INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 5

Nuclear Steam Dump Simulator for Feedforward and Override on a Load Rejection

Pages 112–141

Edited by John D. Hedengren

How to cite this chapter. "Nuclear Steam Dump Simulator for Feedforward and Override on a Load Rejection," 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. 5, pp. 112–141.

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

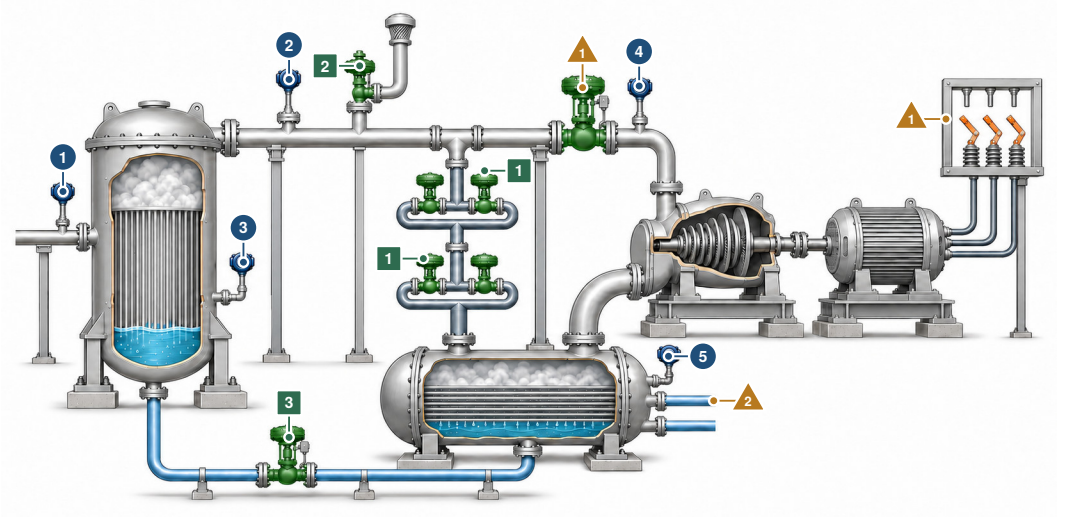
The app for this chapter
https://apmonitor.com/apps/turbine-trip/

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

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

CHAPTER 5

Nuclear Steam Dump Simulator for Feedforward and Override on a Load Rejection



The controller moves 1 steam dump valves to the condenser 2 relief valve to the air 3 feedwater valve

The controller reads 1 average coolant temperature 2 steam pressure 3 steam generator water level 4 steam pressure inside the turbine 5 condenser pressure

What disturbs it 1 breaker opens, turbine valve shuts 2 cooling water temperature

When the breaker opens and the turbine valve shuts, the steam has to go through the dump valves into the condenser, and the controller learns of it first from the pressure inside the turbine. Illustration, not to scale.

IN PLAIN TERMS Where the Steam Goes When the Grid Is Lost

The main idea. A nuclear plant makes steam that spins a turbine, which drives a generator. When the link to the power grid suddenly opens, the turbine valves shut, but the reactor keeps making steam with nowhere to go.

The challenge. The pressure reaches the safety valves in under fifteen seconds, too fast for a thermometer in the water to warn in time. So valves that send the steam to a cooler that turns it back into water must open at once, on a signal from the turbine.

Why it matters. Each time the safety valves open, clean steam is lost to the air and the valves must be tested.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of Section 5.1, the history box of Section 5.2, the try-it boxes of Sections 5.4 and 5.6, the fact box of Section 5.6, 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 5.6, Section 5.7, Section 5.8 and the open challenge

The control question When the generator breaker opens, the steam a pressurized water reactor is still making has nowhere to go, and the secondary reaches its first safety valve in under fifteen seconds. Can a temperature loop on the steam dump be tuned to hold that margin, or does the dump need a signal from the turbine and overrides on its constraints?

Learning objectives

- Compute the heating rate of a shut-in secondary from its energy balance and compare it with the stroke time and the sensor lag.
- Measure the process gain of the dump and explain why a full-load program step saturates any reasonable gain.
- Separate what a load-programmed reference and an explicit lead-lag feedforward each contribute to the anticipation.
- Explain how low-selected overrides cap a shared actuator, and why a controller in front of an interlock beats the interlock.
- Judge when a reactor runback, the reliefs or a larger dump is the right remedy, including for failures that nothing measures.

Prerequisites Sections 1.4.2 (prompt jump and reactivity feedback), 1.4.4 (shrink and swell), 1.4.5 (the clock to a limit), 1.6.2 (velocity form, back-calculation) and 1.6.4 (scenarios and data). Chapter 2 for rod control and the temperature program.

The app https://apmonitor.com/apps/turbine-trip/

See also Volume 3, Chapter 2 for a structure that rejects only the disturbance entering its measurement path, and Volume 3, Chapter 6 for feedforward from an upstream measurement with a slow trim. Volume 4, Chapter 2 for a protective override on a fast transient, and Volume 4, Chapter 3 for a selector over a governor and its limiters.

5.1 Introduction

Almost every disturbance a first control course presents arrives slowly enough to be measured. A feed temperature drifts, a load steps, a composition wanders, and in every case the controlled variable moves, the transmitter reports it, and the loop responds before anything is lost. The lesson a student draws is that a loop which rejects a disturbance badly needs a higher gain or a faster measurement. A generator load rejection on a large steam plant refutes it. The breaker opens, the turbine overspeed protection controller shuts the governor valves inside a second, and about 1600 kilograms per second of steam that was going through a machine is going nowhere, while the reactor making it has no idea anything has changed. The controlled variable does not move for several seconds and the transmitter that reads it has a response time of four.

> **DID YOU KNOW? Heat That Outlasts the Trip**
>
> When a reactor trips, fission stops within a second, but the decaying fission products keep heating the coolant. At Three Mile Island in 1979 that heat was just 6 percent of the fission heat, still large enough that it had to be removed to keep the core from overheating [1].

The misconception this chapter corrects is that one. Feedback performance against a disturbance is bounded by the sensitivity function of the loop, and that bound is set by the phase and the delay rather than by the gain, so a disturbance whose spectrum lies above the achievable bandwidth cannot be attenuated at any gain at all [2]. A block diagram states that result and a transfer function proves it, but neither shows what it costs, because the cost here is a discrete list of consequences a plant either avoids or does not. Measured on the shipped constants, seven times the proportional gain on the dump temperature loop with a fivefold shorter integral time moves the peak steam pressure from 83.37 to 81.70 bar, and an average-temperature RTD made eight times faster than any real instrument gives 82.92 bar with 454 kg of steam still through the safety valves. The same controller handed one additional input reaches 79.74 bar and lifts nothing.

The class of problem is general and is not confined to steam plants. A disturbance arrives fully formed, faster than any measurement of its effect can develop, so the only remedy is anticipation from a signal upstream of the effect, and the constraints that were comfortable in normal operation close in during the transient and have to be handled by something other than the loop that caused them. Tuning rules for feedforward from a measured disturbance exist for exactly this situation and are written in terms of the disturbance path rather than the process path [3], and selective or override structures are the standard answer to a constraint that binds only transiently [4]. Compressor anti-surge is the same shape, because a driver trip collapses the flow through the machine faster than the surge line can be approached gracefully and the anti-surge valve must be opened from a signal that is not the surge margin itself [5], and Volume 4, Chapter 2 treats that machine with a protective override. Generator load rejection on a fossil unit or a gas turbine is the same shape by construction, and so is a batch reactor whose exotherm runs away faster than a jacket temperature can report it. In every case the plant carries a fast anticipatory path, a slow trim, and a set of overrides that speak only when a constraint moves.

Simulators of nuclear steam supply systems are ordinarily built to answer a transient-analysis question rather than a control-architecture one [6, 7], and those built for control work usually stop at the reactor and the temperature program [8]. However, a plant model carrying one control structure cannot separate what the physics decided from what the engineer decided, and on a load rejection that separation is the whole question. This chapter puts a pressurized water reactor secondary side under a generator load rejection in front of a control engineer as an experiment rather than a calculation, and prices each rung of a four-step control ladder in bar of peak steam pressure and in tonnes of steam through the safety valves. The plant reduces published steam generator and station dynamics to teaching scale [8, 9], and steam dump control is a structure Westinghouse units have carried for decades [10]. Four named architectures share one actuator, one sample period, one set of limits and one disturbance, so that the difference between them is wiring and nothing else, and the feedback only structure generates its own reference from a measurement downstream of the event. The scenarios form a sequence in which the failures are structural rather than arithmetic and two of them are out of reach of tuning altogether.

Section 5.2 traces steam bypass from a startup device to a licensed transient and names the control structures the vendors fixed. Section 5.3 gives the balances, the constants and the integration scheme in the detail a reader needs to rebuild the plant. Section 5.4 measures the clock with no controller and the process gain of the dump. Section 5.5 sets out the temperature program, the anticipation, the dump demand with its two overrides, and the reliefs, rods and feedwater. Section 5.6 compares the four architectures on one rejection, shows that tuning does not close the gap, locates the anticipation, and follows the constraint that moves with the weather, the two dump modes and the failures nothing reports. Section 5.7 sets the loop in its plant hardware, and Sections 5.8 to 5.10 state the limitations, discuss what the results mean for a design and summarize the chapter.

5.2 Turbine Bypass and Steam Dump Control

Steam bypass around a turbine is older than nuclear power and began as a startup device rather than a transient one. A boiler brought up to pressure has to put its steam somewhere while the turbine is still on the turning gear, and the earliest arrangements were manual valves to the condenser or to atmosphere sized for whatever the boiler made during warmup. What turned bypass into a control problem was the arrival of units large enough that losing the grid was a plant event rather than an inconvenience. A drum boiler can be fired down in a minute or two. A reactor cannot be switched off, its decay heat continues for days, and the buffer between it and the turbine is only the inventory of the steam generators, so the disposal path has to be sized against the transient rather than against the startup.

The sizing question settled into a split that survives in current practice. Ordinary pressurized water units carry a steam dump of roughly 40 % of rated steam flow, enough to hold the plant through a

turbine trip and bring the average coolant temperature down to its no-load program without lifting the main steam safety valves, and units specified to shed the grid and remain at power carry up to 85 % [10]. Vendor treatments of bypass valve design frame the choice in the same terms and add that a plant expected to cycle needs the larger valve for flexibility rather than for protection [11]. The constraint deciding whether the larger valve is worth buying sits downstream of it, because a condenser sized for turbine exhaust is loaded far beyond its design duty by a large bypass, and guidelines for bypass operation are written around the condenser rather than the valve [12]. Section 5.6 measures that constraint on the shipped model and reaches the same conclusion.

Licensing turned the transient into a required calculation. Loss of external load and turbine trip are grouped together in the standard review plan as decrease-in-heat-removal events, and the acceptance criteria are limits on primary and secondary pressure, on safety valve actuations and on the departure-from-nucleate-boiling ratio, with the applicant required to show the consequences both with and without credit for the steam dump and the reactor control system [13]. That requirement is why a load rejection analysis exists for any unit and why the steam dump is documented as a system with arming permissives and interlocks rather than as a valve. Analyses of real full load rejections followed the requirement rather than preceding it, and the early ones were code-to-plant comparisons on units that had actually shed the grid [7]. Modern work on the same transient is largely about coordination, because a plant asked to ride out a rejection has to move the dump, the rods and the feedwater together rather than sequentially [6], and the pressure to do so comes from grid duty rather than from safety [14].

> **HISTORY** **Riding Out a Turbine Trip Below Half Power (1979)**
>
> In December 1979 a national laboratory working for the United States regulator reviewed a request from the owner of one pressurized water reactor. Its protection system tripped the reactor whenever the turbine tripped above 10 percent power. Because both the reactor and the turbine would accept a 50 percent load rejection, the owner asked to trip the reactor on a turbine trip only at or above 50 percent power. A new permissive, P-9, formed from the power range neutron flux channels, would block the reactor trip below that power. The turbine trip would then ride out as a load rejection, with the reactor power dissipated by a steam dump able to bypass up to 50 percent of rated steam. The review noted that the safety analysis took no credit for the steam dump and recommended approval [15].

The control structure itself was fixed by the vendors and is unusually well documented. A Westinghouse steam dump control system carries two operator selected modes. The load rejection controller works on average coolant temperature against a program derived from turbine first-stage impulse pressure, so it acts continuously and in proportion to the load actually lost. The turbine trip controller works against the no-load program instead, is armed by the P-4 turbine trip permissive rather than by a continuous measurement, and is blocked entirely until that permissive appears. Condenser back pressure and steam line pressure interlock the dump valves shut, and the rod control system carries its own runback logic driven from the same load rejection signal [10]. Every one of those is a control-structure decision rather than a tuning decision, and every one is reproduced in the app of this chapter.

The theory those structures rest on was developed in the process industries under different names. A constraint controller that publishes a limit and is low-selected against a working demand is a selective or override control, and the treatment establishing that such a controller caps a demand rather than adding to it is a process control one [4]. Handover between a selected and an unselected controller is the same problem as anti-windup, and the conditioning technique that solves both by back-calculating the integral against the limit actually applied is the standard method [16]. Systematic procedures for designing a set of selectors against moving active constraints are more recent and treat the selector as a control element in its own right [17]. A steam dump controller is a worked example of all three, which is why it belongs in front of a process control student and not only a nuclear one.

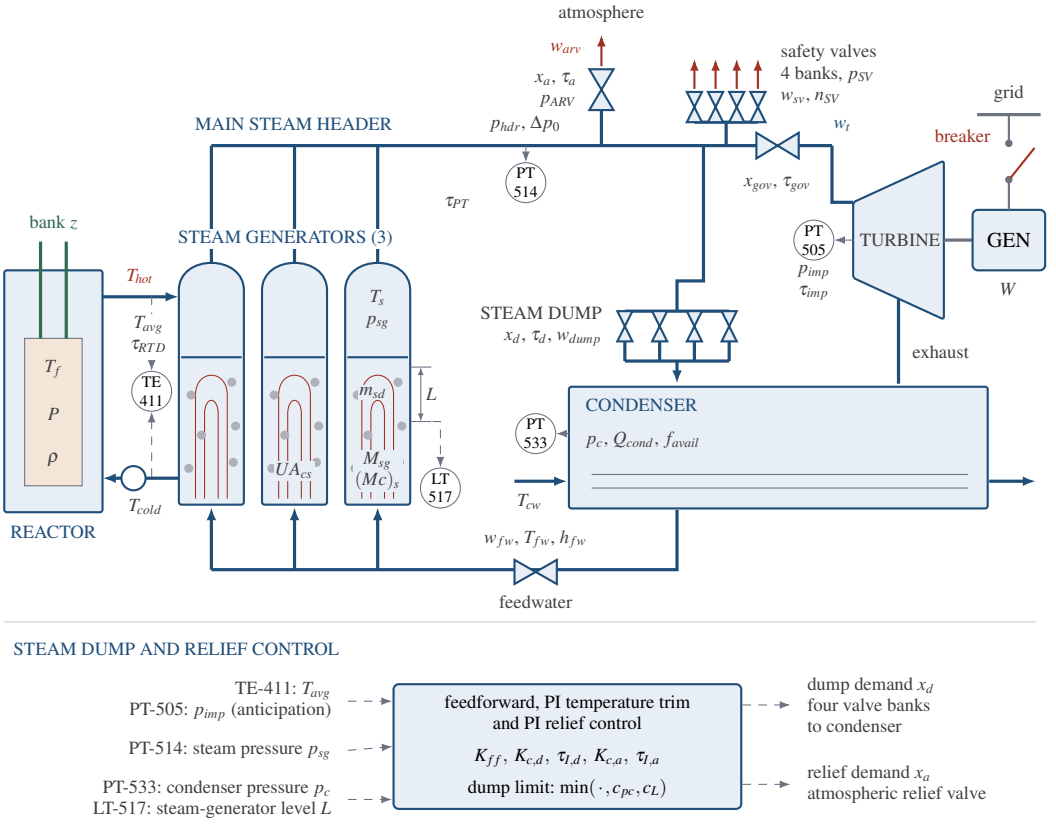


Figure 5.1 Secondary side of unit U-1 during a generator load rejection. The core delivers P across three steam generator tube bundles at conductance UA_{cs}, boiling the inventory M_{sg} at saturation temperature T_s and pressure p_{sg} with steam mass m_{sd} circulating below the water surface. Steam leaves through the header at p_{hdr} into the governor valve x_{gov}, through the four dump banks x_d to the condenser at back pressure p_c, and to atmosphere through the reliefs x_a and the four safety valve banks set at p_{SV}. The dump controller reads turbine impulse pressure p_{imp} as its anticipatory signal, average coolant temperature T_{avg} through RTD TE-411, and condenser pressure and generator level as the two caps c_{pc} and c_L that low-select against its own demand. Dashed lines are signals and the breaker is drawn open.

5.3 Plant Model

Unit U-1 is a Westinghouse-type three-loop pressurized water reactor rated at 2900 MW thermal and 940 MW electrical gross. Three steam generators pass the whole core duty across their U-tube bundles, boiling 1595.0 kg/s of saturated steam at 61.48 bar into a main steam header that runs to one turbine. Four banks of steam dump valves discharge from that header to the condenser, atmospheric reliefs and four banks of main steam safety valves discharge from it to atmosphere, and feedwater returns from the condenser hotwell at 226 °C. That plant is drawn with its instrument tags in Figure 5.1, which also carries every symbol the equations use. The generator breaker is the disturbance and the steam dump valve position is the manipulated variable.

Table 5.1 separates what the loops control from what they move, what upsets them and what the integrator carries. Nineteen states are integrated, and the working units are those of nuclear steam plant practice, power in MW, heat capacity in MW s/°C, temperature in °C, pressure in bar except the condenser end which is in mbar, mass in kg, flow in kg/s, enthalpy in kJ/kg, reactivity in

Table 5.1 Classification of Variables. The Nominal Column Is the Full-Power Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
T_{avg} (TE-411)	CV, measured	310.00 °C	291.7–310.0 °C SP
p_{sg} (PT-514)	CV, measured	61.48 bar	68–84 bar SP
L (LT-517)	CV, measured	50.0 %	35–65 % SP
x_d	MV	0 %	0–100 %
x_a	MV	0 %	0–100 %
v_z rod velocity	MV	0 %/s	±0.5 %/s
w_{fw}	MV	100 %	0–130 %
breaker	DV, measured	closed	open, closed
turbine demand	DV, measured	100 %	0–105 %
T_{cw}	DV, measured	21 °C	4–34 °C
dump banks avail.	DV, unmeasured	100 %	0–100 %
x_a stuck open	DV, unmeasured	0 %	0–100 %
$C_1 \dots C_6$	state	—	≥ 0
D_1, D_2	state	0.028, 0.040	≥ 0
T_f	state	610.0 °C	—
T_{avg}	state	310.00 °C	—
T_s	state	277.01 °C	180–330 °C
M_{sg}	state	135 000 kg	$\geq$ 1000 kg
m_{sd}	state	3190 kg	≥ 0
z	state	68.0 %	0–100 %
x_{gov}	state	0.90	0–1.05
x_d, x_a	state	0 %	0–100 %
p_c	state	53.1 mbar	10–1200 mbar
w_{fw}	state	1595.0 kg/s	≥ 0

pcm ($10^{-5}\,\Delta k/k$) and time in sec. Every controller works in percent of valve travel and percent of transmitter span.

5.3.1 Saturated Properties

The secondary visits 40 to 90 bar and nothing else, so the saturation line is carried as low-order fits rather than as a property library, as shown in Equation (5.1).

$$\begin{aligned} T_{sat}(p) &= 185.0 + 1.88912\,p - 0.006384\,p^2 \\ h_{fg}(p) &= 1930 - 6.0\,p, \qquad h_g(p) = 2831.7 - 0.8\,p \\ h_f(p) &= 1213.7 + 5.2\,(p - 60) \\ \rho_s(p) &= 0.4956 + 0.45663\,p + 8.306\times 10^{-4}\,p^2 \end{aligned} \tag{5.1}$$

where p is in bar, temperature in °C, enthalpy in kJ/kg and steam density in kg/m^3. Saturation temperature is good to about 0.5 °C and latent heat to about 1 % over that band, and $p_{sat}(T)$ is the positive root of the first line rather than a second fit, so the two are exact inverses. The fits carry no pressure clamp, only the saturation temperature state is held between 180 and 330 °C as a numerical guard, and no scenario leaves the fitted band. Feedwater enthalpy is $h_{fw} = 4.19\,T_{fw} + 3.4\times 10^{-4} T_{fw}^2$, which returns 964.3 kJ/kg at 226 °C, and the condenser end uses the Antoine expression $p_c = 10^{8.07131 - 1730.63/(233.426+T)}/750.062$ in bar for T clamped to 5 to 99 °C.

5.3.2 Core Neutronics and Thermal Model

The prompt neutron mode of a light water reactor is about 3 ms and the transient of interest runs for 300 s, so the prompt jump approximation of Section 1.4.2 is applied and only the six delayed

precursor groups are integrated [18], as shown in Equations (5.2) and (5.3).

$$\frac{\mathrm{d}C_i}{\mathrm{d}t} = \beta_i n - \lambda_i C_i, \qquad i = 1 \ldots 6 \tag{5.2}$$

$$n = \frac{\sum_i \lambda_i C_i}{\max(\beta - \rho,\ 0.05\,\beta)} \tag{5.3}$$

where n is fission power as a fraction of rated, C_i is the precursor concentration of group i in the same units, β_i and λ_i are the thermal-fission delayed-neutron fractions and decay constants of U-235, and ρ is net reactivity in $\Delta k/k$. The six fractions are 2.15, 14.24, 12.74, 25.68, 7.48 and 2.73×10^{-4} summing to $\beta = 0.006502$, against decay constants of 0.0124, 0.0305, 0.111, 0.301, 1.14 and 3.01/s [18]. The denominator floor in (5.3) keeps the expression finite on a deep negative excursion. Reactivity is the sum of bank worth and two temperature feedbacks, as shown in Equations (5.4) and (5.5).

$$\rho = \rho_z(z) - \rho_z(z_{ref}) + \alpha_F \Delta T_f + \alpha_M \Delta T_{avg} - \rho_{scram} \tag{5.4}$$

$$\rho_z(z) = W_{tot}\left[f - \frac{\sin 2\pi f}{2\pi}\right], \qquad f = \frac{z}{100} \tag{5.5}$$

where $\Delta T_f = T_f - T_{f,0}$ and $\Delta T_{avg} = T_{avg} - T_{avg,0}$ are departures from the design point, z is control bank position in percent withdrawn, W_{tot} is total bank worth, α_F is the Doppler coefficient, α_M is the moderator temperature coefficient and ρ_{scram} ramps to the scram rod worth over t_{scram} after a trip. Equation (5.5) is the integral S curve, zero fully inserted and W_{tot} fully withdrawn, and its derivative $\frac{\mathrm{d}\rho_z}{\mathrm{d}z} = (W_{tot}/100)(1 - \cos 2\pi f)$ is 42.8 pcm/% at the full-power position $z_{ref} = 68.0\,\%$, where bank worth is 2472 pcm of the 3000 pcm total. That derivative sets the authority of the runback, because 0.5 %/s of bank travel is 21.4 pcm/s of reactivity.

Decay heat is two lumped groups and total core power is their sum with the prompt fission term, as shown in Equations (5.6) and (5.7).

$$\frac{\mathrm{d}D_j}{\mathrm{d}t} = \frac{f_{D,j}\, n - D_j}{\tau_{D,j}}, \qquad j = 1, 2 \tag{5.6}$$

$$P = P_0\left[(1 - f_{D,1} - f_{D,2})\, n + D_1 + D_2\right] \tag{5.7}$$

where $f_{D,j}$ and $\tau_{D,j}$ are the group fraction and time constant, and the weighting in (5.7) makes the equilibrium at $n = 1$ exactly P_0, so the two groups redistribute power rather than adding it. Two thermal nodes carry the heat to the secondary, as shown in Equations (5.8) and (5.9).

$$m_f c_f \frac{\mathrm{d}T_f}{\mathrm{d}t} = f_P P - UA_{fc}\,(T_f - T_{avg}) \tag{5.8}$$

$$\begin{aligned} m_c c_c \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} &= (1 - f_P)P + UA_{fc}(T_f - T_{avg}) \\ &\quad - UA_{cs}(T_{avg} - T_s) \end{aligned} \tag{5.9}$$

where T_f is average fuel temperature, T_{avg} is average coolant temperature, f_P is the fission power fraction that lands in the fuel, and UA_{fc} and UA_{cs} are the fuel-to-coolant and primary-to-secondary conductances. Both sides carry MW. Loop temperatures follow from the core rise, $T_{hot,cold} = T_{avg} \pm \frac{1}{2}\Delta T_{core} P/P_0$, which give 326.50 and 293.50 °C at full power. The primary is incompressible, so there is no pressurizer state.

5.3.3 The Secondary Energy Balance

Equation (5.10) is the whole transient in one line, and every measured number in this chapter follows from it. It is the plant's own form of the clock of Equation (1.8).

$$(Mc)_s \frac{\mathrm{d}T_s}{\mathrm{d}t} = UA_{cs}(T_{avg} - T_s) - \frac{w_{tot} h_{fg}}{1000} - \frac{w_{fw}[h_f - h_{fw}]}{1000} \tag{5.10}$$

where $(Mc)_s$ is the lumped secondary heat capacity of the saturated water, the tube bundles and the shells of all three generators [9, 19] in MW s/°C, w_{tot} is the sum of every steam flow leaving the generators in kg/s, and the factors of 1000 convert kJ/s to MW. The first term is 2900 MW at full power, the second is 1595.0 kg/s at 1561 kJ/kg of latent heat and the third is the 410 MW of sensible heat the feedwater absorbs on its way to saturation, and the three balance to zero. Secondary pressure follows the state through the property fit, $p_{sg} = p_{sat}(T_s)$, and the slope of that fit at the design point is 0.906 bar/°C.

Equation (5.10) is written in saturation temperature because the driving force for the tube bundle is a temperature difference and because writing it this way makes the arithmetic legible. With w_{tot} driven to zero and everything else at the design point the right-hand side is $(2900-410)/1200 = 2.075\,^{\circ}\mathrm{C/s}$, and one degree is 0.906 bar. The rest of the chapter is a consequence of that number.

Pause. Try Exercise 5.1 before reading on.

5.3.4 The Steam Header and Its Valves

Every valve on the header passes flow proportional to header pressure, and the header pressure falls below generator pressure as the square of the total flow. Substituting one relation into the other gives a quadratic in the flow ratio $\xi = w/w_0$, solved in closed form as shown in Equations (5.11) and (5.12).

$$w = K w_0 \frac{p_{hdr}}{p_{hdr,0}}, \qquad p_{hdr} = p_{sg} - \Delta p_0\, \xi^2 \tag{5.11}$$

$$\xi = \frac{-1+\sqrt{1+4ac}}{2a}, \qquad a = \frac{\Delta p_0 K}{p_{hdr,0}}, \quad c = \frac{K p_{sg}}{p_{hdr,0}} \tag{5.12}$$

where $K = K_t + K_d$ is the total flow coefficient of the valves on the header as a multiple of the design value, w_0 is rated steam flow, Δp_0 is the main steam line drop at rated flow and $p_{hdr,0} = p_{sg,0} - \Delta p_0 = 59.98$ bar. Solving (5.11) in closed form matters for more than tidiness. A fourth-order Runge-Kutta stage set asks for the state derivatives four separate times inside one step, and an iterative header solve buried in them would answer each of the four slightly differently.

The two coefficients and the flow split are given by Equation (5.13), and the reliefs and safety valves discharge from the steam lines upstream of the header restriction, so they see generator pressure directly, as shown in Equation (5.14).

$$K_t = \frac{x_{gov}}{x_{gov,0}}, \qquad K_d = c_{av}\frac{r_d\, x_d}{100}$$

$$w_t = w_0 \xi \frac{K_t}{K_t+K_d}, \qquad w_{dump} = w_0 \xi \frac{K_d}{K_t+K_d} \tag{5.13}$$

$$w_{arv} = w_0\, r_a \frac{x_a}{100}\frac{p_{sg}}{p_{sg,0}}, \qquad w_{sv} = w_0\, r_{sv}\, n_{SV}\frac{p_{sg}}{p_{sg,0}} \tag{5.14}$$

where r_d, r_a and r_{sv} are the dump, relief and per-bank safety valve capacities as fractions of rated steam flow, $x_{gov,0} = 0.90$ is the governor position at full load, and c_{av} is the product of the fraction of dump banks in service and the condenser interlock of Equation (5.17). At the design capacities the full dump bank passes 638 kg/s, the reliefs 159 kg/s and each safety valve bank 478 kg/s. The safety valves are four banks staggered by Δp_{SV}, each ramping linearly to full lift one accumulation above its own setpoint, as shown in Equation (5.15).

$$n_{SV} = \sum_{k=0}^{3} \mathrm{clip}\left(\frac{p_{sg} - (p_{SV} + k\,\Delta p_{SV})}{p_{acc}},\ 0,\ 1\right) \tag{5.15}$$

where $p_{SV} = 82.7$ bar is the lowest bank and p_{acc} is the accumulation to full lift. Equation (5.15) carries no latch and no blowdown, which keeps the right-hand side continuous for the integrator at the cost of understating the discharge immediately after a real valve pops and overstating it during reseating.

5.3.5 The Condenser and the Interlock

The condenser is the steam dump's only customer, and it is modeled as a back pressure correlation rather than as a heat exchanger, as shown in Equation (5.16).

$$p_c^* = p_{sat}\left(T_{cw} + \frac{Q}{\dot{m}_{cw}c_p} + TTD_0\left(\frac{Q}{Q_0}\right)^{0.4} + \frac{TTD_d\, Q_{dump}}{Q_0}\right) \tag{5.16}$$

where $Q = Q_0 w_t / w_0 + Q_{dump}$ is the total duty, $Q_{dump} = w_{dump}[h_g(p_{sg,0}) - 140]/1000$ is the dump contribution with condensate leaving at 140 kJ/kg, T_{cw} is the circulating water inlet temperature, $\dot{m}_{cw}c_p$ is its heat capacity rate, TTD_0 is the terminal temperature difference at design duty and TTD_d is an extra terminal difference charged against dump steam. The state follows that target through a first-order lag, $\tau_{cond}\frac{\mathrm{d}p_c}{\mathrm{d}t} = p_c^* - p_c$.

The exponent and the extra terminal difference exist because turbine exhaust and dump steam are not the same commodity. Exhaust has already given up its work and arrives spread across the tube bundle, while dump steam arrives at full enthalpy through a handful of nozzles [12]. Measured on the shipped constants, 40 % of rated flow through the dump is 1686 MW into a condenser rated for 1940 MW, which is 87 % of design duty from a plant making no power at all, and 85 % of rated flow is 3583 MW, or 1.85 times design. The plant protects the condenser with an interlock rather than a controller, and the interlock tightens as it is approached, as shown in Equation (5.17).

$$f_{avail} = \mathrm{clip}\left(\frac{p_{block} - p_c}{p_{block} - p_{derate}},\ 0,\ 1\right) \tag{5.17}$$

where $p_{derate} = 115$ mbar is the back pressure at which the dump begins to be cut back and $p_{block} = 150$ mbar is where it is shut entirely, and a separate latch blocks the dump permanently at $p_{trip} = 165$ mbar, five inches of mercury absolute, where a real unit would also trip the turbine. The feedback runs the wrong way for the controller, because more dump raises Q_{dump} in (5.16), which raises p_c, which lowers f_{avail} and takes the dump away.

5.3.6 Inventory, Circulating Steam and Indicated Level

The generators hold a water inventory and, below the water surface, a mass of circulating steam occupying real volume. The circulation carries gross evaporation less the condensation on the incoming subcooled feedwater for one riser residence time τ_{res}, as Equation (5.18) states [9].

$$m_{sd}^* = \tau_{res}\left[\underbrace{\frac{1000\,UA_{cs}(T_{avg} - T_s)}{h_{fg}}}_{w_{boil}} - \underbrace{\frac{w_{fw}[h_f - h_{fw}]}{h_{fg}}}_{w_{cond}}\right] \tag{5.18}$$

where both bracketed terms are in kg/s. The two mass balances and the level indication follow, as shown in Equations (5.19) and (5.20).

$$\tau_{void}\frac{\mathrm{d}m_{sd}}{\mathrm{d}t} = m_{sd}^* - m_{sd} \tag{5.19}$$

$$\frac{\mathrm{d}M_{sg}}{\mathrm{d}t} = w_{fw} - w_{tot} - \dot{m}_{sd}$$

$$L = \mathrm{clip}\left(50 + \frac{M_{sg} - M_{sg,0}}{k_M} + k_{void}\Delta V_{sd},\ 0,\ 100\right) \tag{5.20}$$

where $\dot{m}_{sd}$ is the derivative in the first line, τ_{void} is the void relaxation time, k_M is inventory per percent of narrow-range span, k_{void} converts sub-surface steam volume to indicated level and $\Delta V_{sd} =$

$m_{sd}/\rho_s(p_{sg}) - V_{sd,0}$ is the departure of that volume from its design value. The transmitter clips at its taps and the protection system reads the same clipped signal, because both are wired to the same instrument. Equation (5.20) is the reason indicated level is not the water inventory. Raising the pressure compresses the voids and the level shrinks on a plant that has not lost a kilogram, and lowering it does the reverse. This is the shrink and swell of Section 1.4.4, which Chapter 4 studies on a drum. Section 5.6 reports which direction this transient actually takes.

5.3.7 Actuators and Instruments

Four valve positions and the rod drive are states, as shown in Equation (5.21).

$$\tau_{gov}\frac{\mathrm{d}x_{gov}}{\mathrm{d}t} = u_{gov} - x_{gov}, \quad \tau_d\frac{\mathrm{d}x_d}{\mathrm{d}t} = u_d - x_d, \quad \frac{\mathrm{d}z}{\mathrm{d}t} = v_z \tag{5.21}$$

with the atmospheric reliefs and the feedwater demand following the same first-order form at τ_a and τ_{fw}. A valve derivative is zeroed whenever its state rests on a travel limit and the command still points outward, which stops a fixed step from carrying the position through a stop and back, and each position is clipped once more after the step. The rod drive is a velocity actuator and is blocked at the insertion limit and at full withdrawal the same way. A dump stroke lag of $\tau_d = 0.9\,\mathrm{s}$ is three seconds to full open, which is the number to compare against the clock.

Five instruments carry first-order lags, each advanced by the zero-order-hold form $\hat{y}_k = \alpha\hat{y}_{k-1} + (1-\alpha)y_k$ with $\alpha = \exp(-\Delta t_c/\tau)$, which cannot go unstable at any sample period. The average coolant temperature RTD at $\tau_{RTD} = 4.0\,\mathrm{s}$ is the most consequential parameter in the application, because it is longer than the margin the transient has, and the turbine impulse pressure transmitter at $\tau_{imp} = 0.30\,\mathrm{s}$ is the one that makes anticipation possible. The other three are the steam pressure transmitter at $\tau_{PT} = 0.50\,\mathrm{s}$, the generator level transmitter at $\tau_{LT} = 2.0\,\mathrm{s}$ and the excore neutron detector filter at $\tau_{NI} = 0.10\,\mathrm{s}$.

5.3.8 Protection

Four protective actions are modeled and none is a control action, in the layering of Section 1.4.5. The reactor trip acts on high neutron flux above 118 % of rated, on overtemperature ΔT when the loop rise exceeds the setpoint of Equation (5.22), or on low-low generator level below 17 %.

$$\Delta T_{set} = \Delta T_{core}\left[K_1 - K_2(T_{avg} - T_{avg,0})\right] \tag{5.22}$$

where $K_1 = 1.13$ and $K_2 = 0.030/^\circ\mathrm{C}$, so the allowed loop rise falls as the average temperature climbs. That is the protection system's proxy for margin to departure from nucleate boiling, and it is the trip a secondary side failing to take heat away eventually reaches. High-high generator level above 78 % isolates feedwater (the turbine trip that accompanies it on a real unit is not modeled), a reactor trip runs automatic main feedwater back to 20 % of rated, and auxiliary feedwater appears as a floor of 6 % of rated flow on low-low level.

5.3.9 Parameters and Initialization

Table 5.2 carries the constants.

No scenario starts from a nominal condition. Given a turbine load demand, the plant settles where the steam the turbine passes at the resulting pressure equals the steam the reactor's own temperature feedback makes at the programd average temperature, which is a fixed point rather than a formula. The initialization iterates forty passes of damped substitution on the load fraction, derives the fuel temperature from Equation (5.8) at zero derivative, solves Equation (5.5) by bisection for the bank position that makes $\rho = 0$, seeds the precursors at $C_i = \beta_i n/\lambda_i$, and sets the inventory so that Equation (5.20) reads exactly 50.0 %. The result at full power is $T_{avg} = 310.00\,^\circ\mathrm{C}$, $T_f = 610.0\,^\circ\mathrm{C}$, $T_s = 277.01\,^\circ\mathrm{C}$, $p_{sg} = 61.48\,\mathrm{bar}$, $p_{hdr} = 59.98\,\mathrm{bar}$, $w_0 = 1595.0\,\mathrm{kg/s}$, $z = 68.0\,\%$, $M_{sg} = 135\,000\,\mathrm{kg}$ and $p_c = 53.1\,\mathrm{mbar}$, and the largest state derivative there is below 10^{-8} in its own units.

Table 5.2 Model Parameters. Five Constants Are Hidden From the Operator and the Remainder Are Visible in the Model Panel With a Slider.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
P_0	rated thermal power	2900	MW	Δp_0	steam line drop at rated flow	1.5	bar
$W_{e,0}$	gross generator output	940	MW	$p_{imp,0}$	first-stage pressure, full load	45.0	bar
f_P	fission power into the fuel	0.974	—	r_d	steam dump capacity	0.40	—
$m_f c_f$	fuel heat capacity	47.5	MW s/°C	τ_d	dump valve stroke lag	0.9	s
UA_{fc}	fuel-to-coolant conductance	9.415	MW/°C	r_a	atmospheric relief capacity	0.10	—
$m_c c_c$	primary heat capacity	850	MW s/°C	τ_a	relief stroke lag	2.5	s
ΔT_{core}	core rise at full power	33.0	°C	p_{SV}	first safety valve bank	82.7	bar
α_F	Doppler coefficient	−2.8	pcm/°C	Δp_{SV}	safety valve bank spacing	0.85	bar
α_M	moderator coefficient	−30.0	pcm/°C	p_{acc}	accumulation to full lift	2.3	bar
W_{tot}	control bank worth	3000	pcm	r_{sv}	capacity of each SV bank	0.30	—
v_{max}	maximum rod speed	0.50	%/s	$x_{gov,0}$	governor position, full load	0.90	—
z_{lim}	rod insertion limit	18	%	τ_{gov}	governor closing lag	0.30	s
ρ_{scram}	scram rod worth	0.045	$\Delta k/k$	Q_0	condenser duty at full load	1940	MW
t_{scram}	scram rod drop time	2.2	s	$\dot{m}_{cw} c_p$	circulating water capacity rate	194	MW/°C
UA_{cs}	primary-to-secondary, 3 SG	87.9	MW/°C	TTD_0	terminal difference at design	3.0	°C
$(Mc)_s$	secondary heat capacity, 3 SG	1200	MW s/°C	TTD_d	extra difference on dump steam	6.0	°C
$M_{sg,0}$	inventory at full power	135 000	kg	τ_{cond}	condenser response time	8.0	s
k_M	inventory per percent span	666	kg/%	p_{derate}	dump cutback back pressure	115	mbar
k_{void}	void volume to level	0.62	%/m^3	p_{block}	dump block back pressure	150	mbar
τ_{res}	riser residence time	2.0	s	p_{trip}	condenser high-pressure trip	165	mbar
τ_{void}	void relaxation time	8.0	s	τ_{fw}	feedwater response time	4.0	s
T_{fw}	feedwater temperature	226	°C	K_{fwl}	feedwater level trim gain	2.2	%/%
T_{nl}	no-load programmed T_{avg}	291.7	°C	τ_{RTD}	T_{avg} RTD response	4.00	s
T_{fl}	full-load programd T_{avg}	310.0	°C	τ_{imp}	impulse pressure transmitter	0.30	s
K_1	overtemperature ΔT constant	1.13	—	τ_{PT}	steam pressure transmitter	0.50	s
K_2	overtemperature ΔT coefficient	0.030	1/°C	τ_{LT}	level transmitter	2.00	s
n_{trip}	high neutron flux trip	118	%	τ_{NI}	excore detector filter	0.10	s

Hidden constants. Bank position at full power $z_{ref} = 68\,\%$. Fast decay-heat group $f_{D,1} = 0.028$ with $\tau_{D,1} = 25\,$s. Slow decay-heat group $f_{D,2} = 0.040$ with $\tau_{D,2} = 1500\,$s. *Visible, not listed above.* Level trips $L_{hi,trip} = 78\,\%$ and $L_{lo,trip} = 17\,\%$, and the 20 % feedwater limit after a reactor trip.

The engine of Section 1.6.1 advances the states by classical fourth-order Runge-Kutta on a fixed 10 ms step, the controller runs on a 50 ms period and the historian keeps twenty samples for every simulated second. The fastest modeled dynamics are the governor valve at 0.30 s and the sixth delayed-neutron group at 0.33 s, which leaves the step well inside its stability limit, and the model draws no random numbers, so a scenario replays identically.

5.3.10 Assumptions

Every simplification below carries the condition under which it fails.

Point kinetics with the prompt jump taken. Equation (5.3) is correct for a transient measured in seconds to minutes and wrong wherever the 3 ms prompt mode matters, such as a rod ejection. There is no spatial kinetics, so the axial offset transient a rapid rod insertion produces is absent, and there is no xenon.

Two thermal nodes and constant reactivity coefficients. A real core carries a radial and axial power distribution, a hot channel and a departure-from-nucleate-boiling margin that depends on all of it, so Equation (5.22) is the correct functional form fed by a lumped model. The moderator coefficient runs from near zero at beginning of cycle to about −50 pcm/°C at end of cycle and the whole character of the transient changes with it, but the model has no burnup.

No pressurizer. The primary is incompressible, so the pressurizer level and pressure transient, a large part of what a control room actually watches during a load rejection, is absent along with the spray and heater actions that respond to it.

One lumped secondary for three steam generators. Equation (5.10) carries a single saturation temperature, while a real three-loop plant has three generators seeing slightly different feedwater

temperatures and steam line resistances, and loop-to-loop asymmetry is a genuine feature of a real event.

The header is quasi-steady. Equation (5.12) has no steam line inventory, no acoustic transient and no choked flow, so it cannot reproduce the pressure wave an instantaneous valve closure sends up a main steam line.

Safety valves are continuous ramps. Equation (5.15) has no latch and no blowdown, so the masses quoted here are indicative rather than exact, where a real spring valve snaps open and stays open until the pressure falls well below its setpoint.

The condenser is a correlation. Equation (5.16) captures why the dump loads the condenser and not how a condenser condenses. The extra terminal difference on dump steam is a fitted allowance, and air in-leakage, tube fouling and spray attemperation are absent.

Feedwater is a lag on a demand. There is no feed pump, no three-element flow loop and no tube-side dynamics, and auxiliary feedwater appears only as a flow floor, so a main feedwater pump trip cannot be studied here.

The numerics are more repeatable than a control system. Freezing the step and the controller period gives the same trajectory on every machine, where a distributed control system delivers scan jitter, a coarser signal resolution and an execution order that shifts.

5.4 The Clock

The whole control problem follows from how fast the secondary heats up with nowhere to put its steam, so that number is measured first with no controller in the loop. The plant is initialized at the full power design point, the governor command is set to zero, the dump and the reliefs are held shut, the rods are frozen and the feedwater is held at its initial flow, and Equations (5.2) through (5.21) are integrated on a 2 ms explicit step until the generator pressure reaches the first safety valve bank.

The instantaneous heating rate at the design point with the governor already shut is 2.075 °C/s, exactly the arithmetic of Equation (5.10), and the peak rate during the real shut-in is 1.899 °C/s at $t = 1.18$ s, lower because the governor valve takes 0.30 s to close and the bundle duty is already falling as T_s rises. In pressure the peak rate is 1.755 bar/s at $t = 1.29$ s, the generator pressure passes the no-load value of 76.00 bar at $t = 9.62$ s, and it reaches the first safety valve bank at 82.7 bar at $t = 14.80$ s.

Three numbers make that clock harder than it looks. The dump valves stroke in 2.7 s to full open, so the valves alone consume a fifth of the budget after they are told. The average coolant temperature RTD has a response of 4.0 s, so a controller reading it consumes another quarter of the budget before its measurement is correct. And the no-load steam pressure is itself 76.00 bar, because the no-load programmed average temperature of 291.7 °C has that saturation pressure, leaving only 6.70 bar between where the plant has to end up and where the safety valves open. The margin fought over is not the 21.2 bar between full power and the setpoint. It is the overshoot above a hot standby condition the plant has to reach anyway.

The reactor is the one comfortable part of the plant. Nothing tells it the turbine has gone and its only route to the news is T_{avg} in Equation (5.4). With a moderator coefficient of −30 pcm/°C on a coolant that is heating up and Doppler feedback of −2.8 pcm/°C on a fuel node with a five second time constant, core power falls from 100 to 83 % of rated by $t = 9.6$ s and to 73 % by $t = 14.8$ s without a rod moving, which is why the unattended case does not trip on flux. Reactivity feedback buys real time here and it is not enough, and the gap it leaves is why the steam dump exists.

> **TRY IT IN THE APP** **1 · Full Power, and Then the Breaker Opens**
>
> Load scenario 1 and press Run. Nobody acts: the dump stays shut and the feedwater stays on manual. When the breaker opens, watch the steam pressure climb from 61.48 bar toward the safety valve line at 82.7 bar, which it reaches in under fifteen seconds. Keep running and watch the steam generator level rise as the feedwater fills them.

5.4.1 The Process Gain of the Dump

The dump has to be tuned against something, and the steady-state gain from valve position to average coolant temperature is measured with a differenced experiment, because the inventory dynamics drift over the minutes the measurement needs and a single run would attribute that drift to the valve. The turbine is run back from 100 to 85 % of load with everything in manual, a stepped run at a fixed dump position is compared against a companion run held with the dump shut, and the difference is read over $60 \le t \le 180$ s, after which the companion run reaches the high-high level trip and is no longer a valid reference.

Read at $t = 120$ s the settled average temperatures are 312.00 °C with the dump shut, 311.18 at 15 % open, 310.39 at 30 and 308.92 °C at 60 %, so the average gain over the first 60 % of travel is −0.0513 °C/%. Solving the same steady state directly with the bank frozen gives an incremental gain falling from −0.0595 °C/% near the seat to −0.0469 °C/% wide open, because header pressure falls as the flow rises. About 0.052 °C/% is therefore the right number to tune against, and its reciprocal of 19 %/°C is unity loop gain, so the shipped $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$ is conservative by any ordinary standard.

The consequence of that gain is a mismatch of scale rather than a tuning number. A hundred percent of dump valve is worth 5.2 °C of average coolant temperature in the steady state, while the temperature program of Equation (5.23) spans 18.3 °C. The actuator is a factor of 3.5 too small for its own reference program, so on a full rejection the dump is saturated by the setpoint change alone and the only question left is how quickly the controller knows to saturate it.

5.5 Controller Design

Four architectures are selectable and share everything except which signals are wired in. Rod control and feedwater control are on automatic in all four, so the comparison isolates the dump, and only the reactor runback belongs to the coordinated structure. A manual mode drives the dump valves, the reliefs, the rod velocity and the feedwater demand directly, and the transfer between manual and automatic is bumpless in both directions, as Section 1.6.2 describes.

5.5.1 The Temperature Program and What Drives It

A pressurized water reactor programs its average coolant temperature against turbine load, as shown in Equation (5.23).

$$T_{ref} = T_{nl} + (T_{fl} - T_{nl})\,\mathrm{clip}\left(\frac{\ell}{100}, 0, 1.05\right) \tag{5.23}$$

where ℓ is a load signal in percent, $T_{nl} = 291.7\,^{\circ}\mathrm{C}$ and $T_{fl} = 310.0\,^{\circ}\mathrm{C}$. The question that decides this whole application is what supplies ℓ.

The two architectures carrying a turbine instrument take ℓ from first-stage impulse pressure, $\ell = 100\,\hat{p}_{imp}/p_{imp,0}$, where $\hat{p}_{imp}$ is the transmitter output at $\tau_{imp} = 0.30$ s. Impulse pressure measures how much steam the machine is passing, so it moves at the speed of the event because it is the event, and the moment the governor valves shut Equation (5.23) falls 18.3 °C. At $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$ that displacement alone asks for 110 % of a valve that stops at a hundred, and the measured peak of the proportional term over the transient is 132 %. The reference is therefore itself a feedforward, and it is the larger half of the anticipation in the plant.

The feedback only architecture has no turbine instrument anywhere, which is the point of it, so it generates ℓ from the only other load-like measurement available. That is the reactor's own power from the excore neutron detectors, $\ell = 100\,\hat{n}$ filtered at $\tau_{NI} = 0.10$ s, and the reactor sits downstream of the disturbance. Equation (5.4) does not move until the coolant has already heated enough to feed reactivity back, so the reference stands still while the plant heats, the loop has nothing to act on but a four second RTD, and the valves arrive late. Building the comparison this way rather than freezing the reference is what makes it a real four-way comparison, because a reference programd from turbine load is indistinguishable from feedforward and a feedback only mode that kept it would not be feedback only.

In turbine trip mode the reference is the no-load program T_{nl} outright, and the mode is armed by the P-4 turbine trip permissive, which requires both that the machine has tripped and that $x_{gov} < 0.05\,x_{gov,0}$. Until P-4 appears the dump is blocked and its controller tracks zero, so it comes off the seat from the right place the instant the permissive arrives [10].

5.5.2 Anticipation

The explicit anticipatory term is a lead-lag on the load signal rather than a proportional term, as shown in Equation (5.24).

$$x_{ff} = \mathrm{clip}\Big(K_{ff}\left[\ell_{lag} - \ell\right]^{+}, 0, 100\Big), \qquad \tau_{ff}\frac{\mathrm{d}\ell_{lag}}{\mathrm{d}t} = \ell - \ell_{lag} \tag{5.24}$$

where ℓ_{lag} is a first-order lag of the same load signal at $\tau_{ff} = 20\,\mathrm{s}$, $[\cdot]^{+}$ is the positive part and $K_{ff} = 1.20\,\%/\%$, which drives the banks fully open for a load loss of 83 % or more. The washout is not a refinement. After the transient the only heat left in the secondary is what the reactor is still making, so a feedforward proportional to load lost rather than to the rate of losing it would hold the valves open indefinitely and freeze the plant. In turbine trip mode the same term degenerates to a fixed bias armed by P-4 and decaying on the same constant, $x_{ff} = b_d \exp[-(t - t_{P4})/\tau_{ff}]$, because that mode knows only that the machine tripped.

5.5.3 The Dump Demand and Its Two Overrides

Every loop uses the velocity-form PID of Section 1.6.2, written in this chapter's symbols in Equation (5.25).

$$\begin{aligned} e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\ D_k &= -aK_c\tau_D\frac{\mathrm{dPV}^{f}}{\mathrm{d}t} \\ s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias} + s_k + I_{k-1}, \qquad u_k = \mathrm{clip}\,(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}e_k + \frac{u_k - c_k}{T_t}\right]\Delta t_c \end{aligned} \tag{5.25}$$

where $a = -1$ selects direct action, in which a rising measurement raises the output, and $a = +1$ selects reverse action. The running sum s_k is S_k of Equation (1.14), the unsaturated demand c_k is v_k of Equation (1.15) and u_{bias} is u_b. Only I_k sees the clip, through back-calculation with tracking constant $T_t = \tau_I$ [16]. No loop here has a slew limit. Derivative action is zero on every loop, because on a transient decided in six seconds a derivative on a four second measurement contributes noise and no anticipation, and the controller period is $\Delta t_c = 50\,\mathrm{ms}$.

The steam dump loop is direct acting ($a = -1$), because hotter coolant calls for more dump, as Section 1.3 notes, and its reference is deliberately offset above the temperature program, as shown in Equation (5.26).

$$\mathrm{SP}_d = T_{ref} + \Delta_d, \qquad \Delta_d = 1.5\,^{\circ}\mathrm{C} \tag{5.26}$$

Rod control and the steam dump are two integrating controllers acting on one measurement. Offsetting the dump's reference by a degree and a half hands the steady state to the rods, because once the bank has brought T_{avg} back to program the dump's error is negative and its output falls on its low limit, while leaving the transient untouched, since eighteen degrees of error does not care about one and a half. Section 5.6 reports what happens with the offset removed.

Two constraint controllers watch the same actuator and each publishes a cap rather than a contribution, as shown in Equation (5.27).

$$\begin{aligned} c_{pc} &= \mathrm{PID}_{+1}(p_{c,\mathrm{SP}},\ p_c;\ K_{c,pc},\ \tau_{I,pc};\ 0,\ \bar{c}) \\ c_L &= \mathrm{PID}_{+1}\big(L_{hi},\ \hat{L};\ K_{c,L},\ \tau_{I,L};\ 0,\ \bar{c}\big) \\ x_d &= \mathrm{clip}\big(x_{ff} + u_{fb},\ 0,\ \min(100,\ c_{pc},\ c_L)\big) \end{aligned} \tag{5.27}$$

where both constraint loops are reverse acting ($a = +1$) because a rising back pressure or a rising level must lower the cap being published, $p_{c,\mathrm{SP}} = 108\,\mathrm{mbar}$, $L_{hi} = 70\,\%$, and u_{fb} is the temperature trim from Equation (5.25). A cap wins whenever it needs to and is invisible otherwise, where a constraint controller that added to the demand would fight the loop it is protecting [4, 17]. Volume 4, Chapter 3 builds a larger selector of the same kind, with every integrator tracking the output actually applied.

Two implementation details in (5.27) matter more than the gains. The first is that the losing controller is never hard-tracked to a fixed value. Both constraint loops keep executing with their own back-calculation against the ceiling they are given, so their integrals stay consistent with the plant and a re-selection neither steps nor stalls. The second is the ceiling itself, given by Equation (5.28).

$$\bar{c} = \mathrm{clip}\big(\max(x_d,\, x_{ff},\, x_d^{unsat}) + \Delta_{ov},\; 25,\; 100\big) \tag{5.28}$$

where $\Delta_{ov} = 25\,\%$, x_d and x_d^{unsat} come from the previous controller scan and x_{ff} from the current one. An unselected constraint controller parked on its own hundred percent clamp while the working demand is twenty percent has eighty percent of dead travel to wind down through before it can bite, so it arrives as a relay rather than as a controller. Equation (5.28) holds it a stated margin above the working point instead, so it is always within one integral time of being useful, and the application reports an override as selected only when its cap is actually holding the unsaturated demand down. The dump limits are transferred into the feedback controller rather than applied after it, with $u_{min} = -x_{ff}$ and $u_{max} = \max(u_{min}, \min(100, c_{pc}, c_L) - x_{ff})$, so the anti-windup term in (5.25) sees the limit the sum will respect. Without that transfer the trim would wind up against a cap it cannot see.

5.5.4 The Reliefs, the Rods and the Feedwater

The atmospheric reliefs are a direct-acting loop ($a = -1$) on generator pressure against $p_{ARV} = 78.0\,\mathrm{bar}$ at $K_{c,a} = 30\,\%/\mathrm{bar}$ and $\tau_{I,a} = 10\,\mathrm{s}$, which places them between the working pressure and the first safety valve bank at 82.7 bar as the last modulating path before spring valves, and they are present only in the coordinated architecture. In the other three the relief loop tracks its manual station. Settings selected before Run reinitialize the plant and controllers exactly as Reset does. With Coordinated selected on either path, the relief integral starts at zero and the reliefs first open near 78.25 bar in the summer, failed-bank and suggested-challenge cases. A switch during an advancing run retains the manual-tracking history. A relief that has failed open is applied as a floor on position rather than as a demand, so the controller can wind fully down and the steam still leaves.

Rod control is a velocity station with a deadband, as in Chapter 2, and the runback is a bias on it, as shown in Equation (5.29).

$$\begin{aligned} \varepsilon &= \hat{T}_{avg} - T_{ref}, \qquad \delta = \mathrm{sgn}(\varepsilon)\max(0,\; |\varepsilon| - \delta_z) \\ v_z &= \mathrm{clip}(-K_z\delta + b_{rb},\; -v_{lim},\; v_{lim}) \\ b_{rb} &= -v_{max}\,\mathrm{clip}\left(\frac{\Delta\ell}{100}, 0, 1\right), \quad v_{lim} = v_{max} \\ &\qquad \text{whenever } \Delta\ell > \Delta\ell_{arm} \end{aligned}$$

where $\delta_z = 1.0\,^{\circ}\mathrm{C}$ is the deadband that stops the bank walking all day, $K_z = 0.060\,\%/\mathrm{s}$ per °C, $\Delta\ell = [\ell_{lag} - \ell]^+$ is the same load-loss signal that drives Equation (5.24), $\Delta\ell_{arm} = 10\,\%$ and $v_{lim} = 0.20\,\%/\mathrm{s}$ otherwise. The runback is the rod half of the feedforward, lifting the speed limit to the mechanical maximum of 0.50 %/s and biasing the drive inward in proportion to the load lost, so the bank moves before T_{avg} has changed by a degree. It stops at the insertion limit $z_{lim} = 18\,\%$ because going further eats into the shutdown margin the safety analysis assumes. Feedwater is a feedforward on total steam flow with a level trim, $w_{fw} = \mathrm{clip}(w_{tot} + K_{fwl}(L_{\mathrm{SP}} - \hat{L})w_0/100,\, 0,\, w_{cap})$, where w_{cap} is 130 % of rated flow normally and 20 % after a reactor trip.

Table 5.3 Scenarios and What Each Exercises

#	Scenario	Kind	What it exercises
1	Breaker opens, nobody acting	Demonstration	the clock, and reactivity feedback alone
2	Dump by hand on a 15 % runback	Experiment	the process gain of the actuator
3	Operator opens the dump late	Demonstration	reaction time against a stroke time
4	Feedback only	Demonstration	a fast disturbance and a slow measurement
5	Feedforward plus feedback	Experiment	anticipation from an upstream signal
6	Coordinated	Experiment	overrides, reliefs and the runback
7	50 % rejection, two dump modes	Experiment	armed versus continuous anticipation
8	Circulating water at 32 °C	Demonstration	a constraint that moves with the weather
9	One dump bank does not open	Demonstration	unmeasured loss of actuator authority
10	A relief sticks open afterwards	Experiment	a heat sink nobody is controlling
11	85 % dump capacity	Experiment	where the binding constraint moves to
12	Graded challenge	Challenge	the whole structure under three faults

Table 5.4 Measured Outcome of a Full Generator Load Rejection at $t = 5\,\text{s}$, Read Over $0 \le t \le 120\,\text{s}$. The Four Columns Differ Only in Which Signals Are Wired to the Steam Dump.

	Manual	Feedback	FF + FB	Coordinated
Peak p_{sg} (bar)	84.73	83.37	79.74	76.39
Margin to p_{SV} (bar)	−2.03	−0.67	+2.96	+6.31
Safety valve steam (t)	64.69	1.53	0.00	0.00
Relief steam (t)	0.00	0.00	0.00	0.00
Dump steam (t)	0.0	85.1	90.2	64.7
Time to 50 % open (s)	—	12.1	0.9	0.9
Peak T_{avg} (°C)	320.2	319.5	315.4	312.6
Peak p_c (mbar)	53	86	84	80
Time off program (s)	111	105	115	95
Reactor trip	no	no	no	no

5.6 Results

The twelve scenarios appear in Table 5.3 in the order a laboratory session would take them, with the mechanism each exercises. Each one pins down the starting mode, the architecture, the dump mode, the tuning, the disturbances and any timed events, and each carries a completion rule of the kind Section 1.6.4 describes. The challenge is judged on the case as given, so changing a model parameter or disturbance, or firing an extra learner disturbance event, voids the attempt. Every figure below comes from driving the shipped model at its physics step with no user interface attached and reading its own published outputs, with each configuration set after the case is loaded and before Run. Settings made before Run and carried through Reset give identical results. Every accumulated quantity carries the window over which it accumulated.

5.6.1 Four Architectures on One Rejection

Table 5.4 is the central result. The four architectures meet an identical disturbance, scenario 5 with only the architecture changed, with the breaker opening at $t = 5\,\text{s}$ from a settled full-power start, the same $K_{c,d} = 6.0\,\%/^{\circ}\text{C}$ and $\tau_{I,d} = 25\,\text{s}$, the same 40 % dump capacity, the same 4.0 s RTD and rod control on automatic in all four, read over $0 \le t \le 120\,\text{s}$.

Reading Table 5.4 across the row of safety valve masses gives the whole argument. The unattended dump discharges 64.69 t to atmosphere. A properly tuned proportional-integral controller on the same actuator reduces that to 1.53 t, a factor of 42 and still a reportable lift with a valve test attached to it. Handing the identical controller one additional input takes it to zero with 2.96 bar of margin, and the coordinated structure doubles that margin again. The mechanism is one row higher. Feedback alone

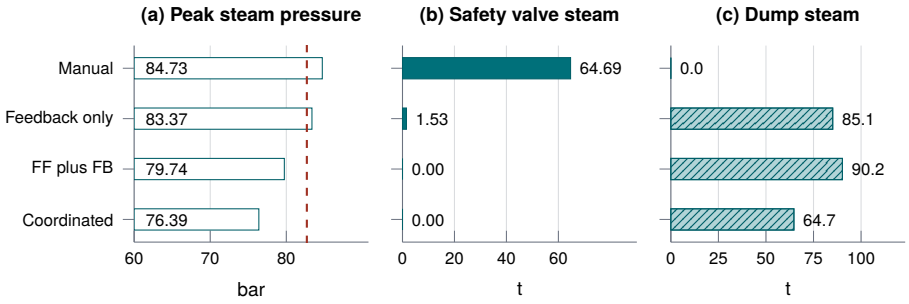


Figure 5.2 Table 5.4 drawn as bars: the full load rejection under the four architectures, read over $0 \leq t \leq 120$ s. (a) Peak steam generator pressure, with the first safety valve bank at 82.7 bar as the dashed line. (b) Steam through the safety valves. (c) Steam through the dump to the condenser. Notice that the two architectures carrying a turbine signal stay left of the dashed line and lift nothing, and that the coordinated one disposes of less dump steam because its runback makes less steam.

does not get the dump banks half open until 12.1 s after the breaker, which is almost exactly as late as a human operator, while both structures carrying a turbine signal are half open at 0.9 s. Figure 5.2 draws three rows of the table as bars.

Two rows read against the intuition they break. The dump steam row is largest for feedforward plus feedback at 90.2 t and smaller for the coordinated architecture at 64.7 t, because the coordinated structure runs the reactor back and has less steam to dispose of. The time off program row is worst for feedforward plus feedback at 115 s, which is not a defect of the measurement, because the counter accumulates whenever $|T_{avg} - T_{ref}|$ exceeds 2 °C in either direction and each architecture is measured against its own reference. A reference programmed from turbine load falls to the no-load value the moment the governor valves shut, so the count starts at once and runs until the coolant has cooled to within 2 °C of hot standby. Over this window every second of it is spent above program. The feedback only reference follows excore power down and stays higher, so the same measure of its run is not a like-for-like comparison. The coordinated architecture is the one that dips below program inside the window, and overcooling is a real failure mode on a plant with a moderator coefficient of –30 pcm/°C, because a colder moderator is positive reactivity.

The safety valve masses in the first two columns depend on the window. In the manual column the discharge is 11.0 t by $t = 30$ s, 34.7 t by 60 s, 64.7 t by 120 s and 75.7 t by 180 s, and it is still accruing at 99.5 t after ten minutes because the safety valves remain the only disposal path in service. The two architectures with a turbine signal discharge exactly nothing at any window.

Figure 5.3 puts the same result in the form the application keeps permanently on screen, because every kilogram per second of steam produced and not disposed of is stored in the secondary as pressure through Equation (5.10). The feedback only panel shows a disposal path collapsed to 60 kg/s one second after the breaker and not exceeding production again until 17.5 s after it, and the safety valve trace shows what pays for the intervening sixteen and a half seconds.

5.6.2 Tuning Does Not Reach It

The strongest form of the lesson is a negative result, because it establishes that the remedy is structural rather than a gain. Table 5.5 sweeps the feedback only architecture over the two things an engineer would ordinarily reach for, with every figure read over $0 \leq t \leq 90$ s.

Seven times the proportional gain buys 1.64 bar and a fivefold shorter integral time buys 0.48 bar. Replacing the four second RTD with an instrument eight times faster than anything a bypass manifold produces buys 0.45 bar, and making the measurement instantaneous buys 0.50 bar, which is the honest asymptote of the instrument remedy. Both remedies together, the gain of 40 with $\tau_I = 5$ s and the 0.5 s RTD, reach 81.30 bar, still 1.56 bar worse than the shipped gain with one wire added.

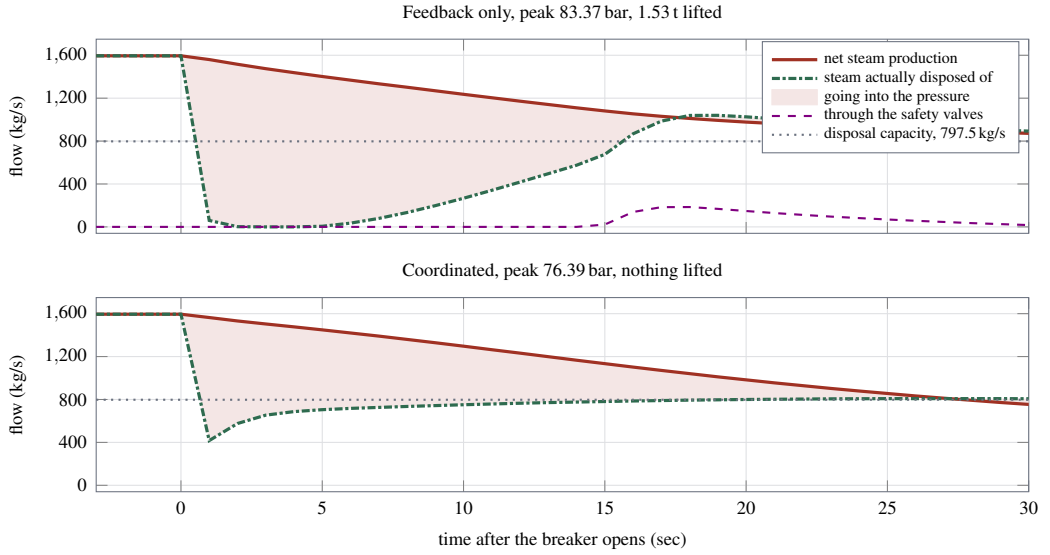


Figure 5.3 The gap between what the steam generators produce and what the plant disposes of, which is the entire transient. Both panels are the same full load rejection with rod control on automatic. Net production is the thick solid line, the steam disposed of the thick dash-dotted line, the flow through the safety valves the dashed line and the disposal capacity the dotted line, and the shaded band between the two thick lines is going into the pressure. Net production is gross evaporation less the condensation on the incoming subcooled feedwater, so it equals disposal exactly at the design point. Feedback only leaves a gap of 1500 kg/s one second after the breaker and does not close it until 17.5 s, 2.6 s after the safety valves begin to lift, while the coordinated structure closes it to 1147 kg/s in the first second and to zero at 27 s. Values are taken at the physics step.

Table 5.5 Feedback Only Under Every Available Remedy. The Bottom Row Is the Same Loop With One Additional Input and No Change of Gain.

Configuration	K_c	Peak p_{sg} (bar)	SV steam (kg)
As shipped, $\tau_{RTD} = 4.0$ s	6	83.37	1529
τ_I raised to 100 s	6	83.65	2226
τ_I cut to 5 s	6	82.89	385
Gain raised	15	82.31	0
Gain raised further	40	81.73	0
Gain raised, $\tau_I = 5$ s	40	81.70	0
RTD eight times faster, 0.5 s	6	82.92	454
RTD instantaneous, 0 s	6	82.87	311
Both remedies together	40	81.30	0
Turbine signal added, gain unchanged	6	79.74	0

The asymptote comes from Equation (5.23) rather than Equation (5.25), because the feedback only reference is generated from excore power and excore power is downstream of the disturbance, so the loop's setpoint stands still for the first several seconds no matter how good its measurement is. What is missing is not speed. It is a signal that moves when the disturbance moves. Volume 3, Chapter 2 reaches the same conclusion on a tank, where each structure rejects the disturbance that enters its own measurement path and leaves the other nearly untouched.

🕮 **Pause.** Try Exercise 5.2 before reading on.

The operator is the same result in human units. With the manual architecture and every dump bank driven fully open at a fixed delay after the breaker, the peak steam pressure is 80.97, 81.54 and 82.47 bar at delays of 1, 5 and 8 s with nothing lifted, then 82.97 bar with 613 kg lifted at 10 s, 83.47 bar with 1997 kg at 12 s and 84.73 bar with 8777 kg at 20 s. The entire difference between a normal day and a reportable event is the two seconds between eight and ten, and no training program makes a two second window reliable. Even acting at one second the peak is 80.97 bar, so a 40 % dump does not make this transient go away. It keeps it under the setpoint.

> ✦ **DID YOU KNOW? One Wire Beats Seven Times the Gain**
>
> From the simulator: raising the gain of the feedback only loop from 6 to 40 lowers the peak steam pressure only to 81.73 bar. Keeping the gain at 6 and wiring in the turbine signal lowers it to 79.74 bar, with nothing through the safety valves (Table 5.5).

5.6.3 Where the Anticipation Actually Lives

The result that surprises practitioners is that the explicit anticipatory term of Equation (5.24) is worth almost nothing at the design tuning. Removing it entirely, with $K_{ff} = 0$ against 1.2 at $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$, changes the peak steam pressure from 79.74 to 79.79 bar and lifts nothing in either case, read over $0 \le t \le 120\,\mathrm{s}$. The reason is the scale mismatch measured in Section 5.4. An 18.3 °C setpoint displacement at 6 %/°C asks for 110 % of a valve that stops at a hundred, so the trim saturates on the reference alone and the summing block has nothing left to add. Detune the trim to $K_{c,d} = 2.0\,\%/^{\circ}\mathrm{C}$, where the same displacement asks for 37 % of valve and the unsaturated demand without the term is only 83 % when the pressure peaks at $t = 27.5\,\mathrm{s}$, and the same term is the difference between 83.40 bar with 1554 kg lifted and 79.75 bar with nothing. The explicit feedforward is not what rejects this disturbance. It is what makes the structure insensitive to its own tuning, which is exactly the property to want on a transient nobody rehearses.

The same term earns its keep outright on a partial event, where the reference step is too small to saturate anything. On the 50 % rejection at $K_{c,d} = 2$, the banks never get past 24 % open with $K_{ff} = 0$ and the plant spends 61 s off program, against 65 % open and 42 s with $K_{ff} = 1.2$, read over $0 \le t \le 300\,\mathrm{s}$.

5.6.4 Attacking the Source Rather Than the Consequence

Of the four things the coordinated architecture adds, the reactor runback of Equation (5.29) does most of the work on a temperate day. Disabling it by raising the arming threshold to 95 % of load lost, the top of its range and above the 92.3 % the load-loss signal reaches, with everything else identical, moves the peak steam pressure from 76.39 bar to 79.10 bar and forces the atmospheric reliefs to discharge 2.3 t where they discharged nothing, read over $0 \le t \le 150\,\mathrm{s}$. The runback is worth 2.71 bar of the 3.35 bar the coordinated architecture gains, it leaves the reactor at 16.5 % of rated power against 31.1 % with the runback disabled, and alone among the four additions it reduces the steam rather than finding somewhere to put it. The condenser override does nothing at all on either run, which is correct behavior rather than a defect, because the peak back pressure is 80 mbar (83 mbar without the runback) against a setpoint of 108 mbar and an interlock that begins cutting the dump back at 115 mbar. An override visible during a normal transient has been set wrongly.

🕮 **Pause.** Try Exercise 5.3 before reading on.

5.6.5 The Constraint That Moves With the Weather

Scenario 8 is the feedforward plus feedback architecture, exactly the configuration with 2.96 bar of margin in Table 5.4, on a day when the circulating water arrives at 32 °C instead of 21 °C. Nothing has failed and nothing alarms. Read over $0 \leq t \leq 150$ s the peak steam pressure is 83.03 bar and 1.2 t goes through the safety valves.

The mechanism is entirely in Equations (5.16) and (5.17). Back pressure peaks at 123 mbar, the interlock takes hold, and at the worst moment the plant is allowed only 78 % of the dump it thinks it has. The controller has its valves at a hundred percent and is getting three quarters of the flow, and there is no flow measurement on a steam dump line to tell it. Switching to the coordinated architecture on the same day gives 79.98 bar with nothing lifted, because the condenser override caps the demand so that back pressure holds at 110 mbar and the plant never enters the interlock band, and because the atmospheric reliefs modulate and take 3.5 t out of the roof. Both initialization paths give these results. The override allows the plant *less* dump than the interlock would have permitted at the worst moment and produces a better transient anyway, because a controller modulates where an interlock ratchets.

> **TRY IT IN THE APP** **8 · A Hot Afternoon on the Intake**
>
> Load scenario 8, with the circulating water at 32 °C, and press Run. Watch the condenser back pressure: it peaks at 123 mbar, the interlock takes hold, and 1.2 t goes through the safety valves. Reset, select the coordinated architecture before the breaker opens and run again: the peak is 79.98 bar with nothing lifted.

Losing a dump bank is worse than losing eleven degrees of circulating water, for the same structural reason. With one of the four banks failing to open, the feedforward plus feedback architecture peaks at 83.31 bar with 3.3 t lifted over the same window, because the controller drives its output to a hundred percent, receives 478 kg/s of capacity instead of 638 kg/s, and cannot tell. The coordinated architecture recovers it to 80.01 bar with nothing lifted, and the recovery comes from the runback and from 3.5 t through the reliefs rather than from dump capacity, which did not exist.

Scenario 11 fits the same unit with an 85 % dump and makes the point from the other direction. The steam pressure becomes a non-event at 68.59 bar, fourteen bar clear of the safety valves. The condenser does not. Solving the fixed point of Equations (5.16) and (5.17) with the banks demanded wide open at the design header pressure gives a settled usable dump fraction of 100 % on 12 °C circulating water, 85 % on 21 °C, 71 % on 26 °C and 55 % on 32 °C, with back pressures of 97, 120, 125 and 131 mbar respectively. A 40 % dump loses at most 0.5 % of its capacity in that range.

At 32 °C an 85 % dump delivers 744 kg/s, only 17 % more than the 638 kg/s a 40 % dump delivers unconstrained on a cool day. The duty explains it. Demanded wide open a 40 % dump puts 1686 MW into a condenser rated for 1940 MW, which is 87 % of design duty from a plant generating nothing, and an 85 % dump puts in 3583 MW, or 1.85 times design. Above roughly 40 % of rated flow the binding constraint stops being valve capacity and becomes condenser duty, and buying more dump valves buys nothing. The coordinated architecture on the same big dump holds back pressure at 112 mbar so the interlock never engages.

5.6.6 The Two Dump Modes and a Partial Rejection

Scenario 7 sheds half the turbine load with the machine still synchronized, so the P-4 permissive never appears. Read over $0 \leq t \leq 200$ s, load rejection mode follows the program for the load actually carried, opens the banks to 99 % and settles T_{avg} at 302.5 °C against a program of 302.0 °C, with a peak of 68.08 bar and 27 s off program. Turbine trip mode never opens the dump at all, peaks at 72.42 bar and spends 69 s off program, with the reactor's own temperature feedback and the rods absorbing the entire event. Nothing is broken, because the mode is doing what it is designed to do, which is nothing until the turbine has actually tripped, and that is why a plant carries both [10].

5.6.7 Two Structures Easier to Assert Than to Measure

Two of the structures described above carry a standard justification that does not survive measurement on this plant, and both are more useful once measured. The application states the measured behavior in its own control hints, its model dialog and its instructional guide, and asserts the first in its test suite.

The first is the level override. Equation (5.20) makes indicated level move with pressure, and the coordinated architecture carries a level cap on the dump for that reason. That cap never selects in any of the twelve scenarios. Measured over $0 \le t \le 300\,\mathrm{s}$ on the full rejection, indicated level falls to 30.6 % in the manual architecture and 34.2 % in the coordinated one and never exceeds 50.2 % in any of the four, with the void term reaching −58.5 % of span while the inventory rises to 174 t. The reason is structural and it is a better teaching point than a swell would have been. Indicated level swells when pressure falls, and on this transient the pressure never falls below where it started, because the no-load steam pressure of 76.00 bar is above the full-power value of 61.48 bar. A load rejection on this plant can only shrink the level, so the nearer of the two level limits is the low-low trip at 17 %, which the manual case approaches to within 13.6 % of span. Re-aiming the same cap at that limit would be worse than leaving it where it is, because a cap on the dump is the wrong actuator for a falling level. Shutting the dump raises the pressure and collapses the voids further, which drives indicated level down rather than up, and low level on this plant belongs to feedwater and to the auxiliary feedwater floor that starts below the low-low trip. An override protecting against a direction the transient cannot travel is correct engineering that happens to be inert on this event, and saying so is more useful than implying that it fires. The application keeps the loop in service, says in the control hint that a load rejection cannot select it, and asserts both halves of that in its test suite: the cap must not select on any of the four architectures, and it must still select and hold when its limit is set below the level being controlled to.

The two scenarios that do reach the high-high level trip at 78 %, 1 and the shut-dump run of 2 (at $t = 182\,\mathrm{s}$), reach it by filling rather than by swelling, and the distinction matters because the two have different cures. In scenario 1 feedwater is on manual at 100 % of rated while gross evaporation collapses to 1325 kg/s by $t = 40\,\mathrm{s}$, so the inventory rises from 135 to 179 t by $t = 62\,\mathrm{s}$, indicated level peaks at 79.4 % at $t = 63\,\mathrm{s}$ and the trip isolates feedwater at $t = 60.4\,\mathrm{s}$. The void term stays between 22 and 39 % of span negative from $t = 30\,\mathrm{s}$ to $120\,\mathrm{s}$, which is the signature of a fill and the opposite of a swell, and the dump is shut in both runs anyway, so a cap on the dump could not have helped.

The second is the reference offset Δ_d of Equation (5.26). Two integrating controllers on one measurement is the classical setup for a limit cycle, and the offset is easy to justify as the cure for one, but no cycle occurs at either setting on this plant. Driven over $300 \le t \le 2400\,\mathrm{s}$ on the coordinated architecture, $\Delta_d = 1.5\,^{\circ}\mathrm{C}$ gives a dump position range of 12.2 % and a steam pressure range of 2.73 bar with one crossing of the mean, and $\Delta_d = 0$ gives 2.6 % and 0.50 bar, also with one crossing. Neither is a cycle. What the offset does do is exactly the structural thing it was put in for, because after forty minutes the dump settles at 1.6 % of valve with $\Delta_d = 1.5\,^{\circ}\mathrm{C}$ and at 49.1 % with $\Delta_d = 0$. The offset is what makes the valves shut and hands the steady state to rod control, at the price of a 1.5 °C standing offset in T_{avg} and a bank driven onto its 18 % insertion limit, where $\Delta_d = 0$ leaves the average temperature exactly on program with the bank at 45.3 % and the dump holding half open indefinitely. That is a real trade between a shut valve and an on-program temperature, and it is a question of which controller owns the steady state rather than a question of stability. Neither of the two settings is pathological and the choice between them is an engineering one, which is a better teaching point than a pathology.

5.6.8 Failures Nothing Reports, and the Challenge

Scenario 10 sticks an atmospheric relief at 40 % open forty seconds after a coordinated rejection has already been handled. The dump closes to 8.0 % and lets the relief do the work, and the average coolant temperature comes back to its program and stays there, which is exactly the robustness that

hides the problem, because the temperature loop cannot tell one steam path from another. The bound depends on where the window starts. At $t = 50\,\text{s}$ the deviation is 11.5 °C, which belongs to the tail of the rejection that began at $t = 5\,\text{s}$ rather than to the relief that stuck at $t = 45\,\text{s}$. From $t = 120\,\text{s}$ to $t = 300\,\text{s}$ the deviation stays within 2.35 °C, worst at $t = 141.5\,\text{s}$, from $t = 180\,\text{s}$ onward within 1.25 °C, and at $t = 300\,\text{s}$ it is 1.04 °C. Over the same five minutes 18.4 t of chemistry-controlled secondary water leaves through the roof and the reactor is driven down to 9.1 % of rated power by a valve nobody is controlling, with nothing out of limits and nothing alarming. The indication that would tell an operator to gag the valve is a mass balance, not a temperature. Tripling the dump valve stroke lag to 2.7 s, so the banks need nine seconds to reach full open instead of three, moves the feedforward plus feedback peak from 79.74 to 80.06 bar and the coordinated peak from 76.39 to 77.02 bar, because the demand arrived at one second.

The graded challenge combines a full rejection with circulating water at 28 °C and one dump bank out of service, and hands the reader turbine trip mode with $K_{c,d} = 2$, $\tau_{I,d} = 90\,\text{s}$ and half the feedforward gain. It is judged at $t = 360\,\text{s}$, and its temperature band applies over the last two minutes. As handed it peaks at 83.93 bar and puts 23.9 t through the safety valves over $0 \le t \le 360\,\text{s}$, and its worst T_{avg} error over the last two minutes is 4.23 °C, so it fails on both counts. The suggested configuration (coordinated, load-rejection mode, $K_{ff} = 1.2$, $\tau_{ff} = 12\,\text{s}$, $K_{c,d} = 8$, $\tau_{I,d} = 20\,\text{s}$, a firmer condenser override, $K_z = 0.08$ and the runback armed at 8 %) finishes at 80.01 bar with nothing through the safety valves, no reactor trip and no condenser trip, with T_{avg} within 1.27 °C of program at six minutes and a peak back pressure of 93 mbar. It passes on both initialization paths with a worst error of 1.98 °C over the judging window, 0.52 °C inside the band. It costs 3.5 t of secondary water through the reliefs and runs the reactor back to 3.8 % of rated power, and both are the price of a dump bank that did not open.

Figure 5.4 shows the feedback-only architecture fourteen seconds into the same rejection Table 5.4 scores. Reproducing that instant numerically gives $p_{sg} = 82.09\,\text{bar}$ against a first-bank setting of 82.70 bar, net production of 1110 kg/s against disposal of 575 kg/s, a dump position of 66.0 % on a demand of 73.5 % with the anticipatory term at zero, $T_{avg} = 317.88\,^{\circ}\text{C}$ with the RTD reading 315.46 °C, core power at 67.2 % and back pressure at 42 mbar, which agrees with every value legible in the image. The first bank begins to lift 0.9 s later, at $t = 19.9\,\text{s}$, and the deficit the inset shades is the quantity Figure 5.3 tracks through the whole event.

5.7 Technology in Practice

The tags of Figure 5.1 are, in a plant, resistance thermometers in the reactor coolant loops, pressure transmitters on the steam lines, the condenser and the turbine's first stage, banks of air-operated dump valves, and the controllers of a non-safety control system. Figure 5.5 draws the dump loop with that hardware in place of the schematic's symbols.

What Measures It

Average coolant temperature comes from resistance temperature detectors, which with thermocouples are the most common temperature sensors in nuclear plants [20]. Steam pressure, condenser back pressure and turbine first-stage pressure come from pressure transmitters, which convert the pressure force to an analog signal such as 4 to 20 mA or, in newer smart transmitters, to a digital serial bus signal [20]. The chapter carries TE-411 as a 4.0 s lag, the impulse pressure transmitter as a 0.30 s lag and PT-514 as a 0.50 s lag. The gap between the first two is the lesson of Section 5.4.

> **TECHNOLOGY** **Sensor: The Resistance Thermometer in the Coolant**
>
> A resistance temperature detector is a fine wire of pure metal, usually platinum or copper, whose resistance rises with temperature. It is not very sensitive, but it is very linear [21]. With the thermocouple it is one of the two most common temperature sensors in nuclear plants, set in contact with the fluid or in a thermowell [20]. A thermowell lets the element be changed without opening the process, but its metal adds a first-order lag,

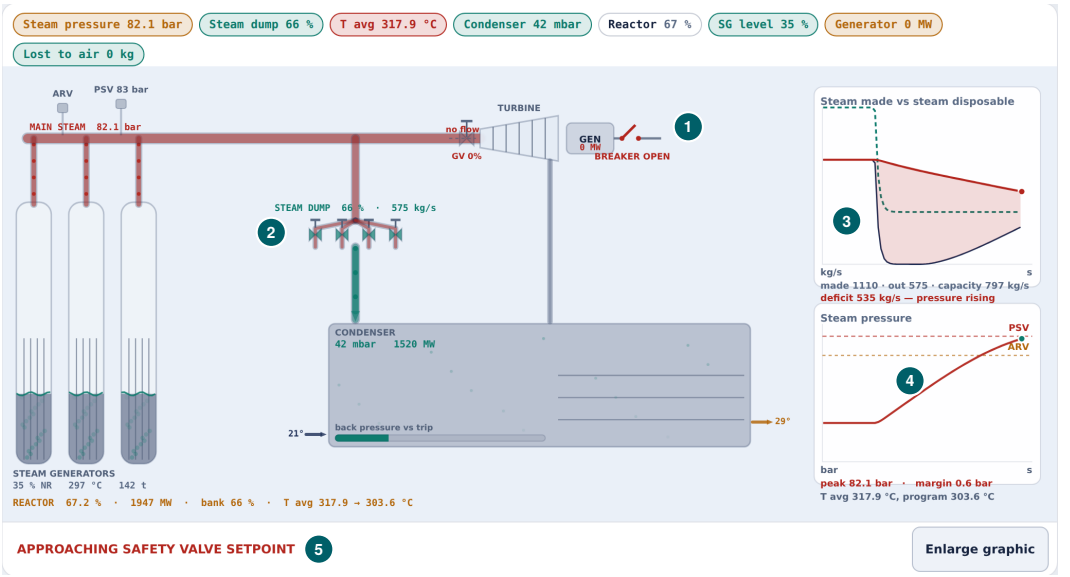


Figure 5.4 The feedback-only architecture at $t = 19$ s, fourteen seconds into a full load rejection and under a second before the first safety valve bank begins to lift, in the plant view of the app. The breaker is open and the governor valves read zero (1), so no steam is reaching the turbine. The dump banks (2) stand at 66.0 %, and the secondary is making 1110 kg/s net and disposing of 575 kg/s of it against a total capacity of 797 kg/s. The inset (3) shades that 535 kg/s deficit, and the pressure is taking it: the steam generators stand at 82.09 bar with the first safety valve bank set at 82.70, a margin of 0.6 bar (4), and the annunciator (5) reads the approach. T_{avg} is 317.88 °C, the reactor has run back to 67.2 %, the condenser is absorbing 1520 MW at 42 mbar, and nothing has yet been lost to atmosphere.

> and an air gap at the bottom of the well adds a second one. On a trend the reading trails the process, and a lag severe enough can make a feedback loop oscillate [21]. TE-411 stands for all of this as a single 4.0 s lag.

What Moves It

The final elements are the four dump banks, the atmospheric reliefs and, through the runback, the control rods. Steam dump to the condenser is listed among the non-safety control functions of a pressurized water reactor, beside turbine control, the main steam atmospheric relief valves and rod control [22]. The main steam safety valves sit outside every controller, which is why Equation (5.15) gives them no command. In the chapter each dump bank strokes as a first-order lag of 0.9 s, an idealization with no dead band and no stiction.

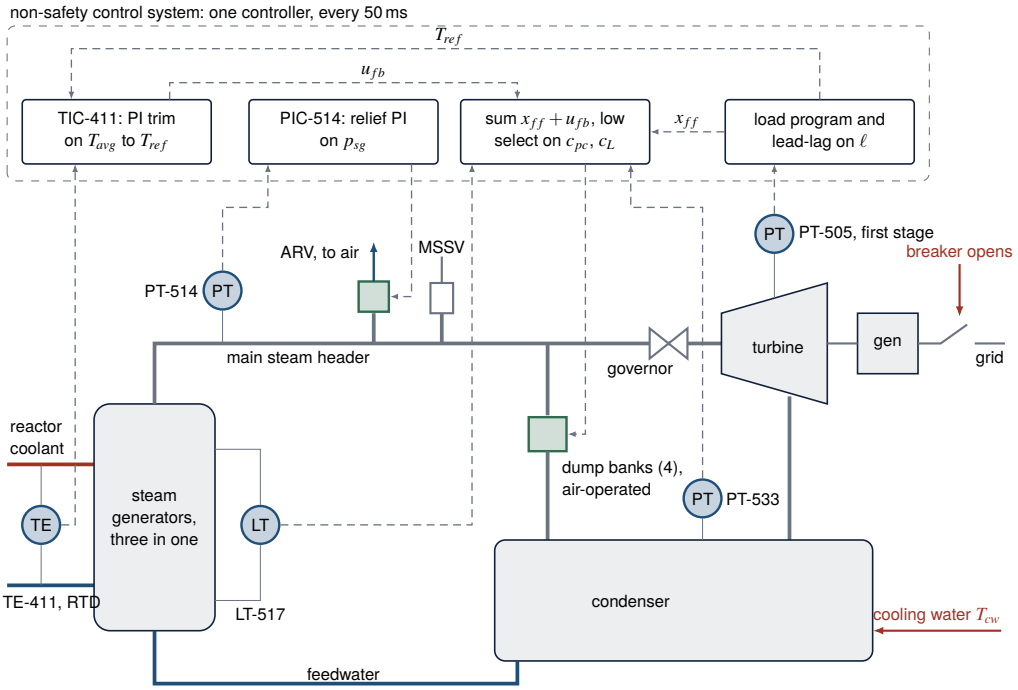


Figure 5.5 The steam dump loop of Figure 5.1 with plant hardware in place of the symbols. Sensors are circles and the two final elements the controller moves, the dump banks and the atmospheric relief, are green squares. Dashed arrows are signals. The temperature trim reads TE-411 between the hot and cold legs, the load program reads PT-505 on the turbine's first stage, and the low selector caps the dump from PT-533 and LT-517. The governor valve and the main steam safety valves (MSSV) take no command from this controller, and the rod runback, the third final element, is not drawn. Notice that PT-505, on the turbine, is the one sensor that moves the instant the breaker opens.

TECHNOLOGY **Final Element: The Air-Operated Steam Dump Valve**

In an air-operated valve, air on a diaphragm pushes the stem against a spring. A converter turns the controller's 4 to 20 mA into air pressure, and a positioner adjusts the air until the stem reaches the demanded position. In some safety-related designs, solenoid valves dump the air and the spring returns the valve to its fail-safe position [22]. Dump valves to the condenser must stroke very quickly on a turbine load rejection, yet hold stable pressure control during startup and shutdown. At one unit, digital positioners on these valves had more air capacity than the analog ones they replaced, so the volume boosters could be removed [22]. The chapter's dump banks need three seconds from shut to full open.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, under a supervisory layer that sends setpoints to the field controllers [23]. The chapter's controller executes every 50 ms over a 10 ms physics step. The structure is a set of function blocks in one controller: a temperature PI whose reference is the load program, a lead-lag on the load signal summed with it, two constraint PI blocks that publish caps, and a low selector that applies the smallest, with the applied limit fed back to every integrator. The runback is a bias on the rod velocity station, and turbine trip mode is armed by the P-4 permissive from the protection system.

> **IN PRACTICE** **A Rejection That Ends at House Load**
>
> Most plants are designed to ride out a limited, rapid load rejection by cutting back reactor power without a reactor trip, on the provision that such an event is rare. A unit able to drop to house load, with the grid disconnected and the generator supplying only the plant's own auxiliaries, typically brings the reactor down to about 20 to 30 percent of rated power and sends the excess heat to the condenser. Such a period is usually limited to a few hours, by constraints that include the capacity of the condensation system [24]. The runback of Equation (5.29) is the reactor's half of that cutback.

What Is Modernizing It

Three developments reach this loop. Digital positioners on control valves diagnose air leakage, friction and dead band, so that degradation is found before an upset, and they allow partial stroke tests on line [22]. On-line monitoring mines the data the control system already records for underlying changes in the relationships between measured parameters. It can show that a sensor has not drifted since its last calibration and so extend the calibration period [20]. And control systems increasingly resemble office computing as they adopt its computers, operating systems and network protocols, which gives them far less isolation from the outside world and a greater need to be secured [23]. Operating nuclear plants are moving from analog to digital control through individual system upgrades [20], so each upgraded dump controller joins that network.

5.8 Limitations

Nothing here is a safety analysis and nothing here is a plant model. Point kinetics, an absent pressurizer, one lumped secondary standing for three steam generators, fitted saturation properties and a correlation where a condenser belongs together rule out any statement about a real unit's peak secondary pressure, its safety valve sizing, its margin to departure from nucleate boiling, or whether some particular plant would ride out some particular load rejection. A licensing calculation for this transient uses a qualified system code against acceptance criteria written for the purpose [13], and the correct use of the numbers in Section 5.6 is as comparisons between control structures on one model, never as absolute predictions.

Three specific consequences follow. The safety valve masses are indicative rather than exact, because Equation (5.15) is a continuous ramp with no latch and no blowdown. The condenser results state the shape of the constraint correctly and its magnitude only to the accuracy of a fitted terminal temperature difference, so they establish that condenser duty becomes binding above roughly 40 % of rated dump flow without fixing the crossover point [12]. And the whole character of the transient depends on the moderator temperature coefficient, which is a constant here and runs from near zero at beginning of cycle to about −50 pcm/°C at end of cycle in a real core, so the reactivity feedback that buys the first ten seconds in Section 5.4 is worth much less early in a fuel cycle than the shipped value suggests.

5.9 Discussion

The engineering conclusion of Table 5.4 is not that feedforward is good. It is that the peak steam pressure on a load rejection is set by the instrument list rather than by the controller settings, and that the decision fixing it is made at design time when somebody chooses whether to run a cable from the turbine first-stage pressure transmitter to the steam dump controller. A factor of seven in gain is worth 1.64 bar, an instrument eight times faster than physically available is worth 0.45 bar, and the cable is worth 3.63 bar. Every disposal path in the plant addresses the consequence of the disturbance and the runback addresses its source, and on the shipped constants the source is worth more than any of the consequences at 2.71 bar of the 3.35 bar the coordinated architecture gains, because the same load-loss signal that opens the dump valves can drive the control rods and the anticipatory measurement is reused rather than duplicated.

Condenser back pressure is the quietest constraint in a steam plant. It appears in no control narrative, it varies with the weather and with tube cleanliness, and through Equation (5.17) it silently rescales the authority of the one actuator the whole transient depends on. Eleven degrees of circulating water temperature turned a transient with 2.96 bar of margin into one that lifted 1.2 t, with nothing failed and nothing alarmed. The remedy is not more valve, because an 85 % dump delivers only 744 kg/s on 32 °C water against 638 kg/s from a 40 % dump on a cool day. The remedy is a controller in front of the interlock, holding the plant off protection so that a ratcheting constraint is never approached, and that principle transfers to any plant where protection and control share an actuator. Methods beyond the implemented loops fit the same model, since a nineteen-state balance carrying a hard dump capacity, a moving condenser constraint and a rod insertion limit is a constrained model predictive control problem in which the low selects of Equation (5.27) turn into inequality constraints and the anticipation of Equation (5.24) turns into an explicit disturbance forecast from the impulse pressure measurement [17].

OPEN CHALLENGE Steam Bypass as a Daily Tool

A plant asked to cut its electrical output for a short time can bypass the turbine and send steam to the condenser while the reactor stays near full power, provided its feedwater, condenser and valves are designed for it [24]. The price is wear. Redirected steam loads the condenser shell and tubes, the bypass lines, the turbine stop and throttle valves and the steam spargers inside the condenser, and more so when bypass is the only means of flexible operation [24]. Progress would be a dump controller that weighs this wear against moving the reactor, the trade the runback of this chapter makes on a single event.

5.10 Conclusion

This chapter studied a pressurized water reactor secondary side under a generator load rejection as a working laboratory for a disturbance that arrives faster than any measurement of its effect can develop. Nineteen states are integrated against saturation property fits and a main steam header solved in closed form, covering six delayed-neutron precursor groups under a prompt jump, two decay-heat groups, fuel and coolant temperatures, a lumped secondary saturation temperature with its water inventory and its sub-surface steam holdup, a control bank position, four valve actuators and a condenser back pressure. Shut in at full power the secondary heats at 2.075 °C/s, passes the no-load pressure of 76.00 bar at 9.62 s and reaches the first main steam safety valve at 82.7 bar at 14.80 s, against dump valves that stroke in three seconds and an average-temperature RTD with a four second response. On one identical rejection the unattended dump peaks at 84.73 bar and discharges 64.69 t to atmosphere within two minutes, a properly tuned temperature loop with no turbine instrument peaks at 83.37 bar and discharges 1.53 t, the same loop with turbine impulse pressure wired to it peaks at 79.74 bar and lifts nothing, and a coordinated structure with two low-selected overrides and a reactor runback peaks at 76.39 bar. Seven times the gain reaches 81.73 bar and an instantaneous RTD reaches 82.87 bar, so neither tuning nor instrumentation closes a gap that only a measurement of the disturbance closes. The reactor runback is worth 2.71 bar of the coordinated architecture’s 3.35 bar advantage, and on 32 °C circulating water the plant’s own condenser interlock removes 22 % of the dump capacity at the moment it is needed.

A fuller model would add a pressurizer so that the primary pressure transient and its spray and heater actions appear, three separately modeled steam generators so that loop asymmetry is visible, latching safety valves with blowdown so that the discharge masses become quantitative, a real three-element feedwater system, and a burnup coordinate so that the moderator coefficient moves through a fuel cycle. What the steam dump plant gives a control engineer is fifteen seconds in which structure beats gain.

KEY POINTS

- When the turbine valves shut, the steam generators keep making steam, and the pressure reaches the first safety valve in under fifteen seconds.
- Seven times the gain on the four second RTD reaches 81.73 bar, while the same loop with turbine impulse pressure wired to it reaches 79.74 bar and lifts nothing.
- The reactor runback is worth 2.71 bar of the coordinated architecture's 3.35 bar advantage, because it reduces the steam rather than finding somewhere to put it.
- A controller in front of the condenser interlock beats the interlock: on 32 °C circulating water the coordinated structure lifts nothing where feedforward plus feedback lifts 1.2 t.

Where else this applies. A protective override on a fast transient also guards the compressor of Volume 4, Chapter 2, and Volume 4, Chapter 3 low-selects limiters over a governor as the dump's overrides do here.

References

[1] 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.

[2] K. J. Åström, "Limitations on control system performance," *European Journal of Control*, vol. 6, no. 1, pp. 2–20, 2000.

[3] J. L. Guzmán and T. Hägglund, "Tuning rules for feedforward control from measurable disturbances combined with PID control: a review," *International Journal of Control*, vol. 97, no. 1, pp. 2–15, 2024.

[4] F. G. Shinskey and J. P. Shunta, "Selective, override, and limit controls," in *Process Control*. Elsevier, 1995, pp. 112–118.

[5] J. T. Gravdahl and O. Egeland, "Modelling and control of surge for a centrifugal compressor with non-constant speed," in *Compressor Surge and Rotating Stall*, ser. Advances in Industrial Control. Springer London, 1999, pp. 143–176.

[6] L. Wang, C. Chen, L. Chen, Z. Li, and W. Zeng, "A coordinated control methodology for small pressurized water reactor with steam dump control system," *Energy*, vol. 282, p. 129008, 2023.

[7] E. Lin, Y. H. Cheng, and J. K. Hsiue, "Analysis of chinshan unit 1 full load rejection transient using RETRAN," *Nuclear Technology*, vol. 54, no. 3, pp. 380–384, 1981.

[8] M. D. Murrell, "Boiler simulation, PWR station dynamics," in *Boiler Dynamics and Control in Nuclear Power Stations*. Emerald Publishing Limited, 1973, pp. 4.1–4.9.

[9] T. W. Kerlin, E. M. Katz, J. Freels, and J. Thakkar, "Dynamic modelling of nuclear steam generators," in *Boiler Dynamics and Control in Nuclear Power Stations 2*. Emerald Publishing Limited, 1980, pp. 249–258.

[10] U.S. Nuclear Regulatory Commission, Human Resources Training and Development, "Westinghouse technology systems manual, section 11.2: Steam dump control system," USNRC HRTD, Rev. 0403, pp. 11.2–1–11.2–8, 2003, aDAMS Accession No. ML11223A294. [Online]. Available: https://www.nrc.gov/docs/ML1122/ML11223A294.pdf

[11] A. Logar, T. Depolt, and E. Gobrecht, "Advanced steam turbine bypass valve design for flexible power plants," in *2002 International Joint Power Generation Conference*. ASMEDC, 2002, pp. 43–49.

[12] D. M. Nightingale, "Design guidelines for the safe operation of steam surface condenser turbine bypass on combined cycle power plants," in *Proceedings of the ASME 2017 Power Conference*. American Society of Mechanical Engineers, 2017, p. V001T05A001.

[13] U.S. Nuclear Regulatory Commission, "Standard review plan 15.2.1–15.2.5: Loss of external load; turbine trip; loss of condenser vacuum; closure of main steam isolation valve; and steam pressure regulator failure," U.S. Nuclear Regulatory Commission, Washington, DC, Tech. Rep. NUREG-0800, Rev. 2, Mar. 2007, aDAMS Accession No. ML070300702. [Online]. Available: https://www.nrc.gov/docs/ML0703/ML070300702.pdf

[14] OECD and Nuclear Energy Agency, *Technical and Economic Aspects of Load Following with Nuclear Power Plants*. Paris: OECD Publishing, 2021.

[15] W. E. Reeves, "Technical evaluation of the proposed deletion of a reactor trip on a turbine trip below 50-percent power for the Beaver Valley nuclear power plant, Unit 1," Lawrence Livermore Laboratory, Livermore, CA, Tech. Rep., Dec. 1979.

[16] R. Hanus, M. Kinnaert, and J.-L. Henrotte, "Conditioning technique, a general anti-windup and bumpless transfer method," *Automatica*, vol. 23, no. 6, pp. 729–739, 1987.

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

[18] R. Goldstein and L. M. Shotkin, "Use of the prompt-jump approximation in fast-reactor kinetics," *Nuclear Science and Engineering*, vol. 38, no. 2, pp. 94–103, 1969.

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

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

[21] 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.

[22] 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.

[23] 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.

[24] 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.

5.11 Exercises

5.1 Concept check · 1 min

Equation (5.10) sets the heat capacity $(Mc)_s$ times the rate of change of T_s equal to the bundle duty less two steam terms. What does the equation describe?

(a) A balance on one delayed-neutron precursor group of the core

(b) An energy balance on the saturated secondary side of the steam generators

(c) An energy balance on the fuel node that receives the fission power

(d) The first-order lag of the steam dump valve position on its demand

5.2 Concept check · 2 min

On the full rejection, the feedback only architecture with $K_{c,d} = 40\,\%/^{\circ}\mathrm{C}$ and $\tau_{I,d} = 5\,\mathrm{s}$ peaks at 81.70 bar, while feedforward plus feedback at the shipped $K_{c,d} = 6\,\%/^{\circ}\mathrm{C}$ peaks at 79.74 bar. Which statement explains the gap?

(a) A high gain on a four second RTD makes the loop oscillate, so its peak cannot fall further

(b) The integral winds up against the full-open limit and holds the valves back after the breaker

(c) Its reference follows excore power, which lags the coolant, so no input moves when the breaker opens

(d) The dump valves stroke more slowly when the demand changes faster, so a higher gain cannot open them sooner

5.3 Concept check · 2 min

On a 21 °C day the coordinated architecture of scenario 6 peaks at 76.39 bar, against 79.74 bar for feedforward plus feedback. Which of the four additions does most of that work?

(a) The condenser back-pressure override, which caps the dump as the back pressure rises

(b) The steam generator level override, which holds the indicated level off its high limit

(c) The modulating atmospheric reliefs, which open above their 78.0 bar setpoint

(d) The reactor runback, which drives the bank in on the same load-loss signal

5.4 Theory and calculation · 6 min

Unit U-1 runs at its design point: 1595.0 kg/s of steam at $p_{sg} = 61.48$ bar, feedwater at 964.3 kJ/kg, a bundle duty of 2900 MW and $(Mc)_s = 1200\,\mathrm{MW\,s}/^{\circ}\mathrm{C}$. Use the saturation fits of Equation (5.1).

(a) Compute the latent and sensible terms of Equation (5.10) and show that the right-hand side is zero. *Checkpoint:* 2490 MW and 410 MW, summing to 2900 MW.

(b) The governor valves shut and every steam flow stops, with everything else held at the design point. Find the heating rate in degrees and, with the slope 0.906 bar/°C, in bar per second. *Checkpoint:* 2.075 °C/s and 1.88 bar/s.

(c) Find the saturation temperature at the first safety valve bank, 82.7 bar, and the time to reach it at the rate of (b). Compare it with the measured 14.80 s and give two reasons for the difference. *Checkpoint:* 297.57 °C, a rise of 20.56 °C, about 9.9 s.

5.5 Theory and calculation · 5 min

Section 5.4 measures the settled average coolant temperature at 85 % turbine load as 312.00 °C with the dump shut and 308.92 °C with it 60 % open.

(a) Find the average process gain of the dump over that travel. *Checkpoint:* −0.0513 °C/%.

(b) Taking the gain magnitude as 0.052 °C/%, find the controller gain that gives unity loop gain, and the loop gain of the shipped $K_{c,d} = 6\,\%/°\mathrm{C}$. *Checkpoint:* 19.2 %/°C and 0.31.

(c) The program of Equation (5.23) falls 18.3 °C when the impulse pressure collapses. Find the proportional demand this asks for at $K_{c,d} = 6$ and at $K_{c,d} = 2$, and the steady temperature change that a fully open dump can make. *Checkpoint:* 110 %, 37 % and 5.2 °C.

5.6 Theory and calculation · 6 min

The control bank of U-1 has a total worth $W_{tot} = 3000$ pcm with the S curve of Equation (5.5), sits at $z_{ref} = 68\,\%$ at full power, and the runback drives it at $v_{max} = 0.5\,\%/\mathrm{s}$. The moderator coefficient is −30 pcm/°C.

(a) Find the bank worth and the differential worth at z_{ref}. *Checkpoint:* 2472 pcm and 42.8 pcm/%.

(b) Find the reactivity rate of a full-speed runback at that position. *Checkpoint:* 21.4 pcm/s.

(c) Turbine-trip mode is armed on a machine that still carries half its load, so the dump aims at the no-load program of 291.7 °C instead of the half-load value of Equation (5.23). Find that half-load value, the reactivity the overcooling adds, and how long the runback would take to insert the same worth. *Checkpoint:* 300.85 °C, +274.5 pcm and about 12.8 s.

5.7 App challenge · 8 min

Load scenario 4, *Feedback only*, in Auto with Steam dump architecture on Feedback only. The breaker opens at 5 s. Use speed 2 and run each case for 60 s, then export the CSV file.

(a) Run the scenario as shipped. Read the Peak steam pressure and Safety valve steam tiles, and find in the file the first row at which `dump_pos_pct` reaches 50 %. *Checkpoint:* 83.4 bar and 1529 kg, within 5 kg. Half open at 17.15 s, 12.1 s after the breaker, within 0.1 s.

(b) Press Reset, set Dump temperature gain to 40 and Dump integral time to 5 s before Run, and run again. Read the same tiles and the last value of `metric_peak_psg_bar`. *Checkpoint:* 81.70 bar, within 0.01 bar, and 0 kg.

(c) The shipped feedforward plus feedback architecture peaks at 79.74 bar. Explain why seven times the gain leaves the loop 1.96 bar short of it. *Checkpoint:* The reference does not move with the breaker.

5.8 App challenge · 8 min

Load scenario 5, *Feedforward plus feedback*, in Auto with Steam dump architecture on Feedforward + FB. Run four cases of 60 s at speed 5, each set before Run from a reload of the scenario, with

Feedforward gain (Advanced) and Dump temperature gain as follows: 1.2 and 6, 0 and 6, 0 and 2, 1.2 and 2. Export the CSV file after each run.

(a) Read the last value of `metric_peak_psg_bar` and the Safety valve steam tile for each case. *Checkpoint:* 79.74 bar, 79.79 bar, 83.40 bar with 1554 kg lifted, 79.75 bar, each within 0.01 bar.

(b) Find the peak pressure the explicit feedforward term saves at each gain. *Checkpoint:* 0.05 bar at $K_{c,d} = 6$ and 3.65 bar at $K_{c,d} = 2$.

(c) Explain the difference with the proportional demand that the program step asks for at each gain. *Checkpoint:* 110 % against 37 % of valve.

5.9 App challenge · 8 min

Load scenario 8, *A hot afternoon on the intake*, in Auto. The circulating water is at 32 °C and the breaker opens at 5 s. Run each case to 150 s at speed 5 and export the CSV file.

(a) Run with Steam dump architecture on Feedforward + FB as shipped. Read Peak steam pressure, Safety valve steam and Peak back pressure, and the lowest value of `cond_derate` in the file. *Checkpoint:* 83.0 bar, 1223 kg within 5 kg, 123 mbar and 0.78.

(b) Press Reset, set Steam dump architecture to Coordinated before Run and run again. Read the same three tiles, Lost to atmosphere and the lowest `cond_derate`. *Checkpoint:* 80.0 bar, 0 kg, 110 mbar, 3519 kg within 10 kg, and 1.00.

(c) Explain why the coordinated structure does better while it lets the condenser take less steam than the interlock allowed. *Checkpoint:* A controller modulates in front of the interlock.

5.10 Application area · 4 min

A plant considers replacing its 40 % dump with one of 85 % of rated steam flow, $w_0 = 1595\,\mathrm{kg/s}$. Dump steam enters the condenser at $h_g(p_{sg,0})$ and leaves at 140 kJ/kg, and the condenser is rated for 1940 MW.

(a) Find the duty of each dump wide open at the design pressure, as a fraction of the condenser rating. *Checkpoint:* 1686 MW, 87 %, and 3583 MW, 1.85 times the rating.

(b) On 32 °C water the larger dump delivers 744 kg/s once the interlock settles. Find the usable fraction and compare the flow with the 638 kg/s of the smaller dump on a cool day. *Checkpoint:* 0.549 of its capacity, 17 % more flow.

5.11 Application area · 3 min

In the manual architecture every dump bank is opened fully at a fixed delay after the breaker. A delay of 8 s peaks at 82.47 bar with nothing lifted, 10 s lifts 613 kg and 12 s lifts 1997 kg. A procedure writer proposes to rely on the control room operator. Give two reasons from the chapter why the dump must be automatic, and why an automatic temperature loop alone is not the answer either.

5.12 Review the agent · 5 min

An AI assistant reviews Table 5.4 and writes: “Feedback only spends 105 s off program against 115 s for feedforward plus feedback, so it holds the temperature program better. Its safety valve lift is only 1.53 t. A plant that values temperature control should prefer feedback only.” Find the flaw in the argument and say how to check it on the app.