. Three-element control at $K_c = 2.0$, $\tau_I = 500$ sec, $K_{ff} = 1$ and an inner loop at 0.6 and 8 sec rides all three events, with 24868 mm s of error at $t = 900$ sec and a worst excursion of −154.2 mm, 70 % of travel and the level inside 14.8 mm for the last five minutes. The middle event carries the mechanism. Feedwater at 110 °C instead of 160 °C raises the subcooling in (4.4) from 408 to 622 kJ/kg, so every kilogram entering the drum condenses 0.363 kg of steam instead of 0.238 kg. Nothing about the flows changed. The shrink is permanently 52 % stronger, and the level is the only measurement in the scheme that reports it.

4.6.7 Comparison With the Predictive Controller

Table 4.7 sets the MPC against three-element control on six scenarios and the challenge, each read at the clock on which its lesson is credited, 600 sec or 900 sec on the challenge, so its integrals do not match those of Table 4.6. Three-element control keeps each case's own outer tuning, switched to three-element before the run in scenarios 10 and 11, and the suggested settings on the challenge. The MPC keeps the settings of Section 4.5.3 throughout, both act through the case's own flow loop and read the same measured steam flow, and both see the same events, instruments and noise seed. The challenge is also run as shipped against the MPC selected as the only change.

Scenario 14 is the case added with the controller. The boiler idles at 40 % load, and from 60 sec the turbine picks up to 100 % at 2 %/s. The void contribution rises from −180.9 mm to −0.5 mm at 600 sec, so with the indication on setpoint the inventory has to fall by as much. Three-element control raises its flow setpoint with the steam flow, to 81.2 % at 89.5 sec, and drives it to zero only once the

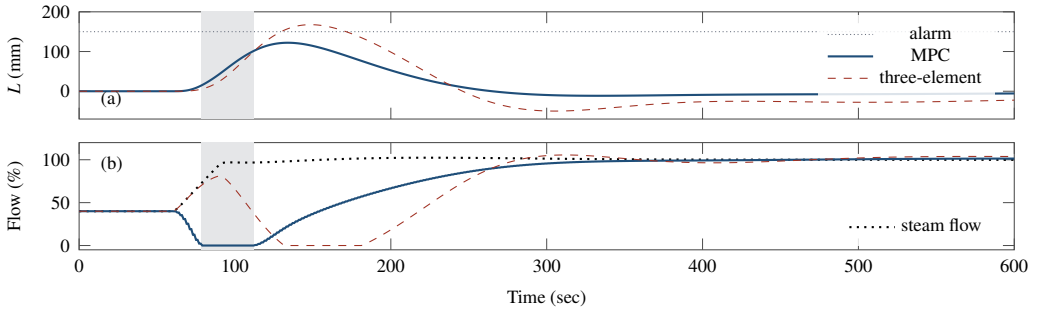


Figure 4.4 Scenario 14 under the predictive controller and under three-element control, with the same seed, over the ten minutes to the lesson clock. The turbine picks up from 40 to 100 % load from 60 sec at 2 %/s, and the shaded span from 78 to 112 sec is the interval in which the MPC holds the feedwater flow setpoint at zero. (a) True drum level under the MPC (thick solid) and under three-element control (dashed), and the 150 mm deviation alarm (thin dotted). (b) Feedwater flow setpoint under each structure, in the same line styles, and the measured steam flow of the MPC run (thick dotted), in percent of rated flow, the latter within 3.1 points of that of the three-element run.

swell reaches the transmitter, from 131.5 to 181.6 sec, after which the drum must be refilled. The MPC forecasts the void growth from the same steam flow and holds its setpoint at zero from 78 to 112 sec, while the steam flow is still rising, so the inventory falls before the voids arrive, as shown in Figure 4.4. The level peaks at +122.1 mm at 133.9 sec against +167.6 mm at 148.4 sec, beyond the 150 mm alarm, the transmitter comes within 127 mm of the high-high trip against 80 mm, and the trough that follows is −11.4 mm against −49.9 mm. The feedforward encodes only the steady mass balance, while the model also carries the void growth that occupies the drum first.

On scenarios 7 to 11 the MPC also lowers both the peak and the IAE, by factors of 1.7 to 2.3 in IAE. At 40 % load it does so with 42.2 against 136.4 % of stem travel, where the full-load gain of 3.5 suits three-element control least. On the pump trip the valve command sits at its upper stop for 64 sec, the 10.7 % of the Limit column, while the setpoint averages 6.7 points above what the tripped pump delivers. On the fast pickup of scenario 11, three-element control reaches the high-high trip at +264.7 mm, while the MPC peaks at +217.0 mm with its band slack reaching 62.8 mm, a plan that priced the band rather than obeyed it. On the detuned loop of scenario 9 it more than halves the IAE with more stem travel, 11.8 against 8.1 %.

The stem column understates what the MPC asks of the valve. Each 2 sec setpoint step draws an immediate command move of $K_{c,F}$ times the step from Equation (4.15), partly returned as the flow follows, and the 4 sec stem filters most of it out. On scenario 14 the command travels 116.7 % against 102.7 % on the stem, still below three-element control's 154.7 %. On the challenge with the suggested flow loop it travels 161.8 % against 74.4 %, more than twice as far, although the two setpoints travel 217.9 and 204.4 % and the two stems 67.3 and 69.7 %, and it also travels further on scenarios 7 and 9.

With the suggested settings both structures pass the challenge, the MPC at 14540 mm s against 24868 mm s. Selected as the only change to the challenge as shipped, the MPC reaches no trip but fails. From 564.1 sec its setpoint sits at zero while the detuned flow loop, $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, closes the valve command from 65.1 to 54.0 %, 1.98 points a minute, and the level ends 106.1 mm above setpoint against the 30 mm allowed. No setpoint in range could have closed the valve faster, so the remedy is the inner loop, as for three-element control, and with the suggested flow tuning the same MPC passes. In Chrome on scenario 14 the mean solve took 238 to 279 µs over five runs and the largest 2.6 to 3.7 ms, measured on a 20-core host at a load average near 2, or 0.014 % of a core at real time.

✦ DID YOU KNOW? No Trip, Still a Fail

From the simulator: on the challenge as shipped, selecting the predictive controller as the only change cuts the stem travel to 28.3 % against 211.1 % for the shipped single-element loop and reaches no trip. The largest level error over the last five minutes is still 106.1 mm against the 30 mm allowed (Table 4.7).

4.6.8 Numerics of the Running Application

Three clocks run, as Section 1.6.1 describes, the 0.05 sec Runge-Kutta step for the physics, the 0.5 sec period for both controllers and ten historian samples per simulated second. A scheduled scenario event is applied at the start of the step that begins at its time, before the controller execution due on that step. Within an execution the order is fixed, the governor ramp first, then the instrument chain with its density error, bias, noise and lag, then the feedforward, then the Kalman filter of Section 4.5.3, then either the controllers or the manual tracker, then the integration, then the latched trips and the travel counters, then the outputs and the log. No controller reads a state directly. The export carries 87 columns at ten samples per simulated second, including the true level, the indicated level and the level the controller acted on as three separate columns, the water and void contributions apart from each other, the feedforward and trim split, and the predictive controller's status, slack, solve time and estimates, which is enough to reproduce every number in Section 4.6 away from the browser. The application also runs nineteen checks on itself, covering the exact equilibrium of the design condition, the sign of each flow against its valve, the collapse of the sub-surface steam under extra feedwater, the expansion of the voids under falling pressure, the growth of the residence time at reduced firing, the exactness of (4.12) at its calibration pressure, the valve characteristic and the unity feedforward. One of them runs the plant open loop for 520 sec and asserts that the level goes down first and up afterward. Seven cover the predictive controller, among them its limits on every scenario, its source boundary, a least-squares check on one unconstrained move and a bumpless forced rejection.

Figure 4.5 catches scenario 7 in the middle of its recovery and carries two results worth stating. The first is the integrated absolute error, 3256 mm s at $t = 300$ sec against the 6521 mm s the same run reports 710 sec after the load change, which is what accumulation after a transient looks like and is the reason every integral in this chapter names its reading instant. The second is the differential-pressure cell error, +0.6 mm on a correctly ranged transmitter, because the drum pressure has fallen 0.42 bar below the pressure the cell was ranged at and Equation (4.12) is exact at one pressure only.

4.7 Discussion

Drum level has an unusual property among control loops. It is a frequent cause of unit trips, and the cause is almost never the tuning [1, 3]. The application is built so that the numbers behind that statement can be produced rather than asserted.

Nothing in scenario 10 is mistuned. The controller does exactly what its measurements support and takes the level to −149 mm because two-element control contains no measurement of feedwater flow. Nothing in scenario 12 is mistuned either. The controller holds the indication perfectly while the water sits 19.7 mm above where it is reported to be, and every self-diagnostic on the transmitter passes. An investigation that opens with the tuning constants finds nothing wrong with them in either case, which is a large part of why these events repeat. The general shape is the most dangerous class of automation fault, a system working correctly on information that is wrong, with the evidence available only outside the loop. In a boiler house that evidence is the gauge glass, which is an independent measurement with a different failure mode.

Trip settings are a second place where a control engineer's numbers should be informing a safety decision. The ±250 mm trips protect tubes and turbine blades and their values come from the boiler maker and the turbine maker, but whether they are reachable is a control question with a measurable answer. Scenario 11 reaches the low-low trip from a swell that stops 21.5 mm short of the high-high trip, and the challenge reaches the high-high trip from a tuning that looks unremarkable on a faceplate.

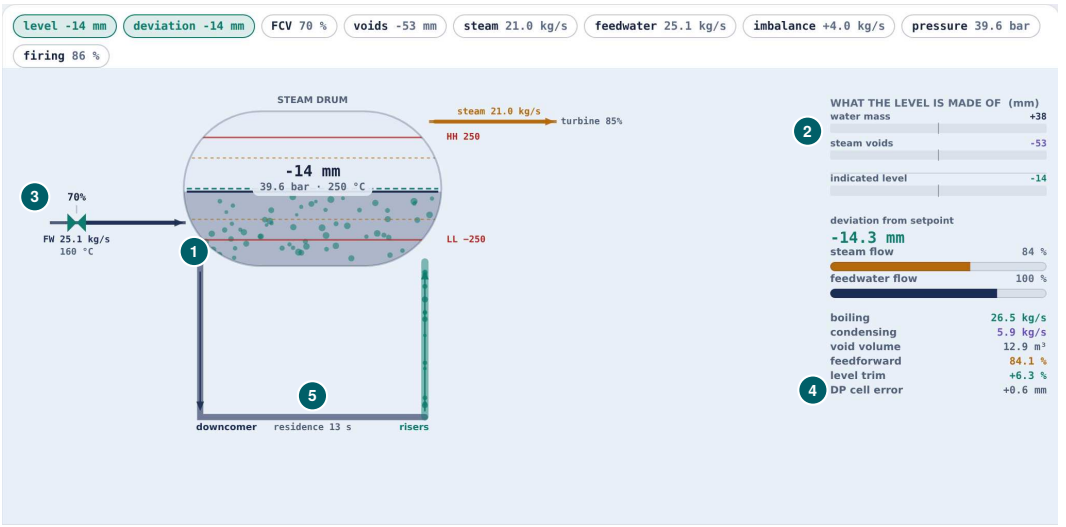


Figure 4.5 Three-element control riding the fifteen-point load rejection of scenario 7, caught at $t = 300$ sec with the disturbance 150 sec old, in the plant view of the app. The steam flow has ramped to 84.1 % of rated and the feedwater is at 100.2 % following an inner setpoint of 99.5 %, so the drum is gaining 4.03 kg/s while the level reads −14.3 mm. The level is decomposed into a water contribution of +38.3 mm against a void contribution of −52.6 mm, recovering from −59.7 mm at 265 sec. The drum pressure has fallen back to 39.58 bar from 40.59 bar at 202 sec. The crop keeps the plant view, whose readouts round these values, and cuts the trends, the control panel and the scorecard. Callouts: (1) steam bubbles held below the water surface, (2) the level split into its water and void terms, (3) feedwater valve FCV-101 with the feedwater flow and temperature, (4) the error of the differential-pressure cell, (5) the downcomer and the risers with their residence time.

The useful product is a table of worst excursion against load, ramp rate and architecture, which states a trip setting and an operating envelope in the same units.

Valve travel belongs on the same page as the error. Scenario 13 spends 2091 % of travel and finishes 103 mm s worse on integrated error than the same loop with the derivative switched off, and scenario 9 posts the best travel figure in the application, 11 %, alongside one of the worst errors. Read on its own, neither number says what happened.

The architecture is the decision with the longest consequence. Moving between the rungs of Figure 4.2 is not a tuning session. It is wiring, a flow element in the feedwater line, a second transmitter and a second controller in the distributed control system. A plant commissioned on two-element control because it was going to be a base-load machine, and now cycled because the grid asks it to be, is running the wrong architecture and no amount of retuning fixes it [18]. Scenario 4 against scenario 8 is that argument with numbers on it, and the judgment that has to be made at design time is a forecast about how the plant will be operated in fifteen years.

> **OPEN CHALLENGE** **Level Loops Tuned for Full Load, Run at Part Load**
>
> Plants asked to follow the load find level loops tuned for another duty. Feedwater heater level controls are typically tuned for full load and may not hold the level at reduced loads, which wears heaters, drain lines and valves, causes water hammer and, in extreme cases, trips the turbine [25]. The effective water level of a steam generator also changes with load, and one fleet's study of deeper load following lists steam generator level control among the items to assess [25]. Progress would be level control whose tuning or structure follows the load, the gain schedule that Section 4.6 argues for.

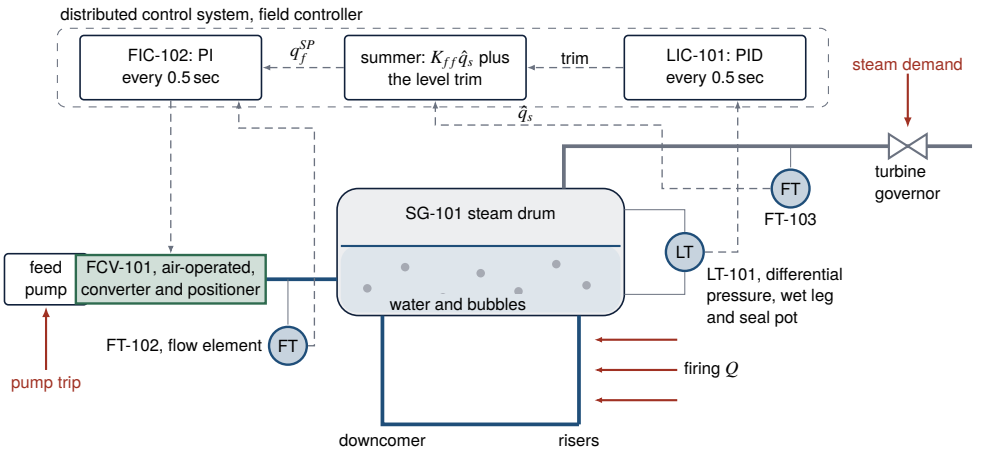


Figure 4.6 The three-element loop of Figure 4.1 with the hardware of a plant. A differential-pressure transmitter (LT) on two drum taps, its upper line a condensate-filled wet leg, reads the level, and flow transmitters (FT) on flow elements read the feedwater and the steam. In the field controller LIC-101 writes a trim that a summer adds to the steam flow feedforward, and FIC-102 holds the resulting feedwater setpoint with the air-operated valve FCV-101, both every 0.5 sec. The turbine governor follows the steam demand and is not moved by the loop. Thick solid lines carry water (blue: the feed line, the downcomer and the risers) and steam (gray: the line to the turbine), thin dashed arrows signals, and red arrows the disturbances.

4.8 Technology in Practice

The tags of Figure 4.1 are, in a plant, a differential-pressure level transmitter on the drum, two flowmeters, an air-operated feedwater valve and the controllers of a distributed control system. Figure 4.6 draws the three-element loop with that hardware in place of the schematic's symbols.

What Measures It

The three elements are three measurements. Drum level comes from a differential-pressure transmitter across two taps on the drum, the most common level method in nuclear plants [26]. Feedwater and steam flow are most often measured the same way, as the pressure drop across an orifice, a Venturi tube or another flow element, a square-root signal that the transmitter or the signal processing makes linear [26]. The chapter carries LT-101 as Equation (4.12) behind a 3 sec lag and both flow transmitters behind 2 sec lags.

> **TECHNOLOGY** **Sensor: Drum Level From a Differential Pressure and a Wet Leg**
>
> Drum level is read as the difference in pressure between a tap in the water and a tap in the steam. Steam condenses in the upper line and fills it, and the constant head of this wet leg is canceled by an intentional zero shift at calibration. Seal pots are standard on steam drums, because a wet leg that has lost water refills slowly and the level reads wrong until it does [27]. The transmitter sends 4 to 20 mA, and the live zero at 4 mA reveals an open circuit [26]. LT-101 adds the density error of Equation (4.12) and a 3 sec lag.

What Moves It

The final element is FCV-101, an equal-percentage valve whose stem follows its command as a 4 sec lag. In pressurized water reactors the feedwater regulating valves are air operated, and they are the valves to which smart positioners have most often been fitted [28]. In safety-related designs, solenoid valves between the positioner and the actuator can dump the air so that the spring returns the valve to its fail-safe position [28]. The inner flow loop shields the level controller from loads in the feedwater

supply and keeps the valve's nonlinear response its own problem [27], which is what the pump trip of scenario 10 measures.

> **TECHNOLOGY** **Final Element: The Air-Operated Feedwater Valve**
>
> In an air-operated valve, air on a diaphragm moves the stem against a spring, and the spring moves it back when the air is released. A converter turns the controller's 4 to 20 mA into 3 to 15 psig, and a positioner adjusts the actuator air until the valve reaches the demanded position. It reduces the hysteresis that packing friction causes and makes up for changes in spring rate and in the pressure on the stem [28]. A sticking valve still moves in steps and ignores small reversals of its signal, and tuning cannot fully remove that [27]. FCV-101 has no stiction and strokes as a 4 sec lag.

Where the Controller Runs

Loops of this kind run in field controllers with cycle times from milliseconds to minutes, under a supervisory layer that sends them setpoints [29]. LIC-101 and FIC-102 execute every 0.5 sec, the scan of a distributed control system, above a 0.05 sec physics step. In the classic three-element arrangement a summing block adds the steam flow signal and the level controller's output, the trim, and passes the sum as a remote setpoint to the feedwater flow controller [27]. The chapter delivers the same sum as the output bias of the level controller, so its integral carries only the trim, and the predictive controller of Section 4.5.3 writes to the same setpoint.

> **IN PRACTICE** **Two Independent Ways to See the Water**
>
> Boiler rules ask for a second, independent view of the drum level. United States rules for marine boilers, which modify paragraph PG-60 of Section I of the boiler and pressure vessel code, require two independent means of indicating the water level, connected directly to the drum, one of them a gauge lit from the emergency supply. The gauge glasses must stay in service whenever the boiler is steaming [30]. The glass is read with care: its water cools and becomes denser than the drum water, so the two columns balance at different heights [27].

What Is Modernizing It

Three developments reach this loop. Digital valve positioners calibrate themselves, diagnose air leakage, friction and dead band, and allow partial stroke tests on line. On the main feedwater valves of one unit they allowed the steam generator levels to be tuned to more stable control [28]. Plants that raise their power rating fit more accurate feedwater flow measurement, because the energy balance that confirms the reactor's thermal power rests on the feedwater flow, temperature and pressure [26]. Model-based control of this level has been proposed since at least 1979 [19], and predictive control has been applied to the steam generator level of a pressurized water reactor [31], the structure that Section 4.6.7 sets against three-element control.

4.9 Limitations

The application demonstrates inverse response arising from a mass and energy balance, a right-half-plane zero that moves with operating point, what each rung of the architecture ladder buys and does not buy, cascade as a loop closed around an actuator, and the drift of a differential-pressure level measurement with operating pressure. It is not a boiler analysis tool.

The circulation is one number, so the low-load region where the application is most interesting is also where the model is least complete, because a real natural-circulation boiler at very low load can circulate unstably and that failure mode is absent entirely. The fuel side is a feedforward and a proportional trim rather than a boiler master, and there is no steam line volume and no superheater, so neither firing rate limits nor the interaction between steam pressure control and level control can be studied. Blowdown, attemperation, the economizer and the drum internals are absent, and separator carryover is represented only by the high-high trip.

Three control-side limitations matter for how the results should be read. Four structures exist, with no automatic transfer to single-element control at low load, although real plants commonly make that transfer below the load at which the feedwater flow measurement stops being trustworthy. The valve has no stiction, no backlash and no positioner of its own, which matters because feedwater valve stiction is a common real cause of level cycling and is exactly the fault an inner flow loop half conceals. There is one feedwater valve and one pump characteristic, while real feedwater systems carry a startup valve, variable-speed drives, minimum-flow recirculation and pumps in parallel, all of which change what the inner loop is controlling. Perfect repeatability is itself a property of the implementation rather than of a plant, and the fixed step, the fixed scan and the fixed seed are the reason for it. Where each individual assumption breaks is set out in Section 4.3.4.

The predictive controller's model has no pressure state, so the immediate expansion of the bubbles when the pressure falls reaches it only through the filter's swell estimate. It assumes the design feedwater temperature, so the heater trip is a model error absorbed by $\hat{d}$, and it ignores the density error of Equation (4.12) and any parameter changed since the last reset. Its single 510 sec lag for the detuned inner loop predicts 2 % of a setpoint step after 10 sec where the loop delivers 15 % through its proportional part. Acting only through the flow setpoint, it fails the challenge as shipped, and its 2 sec steps cost more command travel than three-element control on three cases.

> **OPEN CHALLENGE** **One Automatic Level Controller From Startup to Full Load**
>
> In pressurized water reactors, unplanned trips caused by steam generator level control have been significant contributors to lost availability, and conventional feedwater control cannot cover the whole range from zero to full load. At low power the wrong-way response is strongest and the flow measurements are often unreliable [31]. A 2001 laboratory study of a robust model-based controller reported only preliminary results and called for more work [31]. Progress would be one controller that holds the level from startup to full load with no manual operation.

4.10 Conclusion

This chapter studied a natural-circulation steam generator as a laboratory for a nonminimum-phase plant whose zero depends on the operating point. The simulator integrates five states, the drum water volume, the sub-surface steam mass, the drum pressure, the absorbed heat and the feedwater valve stem, with the indicated level written algebraically as a water term plus a void term so that shrink and swell emerge from the balances rather than from a numerator coefficient. An open-loop feedwater step measures the zero directly. Trough depth scales with the step, from −1.47 mm to −15.21 mm across a factor of ten, while all four crossings back through the starting level fall between 38 and 39 sec, and changing the load leaves the trough near −3.7 mm while moving the crossing from 38 sec at full load to 65 sec at a quarter. On a fifteen-point load rejection the three architectures give −29.8 mm, −50.2 mm and −37.4 mm of peak deviation for 11486, 9094 and 6521 mm s of integrated error, and the same 3.75 kg/s of steam is 2.2 times worse at 40 % load while the same *relative* step is slightly easier there. A feedwater pump trip separates the architectures by a factor of 59 for the same valve travel, and a cell ranged at 28 bar while the drum runs at 40 bar holds an exact indication over water sitting 19.7 mm too high. Slowing the governor ramp from 5 to 0.25 %/s cuts the swell peak from 228.5 to 86.0 mm and does not prevent the low-low trip that follows it. A predictive controller that forecasts the voids from the measured steam flow holds the pickup from 40 to 100 % load to +122.1 mm against +167.6 mm under three-element control, and like three-element control it needs a properly tuned inner flow loop to pass the challenge. A fuller model would add valve stiction, a steam pressure loop, an automatic low-load transfer between architectures and a distributed riser model that admits circulation instability. A nonlinear predictive controller on the five-state balances could hold the two trips as explicit constraints, and an estimator over Equations (4.2) to (4.6) could reconstruct the water inventory and the transmitter density error from the level, the two flows and the pressure. The steam generator measures the price of a right-half-plane zero that moves with the load.

KEY POINTS

- A drum level is water plus the steam bubbles below the surface, so cold water lowers the reading before it raises it.
- The depth of the wrong-way dip scales with the step, while its duration grows from 38 sec at full load to 65 sec at a quarter.
- An inner feedwater flow loop covers a pump trip that no level tuning reaches, a factor of 59 on integrated error.
- A loop holding a biased measurement holds the bias, here 19.7 mm of water from a cell ranged at 28 bar.

Where else this applies. Three-element control also holds the vessel level of Chapter 6, and a zero that moves with load limits the boost converter of Volume 3, Chapter 9.

References

[1] M. V. Kothare, B. Mettler, M. Morari, P. Bendotti, and C.-M. Falinower, "Level control in the steam generator of a nuclear power plant," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 1, pp. 55–69, 2000.

[2] M. G. Na, K. C. Kwon, C. S. Ham, and Z. Bien, "A study on water level control of PWR steam generator at low power and the self-tuning of its fuzzy PI controller," *Fuzzy Sets and Systems*, vol. 74, no. 1, pp. 43–51, 1995.

[3] D. Lindsley, *Power-Plant Control and Instrumentation: The Control of Boilers and HRSG Systems*. London, UK: Institution of Electrical Engineers, 2000.

[4] S. G. Dukelow, *The Control of Boilers*, 2nd ed. Research Triangle Park, NC: Instrument Society of America, 1991.

[5] 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.

[6] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. Chichester, UK: John Wiley & Sons, 2005.

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

[8] W. L. Luyben, "Identification and tuning of integrating processes with deadtime and inverse response," *Industrial & Engineering Chemistry Research*, vol. 42, no. 13, pp. 3030–3035, 2003.

[9] K. L. Chien, E. I. Ergin, C. Ling, and A. Lee, "Dynamic analysis of a boiler," *Transactions of the ASME*, vol. 80, no. 8, pp. 1809–1818, 1958.

[10] K. J. Åström and K. Eklund, "A simplified non-linear model of a drum boiler-turbine unit," *International Journal of Control*, vol. 16, no. 1, pp. 145–169, 1972.

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

[12] H. Kim and S. Choi, "A model on water level dynamics in natural circulation drum-type boilers," *International Communications in Heat and Mass Transfer*, vol. 32, no. 6, pp. 786–796, 2005.

[13] F. G. Shinskey, *Process Control Systems: Application, Design, and Tuning*, 4th ed. New York, NY: McGraw-Hill, 1996.

[14] 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.

[15] K. L. Poon, "Effect of rapid steam take-off on natural circulation and water level in boilers," *Journal of Engineering for Power*, vol. 97, no. 4, pp. 645–654, 1975.

[16] C. Maffezzoni, "Boiler-turbine dynamics in power-plant control," *Control Engineering Practice*, vol. 5, no. 3, pp. 301–312, 1997.

[17] R. J. M. Welsh, "PWR steam generator water level measurement systems," *Measurement and Control*, vol. 26, no. 5, pp. 151–155, 1993.

[18] M. A. Gonzalez-Salazar, T. Kirsten, and L. Prchlik, "Review of the operational flexibility and emissions of gas- and coal-fired power plants in a future with growing renewables," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 1497–1513, 2018.

[19] J. J. Feeley, "Natural circulation steam generator model for optimal steam generator water level control," in *American Nuclear Society Meeting, San Francisco, CA, November 1979*. Idaho National Engineering Laboratory, Jun. 1979, report CONF-791103-30, OSTI 5962048.

[20] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[21] 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.

[22] 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.

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

[24] 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.

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

[26] ——, *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.

[27] 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.

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

[29] 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.

[30] U.S. Coast Guard, "Water-level indicators, water columns, gauge-glass connections, gauge cocks, and pressure gauges (modifies PG-60)," Code of Federal Regulations, Title 46, Section 52.01-110, revised as of October 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[31] A. M. H. Basher and J. March-Leuba, "Development of a robust model-based water level controller for U-tube steam generator," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2001/166, Sep. 2001.

4.11 Exercises

4.1 ○ Concept check · 1 min

In MANUAL at full load the feedwater valve is opened from 70 to 78 %, and the feedwater rises above the steam flow. The drum level first falls. Which explanation is right?

(a) The valve strokes for 4 sec, and the feedwater flow dips while the stem is moving

(b) The level transmitter lag of 3 sec makes the reading run behind the true level

(c) The subcooled feedwater condenses steam below the surface, and the volume lost exceeds the water gained

(d) The extra feedwater raises the drum pressure, which compresses the bubbles below the surface

4.2 ○ Concept check · 2 min

Open-loop feedwater steps of 10, 25, 50 and 100 % are made at full load from the same equilibrium. What does Table 4.3 show about the inverse response?

(a) The trough and the crossing time both grow with the step, so a smaller move gives a faster recovery

(b) The trough is the same for every step, and the crossing time grows with the step

(c) The trough grows with the step, and the crossing time falls because a larger flow refills the drum sooner

(d) The trough grows in proportion to the step, and the crossing time stays near the same value

4.3 Concept check · 2 min

On the load rejection from 100 to 85 %, two-element control at $K_c = 1.6$ dips to −50.2 mm, deeper than the −29.8 mm of single-element control, and has no overshoot. Why is the dip deeper?

(a) The feedforward cuts the feedwater to the new steam flow, so no extra water offsets the void collapse

(b) The feedforward carries the noise of the steam flow transmitter into the valve

(c) The valve-position feedforward is most in error at full load, where it asks for too little valve

(d) The feedforward doubles the effective gain of the level controller, which then overshoots downward

4.4 Theory and calculation · 6 min

Use the property fits of Equation (4.11) at the design pressure of 40 bar.

(a) Compute the subcooling $h_w - h_{fw}$ of 160 °C feedwater and the mass of steam that one kilogram of it condenses on its way to saturation. *Checkpoint:* 408.3 kJ/kg and 0.238 kg.

(b) Compute the ratio ρ_w/ρ_s, and from it the ratio of the steam volume lost to the water volume gained for each kilogram of feedwater. *Checkpoint:* 39.8 and 9.5.

(c) A feedwater heater trip drops the feedwater to 110 °C. Repeat part (a) and state how much stronger the shrink is per kilogram. *Checkpoint:* 622.4 kJ/kg, 0.363 kg, 52 % stronger.

4.5 Theory and calculation · 6 min

At any steady load the drum of SG-101 sits at its design pressure of 40 bar with feedwater equal to steam. Take $\tau_{res,0} = 12$ sec, $q_{s,0} = 25\,\mathrm{kg/s}$, $k_{void} = 0.526$, $A_d = 20\,\mathrm{m^2}$ and $\rho_s = 20.09\,\mathrm{kg/m^3}$, and let the firing fraction Q/Q_0 equal the load fraction.

(a) Use Equations (4.4) and (4.5) to show that the steady sub-surface steam is $m_{sd} = \tau_{res}\, q_s$, and compute τ_{res} at 85, 40 and 25 % load. *Checkpoint:* 12.7, 16.2 and 18.1 sec.

(b) Compute the void contribution to the level, the second term of Equation (4.1), at 85, 40 and 25 % load. *Checkpoint:* −38, −181 and −244 mm.

(c) Compare the change in void contribution for the load rejection from 100 to 85 % with that from 40 to 25 %. Both remove 3.75 kg/s of steam. *Checkpoint:* 38 mm against 63 mm.

4.6 Theory and calculation · 7 min

Level transmitter LT-101 has taps $H = 1.2$ m apart and reports level by Equation (4.12). In scenario 12 it is ranged at $p_{cal} = 28$ bar while the drum runs at 40 bar.

(a) Compute $\rho_w - \rho_s$ at 28 and at 40 bar from Equation (4.11). *Checkpoint:* 816.5 and 778.5 $\mathrm{kg/m^3}$.

(b) The controller holds the indication $\hat{L}$ at zero. Invert Equation (4.12) for the true level L. *Checkpoint:* +19.7 mm.

(c) A correctly ranged cell, $p_{cal} = 40$ bar, sees the drum pressure rise to 42.4 bar with the true level at zero. Compute its error and the sensitivity per bar. *Checkpoint:* −3.5 mm, about 1.5 mm/bar.

4.7 App challenge · 7 min

Load scenario 1, Drive the valve by hand (MANUAL), and change no setting. The feedwater valve steps from 70 to 78 % at 60 sec. Run to 160 sec at the default speed and seed, and export the CSV file.

(a) In column `level_m` (true level, in meters), find the lowest level after the step and the time after the step at which it occurs. *Checkpoint:* −3.95 mm, within 0.05 mm, 19.4 sec after the step, within 0.5 sec.

(b) In the row at 79.4 sec, read the water and void contributions, columns `level_water_m` and `level_void_m`. *Checkpoint:* +7.16 mm and −11.11 mm, each within 0.05 mm.

(c) Find the time after the step at which `level_m` passes back through zero, and compare it with the crossing at a quarter load in Table 4.3. *Checkpoint:* 38.1 sec, within 0.2 sec, against 65 sec.

4.8 App challenge · 9 min

Run the fifteen-point load rejection on the three architectures, each scenario as shipped, at the default seed and a speed of 10×: scenario 3, Single-element on a load rejection, and scenario 5, Two-element: the steam-flow feedforward, each to 800 sec, and scenario 7, Three-element control, to 860 sec. Each reading is then 710 sec after the load change.

(a) For each run read the IAE tile of the scorecard, in m s, and the Valve travel tile. *Checkpoint:* 11.486, 9.094 and 6.521 m s, each within 0.01 m s, for 36, 39 and 21 %.

(b) From each exported file, find the lowest true level, column `level_m`, after the load change. *Checkpoint:* −29.8, −50.2 and −37.4 mm, each within 0.1 mm.

(c) Explain why three-element control cannot bring the trough below about 38 mm whatever its gain.

4.9 App challenge · 8 min

Load scenario 10, Feedwater pump trip, on two-element control. The pump trips at 90 sec and the discharge pressure falls from 55 to 43 bar. Use the default seed and a speed of 10×, and run each case to 990 sec, 900 sec after the trip.

(a) Run the scenario as shipped. From the file, read the feedwater flow `fw_flow_kgps` and the valve stem `fw_valve_pct` at 100 sec, and the lowest true level `level_m` and its time. Read the IAE tile at 990 sec. *Checkpoint:* 10.4 kg/s at 69 % open; −149 mm at 340 sec; IAE 86.51 m s, within 0.05 m s.

(b) Reload the scenario, set Control architecture to Three-element before Run, and repeat. Compute the ratio of the two IAE readings. *Checkpoint:* −7.6 mm, IAE 1.459 m s, within 0.01 m s; a ratio of 59.

4.10 Application area · 3 min

After the pump trip of scenario 10 the discharge pressure is 43 bar against a drum at 40 bar. The valve passes $q_{f,\max} = 58.6\,\mathrm{kg/s}$ wide open at the design differential of 15 bar. What is the most feedwater the valve can now deliver, and what does that mean for a level loop at full load? Name one practical way to keep authority after such a trip.

4.11 Application area · 3 min

A unit was commissioned on two-element control as a base-load machine and is now cycled down to 40 % load every night. Give two results of this chapter that argue for a different architecture at low load, and one behavior of a real boiler at very low load that the model cannot show.

4.12 Review the agent · 5 min

An AI assistant reviews scenario 11, in which the turbine picks up from 25 to 120 % at 5 %/s on single-element control, and replies as follows. “The danger is the swell. The indicated level peaks at +228.5 mm, only 21.5 mm below the high-high trip. Limit the governor to 0.5 %/s and the peak falls to +138.3 mm, so the drum rides through the pickup with no trip and no change to the controller.” Find the flaw in the proposal and say how it could be checked.