INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 9

Multi-Module Steam Header Simulator for Load Allocation and Plantwide Coordination

Pages 229–260

Edited by John D. Hedengren

How to cite this chapter. "Multi-Module Steam Header Simulator for Load Allocation and Plantwide Coordination," 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. 9, pp. 229–260.

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

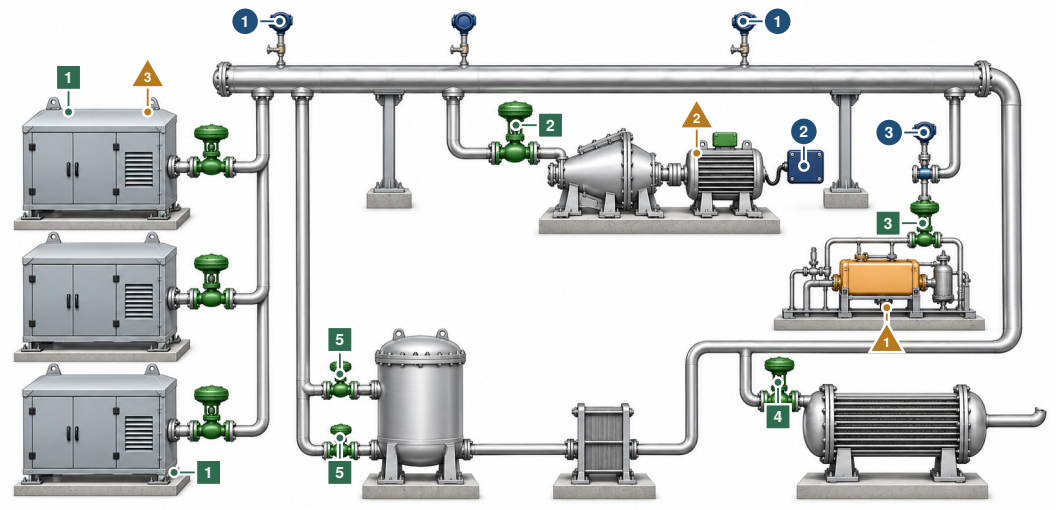
The app for this chapter
https://apmonitor.com/apps/smr-steam-header/

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

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

CHAPTER 9

Multi-Module Steam Header Simulator for Load Allocation and Plantwide Coordination



The controller moves 1 reactor power of each module 2 turbine valve 3 process steam valve 4 bypass valve to the condenser 5 storage charge and discharge valves

The controller reads 1 header pressure, three transmitters 2 generator output 3 process steam flow

What disturbs it 1 process steam request 2 grid load 3 a module trip

Notice that every module and every user meets in the one header pressure, the only number they all see. Illustration, not to scale.

IN PLAIN TERMS **Three Reactors Sharing One Steam Pipe**

The main idea. Three small reactors send steam into one shared pipe that feeds a turbine and a factory. The pressure in the pipe shows whether the reactors make as much steam as the users take.

The challenge. Give each reactor its own controller for that pressure, and the reactors quietly drift apart while the pressure looks perfect. One controller must answer the pressure, and a separate rule must share the work so that no reactor runs nearer its heat limit than the others.

Why it matters. A pipe at the right pressure can hide a reactor pushed to its limit.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of Section 9.1, the history box of Section 9.2, the try-it and fact boxes of Section 9.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 9.6, Section 9.7, Section 9.8 and the open challenge

The control question Three small reactors feed one header that holds 7.55 s of steam. Who should answer a header pressure error, and by what rule should the load be divided when each module's ceiling moves with the header?

Learning objectives

- Compute the time scales of the header, the modules and their ramp limits.
- Explain why independent integral loops on one measurement walk apart.
- Derive margin-proportional allocation and compare it with an equal split.
- Compare master-slave with coordinated control and price the bypass rest.
- Decide which product a runback should shed after a module trip.

Prerequisites Sections 1.6.2, 1.3, 1.4.5 and 1.6.4, and Section 1.6.3 for the predictive controller. Chapter 4 and Chapter 5 follow the steam of one reactor.

The app https://apmonitor.com/apps/smr-steam-header/

See also Volume 6, Chapter 7 and Volume 6, Chapter 8 for plantwide control structure. Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging a fast, costly actuator, and Volume 4, Chapter 12 for allocation among actuators at different prices.

9.1 Introduction

Three reactors that feed one pipe share one number. Header pressure rises when the modules make more steam than the users take and falls when they take more, so it is the plant's statement of whether supply matches demand, and it is also the variable no single piece of equipment owns. Every module pushes it up, every user pulls it down, and the pressure reaches back into all of them at once, because it sets the saturation temperature each steam generator boils against and the flow each user's valve passes. A block diagram of one module and one pressure loop is a correct picture of a plant that does not exist. The plant that exists has three producers with different capability, four consumers with different value, and an inventory of steam between them that lasts 7.55 s.

That configuration produces a failure that a transfer function cannot show. Put an identical, honestly tuned proportional-integral controller on each module, all three holding the same header pressure setpoint, and apply no disturbance at all. The three module powers walk apart, reaching 70.5 MW of spread at 600 s and 91.7 MW at 1200 s, while header pressure stays between 3.385 and 3.401 MPa inside a tolerance band of ±0.08 MPa. Nothing on the header alarms and nothing oscillates. The only signs are per-module ones, a margin-spread score that turns red at 135 s with the modules 20.5 MW apart and a module coolant temperature warning at 271 s, by which time they are 48 MW apart. The misconception this chapter corrects is the belief that a controlled variable held on setpoint is evidence that the plant is under control, when what is under control is one measurement and the megawatts behind it are free to redistribute until two of the three modules sit against their own thermal limits.

The class of problem is general. Several nearly identical producers discharge into one small-inventory buffer, consumers draw from it, and the buffer pressure or level is the only variable they all see. Parallel boilers on a common steam header are the oldest instance, and the plantwide question of which loop owns a shared variable and by what rule the rest of the plant is divided is the subject of a settled literature [1–3]. Parallel compressors on a common discharge manifold pose the same allocation problem with efficiency rather than thermal margin as the currency to be divided [4], and parallel chillers on a common chilled-water loop add a sequencing decision on top of the split [5]. Synchronous generators and grid-forming inverters on a common bus divide load through droop for exactly the reason a header master exists, because more than one integral controller on one frequency or one pressure leaves the division of output indeterminate [6]. In every case the regulation of the shared variable is the easy half. Allocation is the decision that matters, and it is made against constraints that belong to the individual producers and never appear in the shared measurement.

Multi-module reactor plants make the problem urgent. Header pressure control for a plant with several reactors and several loads has been posed and solved as a decoupling problem [7], and

unit-specific load-following across a multi-unit small modular reactor plant has been treated as an optimization over the individual units rather than as a single regulation loop [8]. That work addresses the designer choosing a control structure and presumes that the structure is already accepted as where the answer lies. A student shown a block diagram with one pressure controller has no way to discover why a second one is forbidden. This chapter works with an interactive three-module steam header plant, set out in enough detail to rebuild it from the text, and measures what five control architectures and three allocation rules do to the same plant under the same upsets. Its central idea is the separation of the header problem into a fast regulation component, answered by a bypass valve in seconds, and a slow allocation component, answered by three rate-limited reactors against a capability ceiling that moves as they load. The measured runs show that the two components have different owners and different failure modes, on a plant where a well-regulated variable and a badly run plant coexist and can be watched together. The reactor model and the allocation arithmetic are standard.

> ✦ **DID YOU KNOW? Twelve Reactors and One Refueling Tool**
>
> Small modular reactors are usually planned with several units on one site. One design, of 77 MW electrical modules, can house up to twelve units on a site. Because the site has a single disassembly tool, only one unit can be refueled at a time, a job of ten days [9].

Section 9.2 traces load division on a common header from parallel boilers to multi-module reactor plants. Section 9.3 collects the eighteen-state model, its instrument chain and the constants that reproduce the simulator. Section 9.4 measures the time scales, the coupling between the modules and the capability ceiling that moves with the header, and derives the allocation rule that equalizes margin. Section 9.5 describes the common loop algorithm, the five architectures, the allocator and its spill pass, the two-stage runback and the model predictive controller. Section 9.6 reports the twelve scenarios, from independent control fighting itself to the graded challenge and the predictive controller against coordinated allocation. Section 9.7 sets the plant in its hardware, Section 9.8 states the limitations, Section 9.9 draws out what generalizes, and Section 9.10 summarizes the measured results.

9.2 Background

Dividing one load among several boilers on a common header is one of the oldest problems in industrial process control, and the earliest arrangement was mechanical. A boiler master reduced header pressure to a single air signal that reached every firing rate controller in parallel through a manual bias station on each one. The bias station is the ancestor of every allocation rule that followed, because it decided how the common demand was divided while the master decided only how large the demand was. Only one integral device existed in that scheme, which is the reason it worked and the reason a single pressure master feeding a participation rule is still the recommended arrangement [3]. Where a second was added, one header controller per boiler, the division of firing rate became indeterminate and the plant drifted for reasons no single loop explained.

Dynamic models of the boiler side matured well after the control practice, and the drum boiler acquired a compact and widely reused nonlinear model whose central result is that shrink and swell couple pressure to level and tie the two loops together [10], the inverse response treated in Section 1.4.4 and Chapter 4. The essential feature of that model here is the drum itself, because a recirculating boiler holds a large water and steam inventory and can absorb a load change for a minute or more before its firing rate has to answer. Remove that inventory and the timing of the whole plant changes. A once-through steam generator boils to superheat over its length and holds almost nothing, so feedwater and steam flow are tied together and the unit cannot buffer anything, which is why feedwater control across several once-through generators is treated as a coupled multi-unit problem [11]. Helical-coil geometry is the standard choice inside an integral reactor vessel, and its

thermal-hydraulic analysis shows the moving boundaries between economizer, boiling and superheat regions that a reduced-order model has to lump [12].

Small modular reactors bring the parallel-boiler problem back with the inventory removed. Factory fabrication, staged commissioning and independent dispatch of the individual modules are the economic arguments for the concept, and the same arguments have been made and revised across several deployment cycles [13]. Integral pressurized water reactor designs place a helical-coil once-through steam generator inside a natural-circulation vessel and bring several such modules into one turbine hall, coupled to more than one product through a common steam supply [14]. That coupling is now studied as a system in its own right, with dynamic models built for plants selling electricity and hydrogen from the same steam [15] and for plants that add thermal storage as a further participant on the steam balance [16]. The header in all of those arrangements belongs to none of the products, so multi-reactor, multi-load header pressure control has been posed directly as a decoupling design [7] and multi-unit load following as an allocation across units with individual constraints [8]. Both answers match the practice parallel boilers arrived at, one master and one rule, and both add the constraint the older plants did not have, which is that each producer's capability moves with the shared variable.

> ⧖ HISTORY **Steam Bypass on Demand (2009)**
>
> From 2009 a surplus of baseload generation in one Canadian province led its system operator to require nuclear plants to cut their output when generation ran high and demand ran low. Selected units kept their reactors at full thermal power and cut generator output by up to 300 MW of electricity for short periods, sending the excess steam to the condenser through the steam bypass system. That system had been designed for rare loss of load events and not for frequent use, so its capability had to be evaluated first. Because the plants had several units each, the operating staff could choose which unit to run at reduced power, sparing units with equipment problems. A review published in 2018 counted more than 1500 such operations without an event at one eight-unit station [17].

9.3 Process Model

The simulated plant is three integral pressurized water reactors of 250 MW rated thermal power, each a single vessel standing in a pool with no primary coolant pumps and a helical-coil once-through steam generator wound around the riser inside the vessel. Each module discharges through a main steam isolation valve into one header, which supplies a turbine generator, an industrial process-steam customer, a bypass to the condenser, a thermal storage loop and a relief valve, as shown in Figure 9.1 with every symbol the equations below use. The three modules are deliberately not identical, because a plant one year after commissioning has modules at different burnup with steam generators in different condition, and that difference is what an allocator has to divide.

Table 9.1 classifies what the control system interacts with. The plant is expressed in the units of steam-plant operation, pressures in MPa, temperatures in °C, flows in kg/s, thermal and electrical power in MW, enthalpies in kJ/kg, storage in MWh thermal and time in seconds.

🕮 **Pause.** Try Exercise 9.1 before reading on.

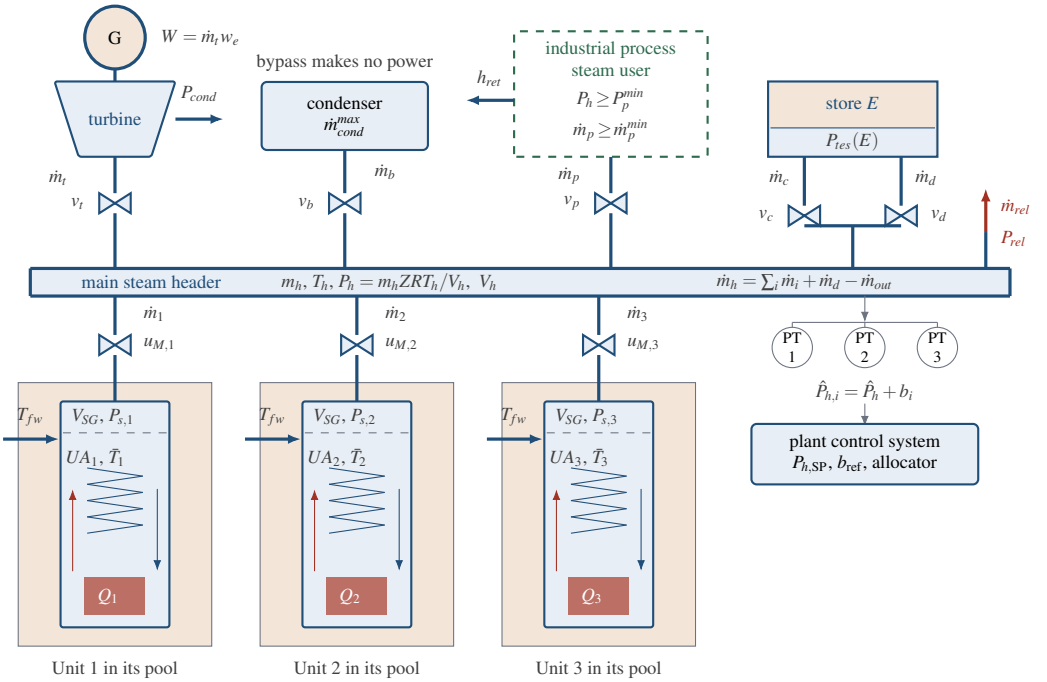


Figure 9.1 Three integral pressurized water reactors on one main steam header, with every symbol used in the equations. Core power Q_i heats a primary loop of average temperature $\bar{T}_i$, which passes $UA_i(\bar{T}_i - T_{sat}(P_{s,i}))$ into a once-through steam generator of steam-side volume V_{SG} at pressure $P_{s,i}$. Flow $\dot{m}_i$ reaches the header through an isolation valve of fractional opening $u_{M,i}$. The header holds mass m_h at temperature T_h in volume V_h, and its pressure P_h sets both the saturation temperature in every steam generator and the flow through every user's valve. Three pressure transmitters read the same header with different calibration biases b_i.

9.3.1 Steam Properties

Saturation temperature, saturated vapor enthalpy and saturated liquid enthalpy are quadratic fits over 1.5 to 6 MPa, the band this header visits, as shown in Equation (9.1) with P in MPa and $x = \ln P$.

$$\begin{aligned} T_{sat} &= 180.30 + 41.81\,x + 6.348\,x^2 \\ h_g &= 2767.5 + 21.03\,P - 3.030\,P^2 \\ h_f &= 768.6 + 170.5\,x + 43.5\,x^2 \\ h_f(T) &= 4.19\,T + 3.4 \times 10^{-4}\,T^2 \\ h(P,T) &= h_g(P) + c_p \max(0,\, T - T_{sat}(P)) \end{aligned} \tag{9.1}$$

where c_p is the superheated heat capacity in kJ/kg/K. The fits reproduce saturation temperature to better than 0.15 °C and saturated vapor enthalpy to about 2 kJ/kg, and every lookup clamps its argument to 0.5 to 8 MPa so that an excursion outside the fitted range returns the edge value rather than an extrapolated one. The equation of state is Equation (9.2), and it serves both the header and each steam generator plenum.

$$P = \frac{mZR(T + 273.15)}{1000\,V} \tag{9.2}$$

where Z is a compressibility factor, $R = 0.4615\,\mathrm{kJ/kg/K}$ is the gas constant of steam, m is mass in kg and V is volume in m^3.

Table 9.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts, and the Range of a CV Is That of Its Setpoint.

Symbol	Role	Nominal	Range
P_h	CV, measured	3.4012 MPa	2.90–3.75 MPa
W	CV, measured	132.7 MW	20–150 MW
$\dot{m}_p$	CV, measured	55.7 kg/s	0–95 kg/s
b (bypass)	CV, mid-ranged	3.0 %	0–20 %
$Q_{i,dem}$	MV	190.0/176.7/163.4 MW	50–260 MW
v_t, v_p	MV	70.6, 54.7 %	0–100 %
v_b	MV	3.0 %	0–60 %
v_c, v_d	MV	0 %	0–100 %
grid load	DV, measured	133 MW	40–150 MW
process request	DV, measured	56 kg/s	0–95 kg/s
f, T_{fw}	DV, measured	60 Hz, 149 °C	±0.6, 100–200
$m_{f,i}$ fouling	DV, unmeasured	1.00/0.93/0.86	0.55–1.05
b_i spread	DV, unmeasured	0.015 MPa	0–0.05 MPa
trip i	DV, event	none	any module
Q_i	state	190.0/176.7/163.4 MW	0–300 MW
$\bar{T}_i$	state	302.9/302.1/301.3 °C	120–400 °C
$P_{s,i}$	state	3.802/3.752/3.704 MPa	0.6–8 MPa
m_h, T_h	state	1724 kg, 296.8 °C	$\geq$ 20, 150–430
E	state	112.5 MWh	0–250 MWh
x_k, five valves	state	valve positions	0–100 %
W_{lag}	state	132.7 MW	≥ 0

9.3.2 Module Side

Each module carries three states. Core thermal power follows a rate-limited demand through one first-order lag, as shown in Equation (9.3).

$$\tau_Q \frac{\mathrm{d}Q_i}{\mathrm{d}t} = Q_{i,dem} - Q_i, \qquad \left|\frac{\mathrm{d}Q_{i,dem}}{\mathrm{d}t}\right| \leq \frac{r_i}{100}\frac{Q_{rate}}{60} \tag{9.3}$$

where τ_Q is the closed-loop core power time constant in seconds, r_i is the module's ramp rate in percent of rating per minute and Q_{rate} is the rated thermal power in MW. Reactivity control, rod worth, xenon and the moderator and Doppler feedbacks are all folded into τ_Q and the rate limit. That rate limit is the constraint the whole application is organized around, because the three modules together can move only 31.3 MW of core power per minute and no controller gain changes the number.

The primary loop is the module's integrator, as shown in Equation (9.4).

$$C_{p1}\frac{\mathrm{d}\bar{T}_i}{\mathrm{d}t} = Q_i - Q_{SG,i}, \qquad Q_{SG,i} = UA_i\left(\bar{T}_i - T_{sat}(P_{s,i})\right) \tag{9.4}$$

where C_{p1} is the heat capacity of coolant, vessel and internals in MJ/°C, $\bar{T}_i$ is the average primary coolant temperature and $UA_i = UA_{SG}\,m_{f,i}$ is the steam generator conductance in MW/°C with $m_{f,i}$ the per-module fouling multiplier. Both sides carry units of MW. The heat sink in (9.4) is the saturation temperature of the module's own steam, which follows the header, so raising header pressure reduces every module's heat removal at the same instant. That single term is why no disturbance on this plant is local, and the ratio $C_{p1}/UA_1 = 180/3.43 = 52.5\,\mathrm{s}$ is the slowest time constant the control system can reach.

The steam side of a once-through generator holds almost nothing, and its pressure is the third state, as shown in Equation (9.5).

$$\begin{aligned} \dot{m}_{gen,i} &= \frac{1000\,Q_{SG,i}}{\max\big(200,\ h(P_{s,i},T_{st,i}) - h_f(T_{fw})\big)} \\ \frac{\mathrm{d}P_{s,i}}{\mathrm{d}t} &= \frac{(\dot{m}_{gen,i} - \dot{m}_i - \dot{m}_{sv,i})\,ZR(T_{st,i}+273.15)}{1000\,V_{SG}} \end{aligned} \tag{9.5}$$

where $\dot{m}_{gen,i}$ is the steam made in kg/s, V_{SG} is the whole steam-side volume of coil, plenum and lead pipe in m^3, $\dot{m}_i$ is the flow to the header and $\dot{m}_{sv,i}$ is safety valve flow. Feedwater flow is set equal to $\dot{m}_{gen,i}$, so there is no feedwater control and no drum level. Outlet superheat is a power law referenced to a design duty rather than an area calculation, as shown in Equation (9.6).

$$\begin{aligned} T_{st,i} &= T_{sat}(P_{s,i}) + \mathrm{clamp}(\theta_i,\,6,\,140) \\ \theta_i &= \Delta T_{SH}\left(\frac{Q_{SH}}{\max(Q_{SG,i},5)}\right)^{0.55} \end{aligned} \tag{9.6}$$

where ΔT_{SH} is the superheat at the reference duty Q_{SH} in MW and θ_i is the unclamped superheat. Superheat falls as duty rises because feedwater flow tracks steam flow against a fixed heat transfer area, and the clamp keeps the correlation inside its fitted band.

Flow from a module to the header is driven by the pressure difference between them, as shown in Equation (9.7).

$$\dot{m}_i = u_{M,i}\,K_{line}\sqrt{10^6\,(P_{s,i}-P_h)\,\rho_i}, \qquad \rho_i = \frac{1000\,P_{s,i}}{ZR\,T^K_{st,i}} \tag{9.7}$$

where $u_{M,i}$ is the fractional isolation valve opening, K_{line} is the line flow coefficient, ρ_i is steam density in $\mathrm{kg/m}^3$ and $T^K_{st,i}$ is the outlet temperature in kelvin. The square root makes the split between modules a hydraulic outcome before it is a control decision, and at the design point each module sits 0.30 to 0.40 MPa above the header. Main steam safety valves are shut in every normal condition and open linearly above their set pressure, as shown in Equation (9.8).

$$\dot{m}_{sv,i} = K_{sv}\max(0,\ P_{s,i} - P_{sv}) \tag{9.8}$$

They are the only outlet an isolated steam generator has for the decay heat still arriving in front of it, which is why an isolated module's pressure settles just above P_{sv} rather than running away.

A scram collapses core power to decay heat and shuts the isolation valve, as shown in Equation (9.9).

$$\begin{aligned} Q_{i,dem} &= 0.065\,Q_{i,0}\left(1 + \frac{t - t_{trip}}{10}\right)^{-0.2} \\ u_{M,i} &= \max\left(0,\ 1 - \frac{t - t_{trip}}{\tau_{MSIV}}\right) \end{aligned} \tag{9.9}$$

where $Q_{i,0}$ is the power the core was making at the trip and τ_{MSIV} is the closure time in seconds. The decay expression is a single-term fit to the Way and Wigner shape of Equation (1.9), giving 6.5 % at once, 4 % after ninety seconds and 2.6 % after a quarter of an hour, and τ_Q is replaced by the faster τ_{scram} while the trip flag is set. A module also trips itself if $\bar{T}_i$ exceeds $\bar{T}_{trip}$.

9.3.3 Header and Users

The header is two states, a mass and a temperature, and its pressure follows from (9.2). The balances are Equations (9.10) and (9.11).

$$\frac{\mathrm{d}m_h}{\mathrm{d}t} = \sum_i \dot{m}_i + \dot{m}_d - \dot{m}_{out} \tag{9.10}$$

$$\begin{aligned} m_h(c_p - ZR)\frac{\mathrm{d}T_h}{\mathrm{d}t} = &\sum_i \dot{m}_i(h_i - u_h) + \dot{m}_d(h_d - u_h) \\ &- \dot{m}_{out}(h_h - u_h) - \dot{Q}_{loss} \end{aligned} \tag{9.11}$$

where h_i is the enthalpy of the steam arriving from module i, h_d is the enthalpy of storage discharge steam, $h_h = h(P_h, T_h)$ is the header enthalpy, $u_h = h_h - ZR(T_h + 273.15)$ is the header specific internal energy in kJ/kg, $\dot{Q}_{loss} = UA_h(T_h - T_{amb})$ is the pipe heat loss in kW with UA_h in kW/°C, and $\dot{m}_{out}$ is the total draw of Equation (9.14). Both sides of (9.10) are in kg/s and both sides of (9.11) are in kW.

Every user is a valve on the same pressure, in the critical-flow form of Equation (9.12).

$$\dot{m}_k = \frac{x_k}{100} K_k \frac{P_h}{\sqrt{T_h + 273.15}}, \qquad k \in \{t, p, b, c\} \tag{9.12}$$

where x_k is the valve position in percent and K_k is the coefficient of the turbine governor, the process valve, the bypass or the storage charge valve. Equation (9.12) is the second coupling term in the model, because it makes every user's draw proportional to header pressure at a fixed valve position. The users therefore self-regulate a little, and the industrial customer's flow falls exactly when its pressure requirement is already in trouble. Storage discharge is driven the other way, as shown in Equation (9.13).

$$\begin{aligned} \dot{m}_d &= \frac{x_d}{100} K_d \max(0,\ P_{tes} - P_h) \\ P_{tes} &= P_{tes}^{lo} + \left(P_{tes}^{hi} - P_{tes}^{lo}\right)\frac{E}{E_{max}} \end{aligned} \tag{9.13}$$

where E is the stored energy in MWh, so the delivery pressure falls with the charge and a nearly empty tank cannot push steam into a healthy header at all. The relief valve is the linear form of (9.8) about P_{rel}, and the total draw is Equation (9.14).

$$\dot{m}_{out} = \dot{m}_t + \dot{m}_p + \dot{m}_b + \dot{m}_c + \dot{m}_{rel} \tag{9.14}$$

Generator output is a polytropic expansion to the condenser through one combined efficiency, followed by a first-order output lag, as shown in Equation (9.15).

$$\begin{aligned} w_e &= \eta_t\, c_p(T_h + 273.15)\left[1 - \left(\frac{P_{cond}}{P_h}\right)^{\kappa}\right] \\ \tau_{gen}\frac{\mathrm{d}W_{lag}}{\mathrm{d}t} &= \frac{\dot{m}_t w_e}{1000} - W_{lag} \end{aligned} \tag{9.15}$$

where w_e is the specific work in kJ/kg, η_t absorbs the isentropic, wetness, mechanical and generator losses, and $\kappa = (\gamma - 1)/\gamma$. At the design point (9.15) returns 783 kJ/kg, so 169.5 kg/s through the governor valve makes 132.7 MW. Process heat and condenser loading are Equation (9.16).

$$\dot{Q}_p = \frac{\dot{m}_p\,(h_h - h_{ret})}{1000}, \qquad L_{cond} = \frac{0.9\,\dot{m}_t + 1.35\,\dot{m}_b}{\dot{m}_{cond}^{max}} \tag{9.16}$$

where h_{ret} is the condensate return enthalpy. The coefficients in (9.16) state that a kilogram of bypass steam has done no work on the way and loads the condenser about a third harder than a kilogram of

turbine exhaust, so the bypass is not a free place to put steam. Storage charge and discharge move the energy state through two fixed efficiencies, as shown in Equation (9.17).

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \frac{1}{3600}\left[\frac{\dot{m}_c\,(h_h - h_f(T_{fw}))\,\eta_c}{1000} - \frac{\dot{m}_d\,(h_d - h_f(T_{fw}))}{1000\,\eta_d}\right] \tag{9.17}$$

where η_c and η_d are the recovery efficiencies and 3600 converts MW to MWh/s. Storage steam is barely superheated at $\Delta T_{SH,d}$ above saturation, because a kettle boiler makes saturated steam, so a hard discharge cools the header. Each of the five valves follows its command through a first-order stroking lag, as shown in Equation (9.18).

$$\tau_{V,k}\,\frac{\mathrm{d}x_k}{\mathrm{d}t} = u_k - x_k \tag{9.18}$$

9.3.4 Instrument Chain

The controller reads three instruments, each a first-order filter of the true value written with the exact zero-order-hold update of Equation (9.19).

$$\hat{y}_k = \alpha\,\hat{y}_{k-1} + (1-\alpha)\,y_k, \qquad \alpha = \mathrm{e}^{-\Delta t/\tau} \tag{9.19}$$

with $\tau = \tau_{PT}$ for header pressure, τ_{WT} for generator output and τ_{FT} for process steam flow. Optional colored noise is added to the header pressure before the filter, produced by passing a Gaussian sequence through a first-order filter and rescaling so that the stationary standard deviation equals the specified value. The three module controllers do not read the same number, as shown in Equation (9.20).

$$\hat{P}_{h,i} = \hat{P}_h + b_i, \qquad b = [\,\sigma_b,\ 0,\ -\sigma_b\,] \tag{9.20}$$

where σ_b is the transmitter calibration spread. Its default of 0.015 MPa is four tenths of one percent of a 3.4 MPa reading, which is a good calibration and not a fault, and Section 9.6 measures what it does.

9.3.5 Parameters, Initialization and Numerics

The constants are listed in Tables 9.2 and 9.3, of which four are hidden and never displayed. Every other constant is exposed as an adjustable parameter, and every disturbance in Table 9.1 as a slider or a discrete event.

Every scenario opens at a genuine steady state, and finding it is not a marching problem, because raising one module's demand raises its steam generator pressure, which raises the header, which reduces every other module's flow and every user's draw at the same time. The equilibrium is solved as the two residuals of Equation (9.21) in the two header states, with each module's steam pressure found by bisection inside.

$$\begin{aligned} r_1(P_h,T_h) &= \sum_i \dot{m}_i + \dot{m}_d - \dot{m}_{out} \\ r_2(P_h,T_h) &= \sum_i \dot{m}_i\,(h_i - u_h) + \dot{m}_d\,(h_d - u_h) \\ &\quad - \dot{m}_{out}\,(h_h - u_h) - \dot{Q}_{loss} \end{aligned} \tag{9.21}$$

The outer solve is a damped Newton iteration on a numerical 2×2 Jacobian, relaxed to 55 % for the first four passes and converged to $|r_1| < 10^{-10}$ and $|r_2| < 10^{-7}$. The inner bisection drives $\dot{m}_{gen,i} - \dot{m}_i - \dot{m}_{sv,i}$ to zero over 70 halvings, which is monotone in $P_{s,i}$ and therefore safe, and with the isolation valve shut the root lands just above P_{sv}, where an isolated module belongs. The same solver supplies the manual aid that sets the three demands to those that put the header on setpoint, using a secant iteration on total demand rather than a fixed gain because the header gain moves with the

Table 9.2 Module and Header Parameters. Hidden Constants Are Never Displayed.

Symbol	Meaning	Value	Unit
Q_{rate}	module rated power	250	MW
Q_i^{lic}	licensed power, U1/U2/U3	250/240/232	MW
r_i	ramp rate, U1/U2/U3	5.0/3.0/4.5	%/min
UA_{SG}	clean SG conductance	3.43	MW/°C
C_{p1}	primary heat capacity	180	MJ/°C
τ_Q	core power lag	14	s
τ_{scram}, τ_{MSIV}	scram, valve closure	1.6, 6	s
$\bar{T}_{max}$, $\bar{T}_{trip}$	coolant limit, trip	315, 324	°C
V_{SG}	SG steam-side volume	30	m^3
K_{line}	steam line coefficient	0.032473	—
P_{sv}	safety valve set pressure	4.60	MPa
ΔT_{SH}	superheat at reference duty	57	K
V_h	header volume	120	m^3
P_h^0, P_{rel}	design, relief pressure	3.40, 3.85	MPa
UA_h, T_{amb}	heat loss, ambient	2.2, 25	—
Z, c_p	compressibility, heat cap.	0.90, 2.65	—
δ_P	pressure tolerance band	0.08	MPa
$\tau_{PT}, \tau_{FT}, \tau_{WT}$	transmitter responses	0.4, 1.2, 0.8	s
hidden			
K_{sv}	safety valve capacity	150	kg/s/MPa
Q_{SH}	reference duty for (9.6)	176.5	MW
K_{rel}	relief valve capacity	400	kg/s/MPa
f_0	nominal frequency	60	Hz

operating point. It is 6.8 kPa per MW at the design point and 13.6 kPa per MW on a plant at half load held at the same pressure.

The engine of Section 1.6.1 integrates the model by fourth-order Runge-Kutta at a fixed 0.05 s step, chosen because the fastest state is a steam generator pressure with a time constant near 1 s. The control system executes every 0.25 s, a realistic distributed control system scan, the historian logs four samples per simulated second, and the noise seed is 20260906. Every state is clipped after each step to the range given in Table 9.1.

9.3.6 Assumptions

The model is detailed on the steam balance and thin everywhere else, and each simplification fails somewhere.

No neutronics. Equation (9.3) is a first-order lag on a rate-limited demand, so reactivity control, rod worth, xenon and the moderator and Doppler feedbacks are all inside one 14 s time constant. That is adequate for a load-sharing study and useless for anything about reactivity, so the application cannot show a reactivity-limited maneuver or a xenon-constrained restart. *One lumped steam generator conductance.* A real helical-coil once-through unit has economizer, boiling and superheat sections whose boundaries move with load [12], and its response to a feedwater step is genuinely different from the single lag of (9.4). *Feedwater flow is not independent.* It equals the steam made, so there is no feedwater control, no drum level because there is no drum, and no feedwater pump as a constraint, and coupled feedwater control across several once-through generators is itself a substantial problem [11]. *Natural circulation is asserted rather than modeled.* There is no primary flow state and therefore no flow-limited power, which is a real constraint on a natural-circulation module.

Steam is an ideal gas. Equation (9.2) carries $Z = 0.90$ and a constant superheated heat capacity, and outside roughly 1.2 to 6.5 MPa the property fits are extrapolation and every lookup clamps. *The header is a single lumped volume*, with no pressure drop, no distinct conditions at each take-off and

Table 9.3 User, Storage, Valve, Runback and Economic Parameters. Where a Row Carries Several Symbols the Values Are Listed in the Same Order and the Units Are Those Named in the Meaning Column or Given in the Text.

Symbol	Meaning	Value	Unit
K_t, K_p, K_b	governor, process, bypass valve	1685.31, 714.98, 702.21	—
η_t, κ	turbine efficiency, exponent	0.683, 0.23	—
P_{cond}, τ_{gen}	condenser pressure, output lag	0.007, 1.5	—
$\dot{m}_{cond}^{max}$	condenser steam capacity	200	kg/s
h_{ret}	condensate return enthalpy	440	kJ/kg
P_p^{min}, $\dot{m}_p^{min}$, t_{grace}	industrial user limits	2.90, 34, 45	—
E_{max}, $E(0)/E_{max}$	storage capacity, initial charge	250, 45	—
P_{tes}^{lo}, P_{tes}^{hi}	delivery pressure empty, full	3.05, 5.40	MPa
K_c^{chg}, K_d	charge, discharge coefficient	280.89, 42.5	—
η_c, η_d	charge, discharge recovery	0.96, 0.94	—
$\Delta T_{SH,d}$	superheat of storage steam	10	K
$\tau_{V,t}$, $\tau_{V,b}$, $\tau_{V,p}$, $\tau_{V,c}$	valve stroking times	0.6, 2.0, 4.0, 5.0	s
R, W_{rated}	governor droop, rating	5 %, 150	—
δ_{fast}, δ_{trim}	runback pick-up, trim band	0.10, 0.02	MPa
K_{shed}, $K_{shed,p}$	runback droop gains	220, 260	—
K_{trim}, $K_{trim,p}$	runback trim rates	0.60, 7	—
c_E, c_Q	electricity, steam price	42, 28	$/MWh
c_{int}, c_{trip}	interruption penalty, loss cost	9000, 40 000	—

no condensate, and two users at opposite ends of a real header do not see the same pressure, which would make the multi-controller failure of Section 9.6 worse rather than better. *The turbine has no shaft dynamics*, so a grid load rejection here is a steam-balance event rather than the mechanical one it would really be. *Storage is a state of charge*, with no thermocline, no temperature limits and no self-discharge, so the application overstates how neatly a store can be dispatched. *Two constraints stop a module, and the economics are four prices.* Only the average coolant temperature limit and the licensed power appear, where a real protection system carries a dozen more, and the prices exist to make the trade between shedding electricity and shedding a customer concrete rather than to value a contract.

9.4 Coupling and Capability

9.4.1 A Fast Buffer Between Slow Plants, and Six Actuators

Three time scales decide everything the control system can do, and all three were measured on the shipped model at the design point. The header holds 1724 kg of steam and passes 228.2 kg/s, which is 7.55 s of residence, and the steam side of a module holds about 470 kg against an average 76.1 kg/s, which is 6.18 s, so the inventory between a reactor core and a customer's reboiler is under

fourteen seconds of throughput. Against that, the primary loop time constant of (9.4) is 52.5 s and the aggregate ramp capability of (9.3) is 31.3 MW/min. The coupling is an order of magnitude faster than anything thermal that could answer it, and that ratio is the whole control problem. A recirculating boiler with a drum would hold thirty times the module inventory and buffer the header for a minute, which is why the parallel-boiler practice of Section 9.2 could tolerate a slow master [10].

Header pressure has six actuators, and every architecture that follows is a decision about which of them answers a given error. Stepping each actuator from the design point in manual mode, 10 s into a run, and running to a new steady state gives the steady gain and the time to 63 % of the final pressure change, counted from the step. A station moved before Run only sets the steady state the plant starts from. Unit 1 from 190 to 210 MW moves the header by $+0.1367$ MPa with $t_{63} = 195$ s, and all three demands together by 30 MW move it by $+0.2052$ MPa. The governor valve from 70.6 to 78.0 % moves it by -0.2578 MPa in 32 s, and the bypass from 3 to 13 % by -0.1501 MPa in 36 s.

Two readings of those numbers matter. The header gain is 6.84 kPa per MW of module duty from the single-module step and 6.84 kPa per MW from the three-module step, so it is a property of the header rather than of which module moved. The two valve times are about 4 s apart, because a pressure change on this header settles over about half a minute whatever valve moved, so what separates the actuators is the speed at which the flow itself changes. A bypass valve strokes in 2 s and passes 1.00 kg/s per percent, so closing a bypass resting at 3 % puts 3 kg/s back into the header inside two seconds. A reactor demand takes 96 s to move 20 MW at 12.5 MW/min before the 14 s power lag and the 52.5 s primary loop have even started, and the measured t_{63} of 195 s is set by those three in series. Everything fast on this plant is a valve and everything free is a reactor.

The single-module step also measures the coupling directly. Raising Unit 1 from 190 to 210 MW raises its own steam generator pressure by 0.204 MPa and, through the header, raises Unit 2 by 0.127 MPa and Unit 3 by 0.128 MPa. The flows of the two untouched modules dip while the header rises and then return to within 0.01 kg/s of where they started, because in steady state a module's steam production is its own heat balance and nothing else. Their margin does not return. The saturation temperature underneath each has risen with its steam pressure, so (9.4) settles at a higher $\bar{T}_i$ for the same duty, and Unit 2 loses 1.95 °C of margin while Unit 3 loses 1.99 °C without either being touched. Loading one module spends thermal margin belonging to all three.

9.4.2 The Ceiling Is Not the Rating, and It Moves

What actually stops a module is the average coolant temperature limit, not the licensed power, as shown in Equation (9.22).

$$Q_i^{max} = \min\Big(Q_i^{lic},\ UA_i\,(\bar{T}_{max} - T_{sat}(P_{s,i}))\Big) \tag{9.22}$$

The physical ceiling in Equation (9.22) is distinct from what the controller can estimate. The allocator and the demand clamps use

$$\begin{aligned}\widehat{Q}_i^{max} = \max\big(0, \min\big(&Q_i^{lic},\\ &\widehat{Q}_{SG,i} + UA_{i,0}(\bar{T}_{max} - \bar{T}_{i,meas})\big)\big).\end{aligned} \tag{9.23}$$

The primary heat balance estimates $\widehat{Q}_{SG,i}$ from the commanded core response and the change in measured average coolant temperature. Conductance, primary heat capacity and core response times are held at their declared design values. Hidden fouling and steam generator pressure never enter the allocator or clamp. The same estimate limits all five structures. With changed conductance, the estimate can differ from the physical ceiling during a transient and approaches it as the measured temperature approaches its limit. In a frozen-instrument regression, a hidden fouling and parameter change moved the former clamp by 57.31 MW and moves the corrected clamp by zero.

At the design point the three ceilings are 232, 218 and 204 MW against licensed powers of 250, 240 and 232 MW, and the plant has 123 MW of thermal reserve. Every ceiling falls as the header

rises, and it falls faster for a fouled module because a smaller UA_i converts the same duty into a larger temperature rise. The consequence is a step test that stops short of its own request. Asking Unit 1, licensed for 250 MW, to make 230 MW produces a demand ramp at 12.5 MW/min that overshoots to 222.6 MW at 240 s with $\bar{T}_1 = 311.4\,^{\circ}\mathrm{C}$ and then settles back to 217.3 MW with $\bar{T}_1$ resting exactly on the 315.00 °C limit, reading 217.5 MW at the 480 s instant the scenario removes the step. Loading it raised its own steam generator pressure from 3.802 to 4.081 MPa and the header from 3.401 to 3.588 MPa, and every one of those kilopascals raised the saturation temperature the primary has to work above. The module ends 12.7 MW short of the request and 32.7 MW short of its license, held there by a constraint that appears in no slider.

9.4.3 Why Allocation by Available Duty Equalizes Margin

The allocation rule follows from (9.4) and (9.22) directly. In steady state $Q_i = UA_i(\bar{T}_i - T_{sat})$, so the margin of module i is Equation (9.24).

$$\bar{T}_{max} - \bar{T}_i = \bar{T}_{max} - T_{sat} - \frac{Q_i}{UA_i} = \frac{Q_i^{max} - Q_i}{UA_i} \tag{9.24}$$

where Q_i^{max} takes the conductance branch of (9.22). Dividing a total demand D in the ratio of the ceilings, $Q_i = D\,Q_i^{max} / \sum_j Q_j^{max}$, and substituting into (9.24) gives Equation (9.25).

$$\bar{T}_{max} - \bar{T}_i = (\bar{T}_{max} - T_{sat}) \left(1 - \frac{D}{\sum_j Q_j^{max}} \right) \tag{9.25}$$

The right side of (9.25) carries no index i, so the rule leaves every module the same number of degrees from its own limit, which is the only sense in which a load share on this plant can be called fair. The equality is exact only where every module boils against the same saturation temperature, and because the three generators sit at slightly different pressures the measured margins under coordinated control at the design demand are 12.8, 12.9 and 13.0 °C, a spread of 0.25 °C. Evaluating the algebra on Unit 1 gives $(232/3.43)(1 - 530/654) = 12.8\,^{\circ}\mathrm{C}$, which matches the measurement. The rule also has a failure mode built into (9.22), because a module whose licensed power binds no longer has a ceiling proportional to its conductance and the equality is lost.

9.5 Control Architectures

9.5.1 The Common Loop Algorithm

Every loop is the shared velocity-form proportional-integral controller of Section 1.6.2, executed every $\Delta t_c = 0.25\,\mathrm{s}$ and written in this chapter's symbols as Equations (9.26) to (9.29).

$$e_k = a\left(y_{\mathrm{SP},k} - \hat{y}_k\right), \qquad P_k = K_c e_k \tag{9.26}$$

$$s_k = s_{k-1} + P_k - P_{k-1} \tag{9.27}$$

$$c_k = u_{bias} + s_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k, u_{min}, u_{max}) \tag{9.28}$$

$$I_k = I_{k-1} + \frac{K_c}{\tau_I} e_k \Delta t_c + \frac{\Delta t_c}{T_t} (u_k - c_k) \tag{9.29}$$

where $a = \pm 1$ is the action, s_k is the running sum S_k of Equation (1.14), accumulated without limit, c_k is the demanded output v_k of Equation (1.15), u_k is the output applied, and T_t is the anti-windup tracking time, which is τ_I on every loop here. No loop in this application has a derivative time and none is exposed, so the 0.3 s derivative filter configured on the two user loops is inert. Anti-windup is the back-calculation of Section 1.6.2, and s_k is never truncated [18]. The reactor ramp limit of (9.3) is applied after the master and valve position controller outputs and lies outside this mechanism, so those integrals keep integrating while the reactors ramp toward their targets, bounded only by the sums of the module floors and ceilings. Manual tracking and the copy of each automatic output to its

Table 9.4 The Eight Implemented Loop Types, Their Pairings, Directions and Default Tunings. Action +1 Is Reverse Acting and −1 Is Direct Acting. Sums Run Over the Modules Still in Service.

CV → MV	Limits	K_c	τ_I (s)	a
$P_h \to D$ (master)	$\sum Q_i^{min}$–$\sum Q_i^{max}$	320	240	+1
$P_h \to Q_{i,dem}$ (per module)	50–Q_i^{max}	950	120	+1
$P_h \to v_b$ (bypass)	0–60 %	260	25	−1
$b \to D$ (position)	$\sum Q_i^{min}$–$\sum Q_i^{max}$	3.0	240	+1
$W \to v_t$	0–100 %	0.35	20	+1
$\dot{m}_p \to v_p$	0–100 %	0.50	15	+1
$P_h \to v_d$ (discharge)	0–100 %	700	40	+1
$P_h \to v_c$ (charge)	0–100 %	600	90	−1

manual station work as Section 1.6.2 describes, every loop not in service in the selected architecture is tracked the same way, and so bumpless transfer and architecture switching both fall out of the tracking. The direction of each loop follows the sign of its process gain and appears as the ISA action flag in Table 9.4, in the convention of Section 1.3.

The two pressure-to-demand loops in Table 9.4 are reverse acting because more reactor power raises the header, and the bypass loop is direct acting because opening the bypass lowers it. The per-module loops and the master work against the same process gain. Section 9.4 measures 6.84 kPa per MW from the single-module step and 6.84 kPa per MW from the three-module step, because the header answers the total duty and not the identity of the module that supplied it, so a per-module controller needs the master gain and no more. The default per-module gain of 950 MW/MPa against a master gain of 320 MW/MPa is three times that. Whatever value is chosen, three such controllers act on the same error at the same instant, so the loop gain the plant sees is three times the loop gain any one of them was tuned for. That multiplication, and not the numerical default, is the structural fault.

9.5.2 Five Architectures

The architecture selector chooses among five answers a plant can give to the question of who moves when header pressure moves.

Manual. The operator sets three module demands and five valve positions, and two manual aids compute the demands that put the header on setpoint and the valve positions that deliver the user setpoints at the present pressure.

Independent. Each module runs its own controller from Table 9.4 on its own transmitter reading $\hat{P}_{h,i}$ of (9.20). Three integral controllers share one measurement, so the loop gain is three times what any of them was tuned for, and because the three transmitters disagree the three integrators chase three targets that cannot all be reached.

Master-slave. One pressure master produces a single total thermal demand D and a participation rule divides it, so only one integrator sees the pressure error and the conflict stops. The reactors remain inside the pressure loop, so the speed of the loop is the speed of the reactors.

Coordinated. Three layers act on different time scales. The bypass controller takes the pressure error directly and its valve strokes in 2 s, a valve position controller then moves the total demand D to walk the bypass back toward its resting position b_{ref} over about four minutes, and the allocator divides D by available duty. Where thermal storage is in service the discharge controller is offered the setpoint 0.04 MPa low and the charge controller 0.04 MPa high, so the store acts only outside a dead band and never fights the master. The valve position controller watches the bypass controller's unsaturated demand c_k of (9.28) rather than the clamped valve, because a bypass clamped shut at zero would otherwise report zero no matter how much steam was missing and the slow layer would never learn it needs to make more. Volume 4, Chapter 10 builds the same kind of hierarchy for a

power block, with a turbine bypass that acts in seconds and discards what it removes, mid-ranged against slower actuators that cost nothing.

Pause. Try Exercise 9.3 before reading on.

MPC. One optimization takes the place of the three coordinated layers, setting every second the ramp of each module and the bypass command over the next two minutes, with the coordinated loops held in reserve as its fallback. Section 9.5.5 writes it out.

9.5.3 The Allocator and Its Spill Pass

Given a total demand D and a participation shape s_i, the allocator computes the per-module target of Equation (9.30).

$$Q_i^{tgt} = \mathrm{clamp}\left(\frac{D\, s_i}{\sum_{j \in \mathscr{F}} s_j},\; Q_i^{min},\; Q_i^{max} \right), \qquad i \in \mathscr{F} \tag{9.30}$$

where $\mathscr{F}$ is the set of modules that have not tripped, $Q_i^{min} = \min(0.20\, Q_{rate},\, Q_i^{max})$ and Q_i^{max} is (9.22). The clamping discards demand, so the calculation runs for up to four passes, each of which sums the spill $\sum_i (Q_i^{raw} - Q_i^{tgt})$, removes the modules now against a limit from $\mathscr{F}$, and redistributes the remainder over those that still have room. A share a module cannot take is not a share, it is a shortfall some other module has to be told about, and an allocator that does not spill is the second most common way plantwide load sharing fails after not having one at all. Three shapes are implemented. Equal gives $s_i = 1$. Margin-proportional gives $s_i = Q_i^{max}$, which is the rule of (9.25). Base and swing holds the base-loaded modules at a fixed demand and gives the first module still running the whole of the variation, which puts the entire ramp-rate burden on one reactor. The target is applied through the rate limit of (9.3), so the allocator decides the destination and the ramp rate decides the path. Volume 4, Chapter 12 divides one demand among three actuators with different prices and limits, the same allocation problem on a vehicle.

9.5.4 The Two-Stage Runback

Neither the architecture nor the allocator creates steam, so a plant that has lost a module needs a way to stop selling something. The runback is two proportional stages on the same measurement, both one-way, as shown in Equations (9.31) and (9.32).

$$\begin{aligned} d_{fast} &= \max\left(0,\; P_{h,\mathrm{SP}} - \delta_{fast} - \hat{P}_h\right), \\ d_{slow} &= \max\left(0,\; P_{h,\mathrm{SP}} - \delta_{trim} - \hat{P}_h\right) \\ \Delta W &= \mathrm{clamp}\left(-K_{shed}\, d_{fast} + \Sigma_{trim},\; -(W_{rated} - 20),\, 0\right) \end{aligned} \tag{9.31}$$

$$\Sigma_{trim} \mathrel{-}= K_{trim}\, d_{slow}\, \Delta t_c \tag{9.32}$$

where ΔW is a bias added to the generator load setpoint. The fast stage is proportional and instantaneous, so it arrests a falling header and returns the load as the pressure recovers, and it is the only thing on the plant that can answer a module trip inside the industrial customer's 45 s ride-through. The slow trim accumulates while the header is outside a narrow dead band and never restores load, because after a trip the surviving reactors reach their ceilings with the pressure still short of setpoint and no master action closes a gap where the steam does not exist. A runback is a latch, and recovering the megawatts is an operator decision made on the load setpoint. The process variant is identical with $K_{shed,p}$ and $K_{trim,p}$ acting on the process flow setpoint, floored a little above $\dot{m}_p^{min}$ so that the runback cannot itself shut the customer down. Chapter 5 meets the opposite event, a load that disappears, with a reactor runback among its last layers.

9.5.5 Model Predictive Control

The fifth architecture is a model predictive controller (MPC), a receding-horizon quadratic program [19, 20] in the arrangement of Section 1.6.3, posed every $T_s = 1$ s, each fourth control scan,

over $N = 120$ steps grouped into twelve move blocks that lengthen from a single step near the present to 75 steps at the far end. It decides four inputs, the ramp rate v_i in MW/s of each running module's demand and the bypass command u_b, and between decisions each demand advances by $v_i \Delta t_c$ per scan. Because the decision is a rate and not a level, the limit of (9.3) becomes the exact hard bound $|v_i| \le \bar{v}_i = r_i Q_{rate}/6000$ on every step of the plan whatever the block lengths, and u_b is held hard to 0–60 %.

The controller's model is the balance of Section 9.3 written out again with the declared values of Tables 9.2 and 9.3, the declared fouling multipliers 1.00, 0.93 and 0.86 and 149 °C feedwater, never the run's, so a fouled generator or an edited parameter is model mismatch. Its 17 states are the three demands, core powers, coolant temperatures and steam generator pressures of (9.3) to (9.7), the header mass and temperature of (9.10) and (9.11), the bypass position, the header transmitter filter of (9.19), and one users' draw w_u in kg/s standing for the turbine, process and storage flows and held constant over the horizon. Each decision linearizes it by central differences about the present estimate and discretizes it with an exact hold and its affine term. The turbine and process loops, the run's fouling and feedwater temperature, transmitter noise and bias, the runback, the trips and the condenser are left out. Stepped open loop from the design point for 130 s, the model overstates the header's response by about a third, 0.075 MPa against the plant's 0.056 MPa for +10 MW on Unit 1 and −0.190 against −0.151 MPa for +5 % of bypass, because the users' valves pass more steam as the header rises until their own loops close them back through (9.12).

Each plan minimizes Equation (9.33).

$$\begin{aligned} \mathscr{J} = \sum_{k=1}^{N} \Big[& q_P \big(\hat{P}_{h,k} - P_{h,\mathrm{SP}}\big)^2 + q_s \sum_i \Big(Q_{i,k} - \sigma_i \sum_j Q_{j,k} \Big)^2 \\ & + \sum_i \frac{\lambda\, v_{i,k}^2}{\bar{v}_i^2} + q_b \big(u_{b,k} - b_{ref}\big)^2 \Big] + \sum_{\ell} \big(w_\varepsilon \varepsilon_\ell^2 + w_\varepsilon \varepsilon_\ell \big) \\ & + \sum_j \Big[\sum_i \frac{\lambda\, \Delta v_{i,j}^2}{\bar{v}_i^2} + s_b\, \Delta u_{b,j}^2 \Big] \end{aligned} \tag{9.33}$$

Here $\hat{P}_{h,k}$ is the predicted transmitter reading, Δ is a change between blocks, and σ_i divides the running total in the ratio of each module's declared available duty $UA_i(\bar{T}_{max} - T_{sat}(\hat{P}_{s,i}))$ capped at its licensed power, the rule of (9.25) evaluated on declared values. The weights are $q_P = 10^4$ per MPa^2, $q_s = 0.01$ per MW^2, $q_b = s_b = 0.02$ per %^{2} and the ramp price $\lambda = 3$, where $\lambda = 1$ would price a full-rate ramp held for a second like 0.01 MPa of header error held for a second. One slack ε_ℓ per softened output, judged every second step and priced at $w_\varepsilon = 10^4$ per °C or MW both quadratically and linearly, softens two limits on each running module, the predicted coolant temperature $\bar{T}_{i,k} \le \bar{T}_{max}$ computed with the declared conductance and the demand between $0.2\,Q_{rate}$ and Q_i^{lic}. A temperature estimate already above $\bar{T}_{max}$ is brought back along $\bar{T}_{max} + (\hat{\bar{T}}_{i,0} - \bar{T}_{max})\,\mathrm{e}^{-kT_s/60\,\mathrm{s}}$, where $\hat{\bar{T}}_{i,0}$ is the present estimate. Each plan is a dense problem in 48 decisions solved by a dual active-set method [21], in a mean 3.80 ms in Chrome on scenario 11, 0.38 % of one core at real time.

A steady-state Kalman filter on the same model, linearized at the declared design point at the 0.25 s scan, updates the estimate every scan under every architecture from the filtered header transmitter alone, the only instrument the pressure loops read, with a noise standard deviation of 0.002 MPa and w_u a random walk of variance $\rho_K (0.05\,\mathrm{kg/s})^2$ per scan, $\rho_K = 1$, which removes steady offset [22]. Neither the MPC nor Coordinated has a load feedforward. One tuning pass on scenario 11 over nine values of λ from 0.01 to 100 and six of ρ_K from 1 to 10^5 admitted only settings that rejected no plan, lifted no relief, kept the customer, tripped no module beyond the scripted one and moved neither the modules nor the bypass more than Coordinated did on the same day, and $\lambda = 3$ with $\rho_K = 1$ was the only one admitted.

Table 9.5 Scenarios and What Each Exercises

#	Scenario	Mechanism	Kind
1	Three reactors and one pipe	header gain, coupling, actuator speed	Experiment
2	One module, one step	rate limit, moving ceiling	Experiment
3	Independent pressure control	three integrators, one measurement	Demonstration
4	Master–slave load sharing	one integrator, reactor-limited speed	Experiment
5	Coordinated allocation	mid-ranging, the resting position	Experiment
6	A large turbine load ramp	fast actuator in the loop	Experiment
7	The chemical plant starts another train	interaction on a common pressure	Experiment
8	The store absorbs the ramp	a fourth participant, and its limits	Experiment
9	Equal megawatts is not equal margin	allocation against a constraint	Experiment
10	Unit 1 scrams	N−1 dispatch and the runback	Demonstration
11	Three reactors and one plan	one optimization against three layers	Experiment
12	Ride through the trip	the dispatch decision, one event	Challenge

The coordinated loops track the applied commands every scan with no preload, so a plan that the solver does not return optimal, finite and inside its bounds is discarded, and Coordinated then holds the demands and the bypass, taking them over with no bump and keeping them until the solver again returns a valid plan. The scram, decay heat and isolation of (9.9), the high-temperature trip, the relief and the runback read no architecture and act above the plan, which after a scram carries the tripped module's decay heat and shut valve as known inputs, pins its ramp at zero and goes on planning the survivors. Every applied demand also passes through the clamp of (9.30) at the ceiling of (9.22), the same clamp the other architectures apply, and that clamp is the estimate of Section 9.4.2, built from the measured coolant temperatures, which follow a fouled module where the plan's declared model does not. The simulated plant, its instruments and the controller's model are thus three separate things.

9.6 Results

Twelve scenarios are arranged as one laboratory sequence, listed in Table 9.5 with the mechanism each exercises. Each fixes the starting mode, the architecture, the split rule, the tuning constants, the disturbances and any timed events, and every figure below was measured by driving the shipped model headlessly from those definitions at the 0.05 s physics step. Each case is credited by the completion rule of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

9.6.1 Independent Control With No Disturbance

Scenario 3 holds the plant at the design point with every module running its own pressure controller and nothing disturbing it. The three module powers separate immediately and do not come back, as shown in Figure 9.2. At 600 s the modules are making 132, 202 and 198 MW, a spread of 70.5 MW, and at 1200 s they are at 120, 212 and 198 MW, a spread of 91.7 MW, while header pressure over the whole run stays between 3.385 and 3.401 MPa, a band of 0.0162 MPa inside a tolerance of ±0.08 MPa. The scorecard sees the divergence only through the margin spread, which passes 8 °C at 135 s.

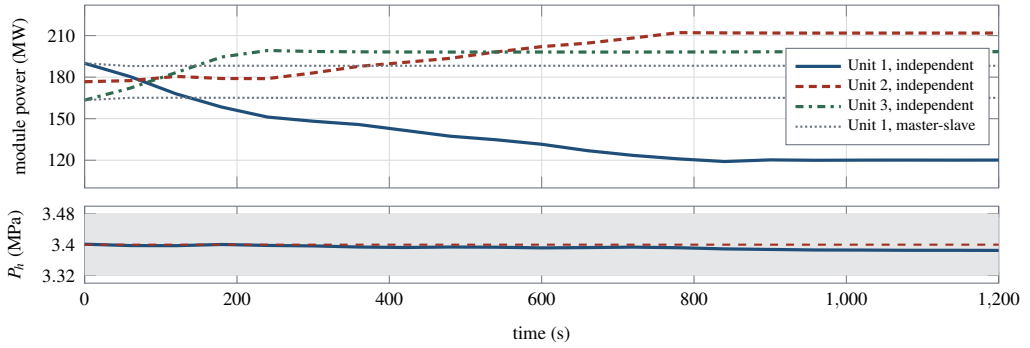


Figure 9.2 Independent per-module pressure control diverging with no disturbance applied, scenario 3. Upper panel: Unit 1 solid, Unit 2 dashed and Unit 3 dash-dotted under independent control, and Units 1 and 3 under master-slave as gray dotted lines. The three module powers separate to 70.5 MW at 600 s and 91.7 MW at 1200 s and stop only when Units 2 and 3 reach their thermal ceilings, while master-slave on the same plant holds a constant 23.2 MW spread. The lower panel shows header pressure over the same run (solid) against its setpoint (dashed) and the shaded ±0.08 MPa tolerance band.

> **TRY IT IN THE APP** **3 · Independent Pressure Control, Fighting Itself**
>
> Load scenario 3 and press Run, at 20× or 60× speed. Nothing disturbs the plant. Watch the three module powers on the Modules tab walk apart, 70.5 MW at ten minutes and 91.7 MW at twenty. Then look at the header pressure trend, which never leaves its band.

The mechanism is (9.20). Unit 1 reads the header 0.015 MPa high, so its integrator winds down, Unit 3 reads it 0.015 MPa low, so its integrator winds up, and Unit 2, whose transmitter happens to be correct, holds the balance. The spread peaks at 93.5 MW at 820 s and the powers then stay within about 1 MW of their final values rather than walking on forever, for a reason worth noticing. Header pressure settles at 3.385 MPa, which is exactly the value Unit 1's transmitter reports as 3.400 MPa, so Unit 1's error reaches zero and its integrator stops, while Units 3 and 2 keep winding up until they reach the Q_i^{max} limits of (9.22), at 196 and 758 s, and stop at 198.4 and 211.8 MW. The final state is a plant with two modules pinned against their thermal ceilings, one at 120 MW, and a controlled variable that never left its band.

Pause. Try Exercise 9.2 before reading on.

Master-slave on the identical scenario holds a spread of 23.2 MW and a header band of 0.0018 MPa, and that 23.2 MW is not a defect but what the margin-proportional rule of (9.25) intends, because the three modules genuinely have different capability. Removing the transmitter spread does not repair the independent architecture, it changes the symptom. With $\sigma_b = 0$ the same three controllers drive header pressure between 2.693 and 3.917 MPa, a band of 1.225 MPa, because three integral controllers acting on one error deliver three times the loop gain any of them was tuned for. No tuning repairs either version, and that is the point. One controlled variable gets one integrator, and the split is a separate decision made by a separate mechanism [1, 3]. Volume 6, Chapter 7 makes the same point on a reactor and column, where the control structure and not the tuning decides how the plant answers a production change.

9.6.2 One Integrator Is Not Fast, It Is Only Consistent

Scenario 4 puts the master in charge and asks the grid for 8 MW more electricity at 90 s, which is about 24 MW of core power that does not exist yet, with the reactors inside the pressure loop. The header falls to 3.2771 MPa, 0.123 MPa below setpoint, and re-enters a ±0.02 MPa band for good

Table 9.6 Measured Header Excursions on Three Upsets. Scenario 6 Steps the Generator Down 23 MW at $t = 60$ s and Raises It 35 MW in Steps of About 9 MW a Minute Apart From $t = 480$ s. Scenario 7 Steps Process Steam 56 to 68 kg/s and Back at 600 s. Scenario 8 Repeats Scenario 6 With the Bypass Resting at 2 % and the Store In and Out of Service. Ranges Are Over the Whole Run After $t = 30$ s, to 900, 780 and 900 s. MPC Rows Use Scenario 11's Settings.

Case	P_h min (MPa)	P_h max (MPa)	bypass peak (%)	relief (s)	customer
6, coordinated	3.166	3.449	24.9	0	kept
6, master-slave	2.647	3.874	3.0	147	lost
6, MPC	2.926	3.455	27.5	0	kept
7, coordinated	3.308	3.416	11.2	0	kept
7, master-slave	3.241	3.567	3.0	0	kept
7, MPC	3.291	3.416	15.0	0	kept
8, store in	3.336	3.448	24.1	0	kept
8, store out	3.124	3.448	24.1	0	kept
8, MPC, store in	3.324	3.455	27.3	0	kept

196 s after the step. Nothing is wrong with the tuning, because the plant cannot move faster than 31.3 MW/min of aggregate ramp rate. Selecting the coordinated architecture and repeating the step gives a nadir of 3.3391 MPa and a recovery in 110 s, an improvement of 0.062 MPa on the deviation and 87 s on the recovery, entirely because a 2 s valve is now in the loop ahead of the reactors.

Scenario 4 also shows why loops designed separately on a shared variable interact. Tuning the generator load controller of Table 9.4 hard, at $K_c = 0.6$ and $\tau_I = 5$ s, and repeating the same 8 MW step drives header pressure between 2.888 and 3.990 MPa over a 900 s run, lifts the relief for 198 s of it and loses the customer. The mechanism is (9.12), because a sagging header reduces the flow the governor valve passes at a fixed position, so a stiff megawatt controller answers a sagging header by opening the valve further, which takes more steam, which sags the header further.

9.6.3 Mid-Ranging and What the Resting Position Costs

Scenario 5 runs the coordinated architecture undisturbed. In steady state the bypass sits on its 3.0 % resting position passing exactly 3.00 kg/s, header pressure is 3.4000 MPa, and the three module margins are 12.8, 12.9 and 13.0 °C. Three percent of bypass is 7.0 MW thermal thrown to the condenser continuously, or 2.35 MW of electricity not sold, which is 99 $/h at the default price, and what it buys is 3 kg/s of instant upward response as well as unlimited downward response. Setting it to zero makes the fast path one-sided, because a shut valve can still catch a rising header in two seconds and can do nothing about a falling one. A setpoint step of +0.05 MPa, which leaves the header low, settles inside 0.01 MPa in 26 s with the bypass resting at 3 % and, at a zero resting position, in 61 s once the bypass has closed to 0.5 % at 349 s and in 64 s once it is at 0.3 % after ten minutes, while a −0.05 MPa step settles in 5.35 s either way. Every mid-ranging scheme makes that trade and the resting position prices it [3].

9.6.4 Load and Process Upsets, With and Without a Fast Actuator

Table 9.6 reports scenarios 6 to 8, each run twice on the same disturbances with the architecture or the store as the only change, and once more under the MPC of Section 9.5.5.

The first pair of rows in Table 9.6 is the clearest single result in the application. Coordinated control holds the header between 3.166 and 3.449 MPa through both halves of the ramp. The step down opens the bypass to 24.9 % within seconds, and the valve position controller then spends four minutes walking the reactors down so the bypass can close again, at a cost of 4.35 MWh to the condenser over the 900 s run, and the ramp up is answered almost entirely by the reactors because nine megawatts a minute of electricity is about 27 MW/min of core power against 31.3 MW/min

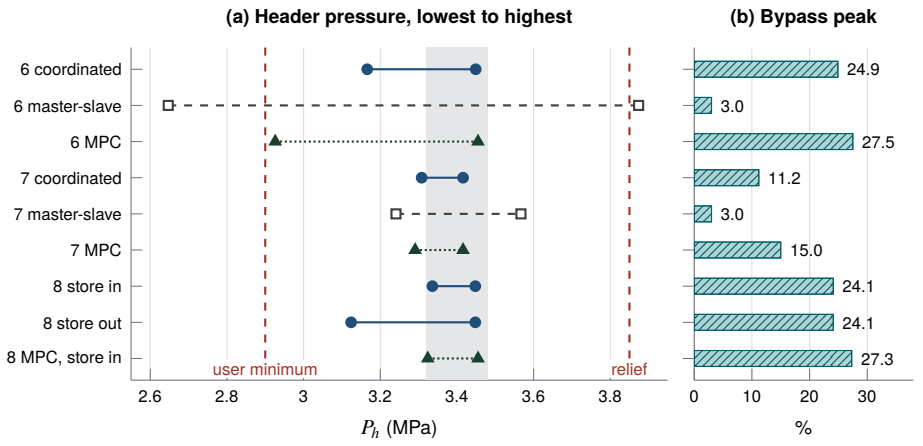


Figure 9.3 Table 9.6 drawn as ranges and bars: scenarios 6 to 8 under coordinated control (solid, circles), master-slave (dashed, squares) and the MPC (dotted, triangles). (a) Header pressure from its lowest to its highest value over each run, against the shaded ±0.08 MPa band, the industrial user's 2.90 MPa minimum and the 3.85 MPa relief pressure (dashed). (b) Bypass peak. Notice that in scenarios 6 and 7 the master-slave run, whose bypass stays at 3.0 %, has the widest range, and that scenario 6 under master-slave crosses both dashed lines: the relief lifts for 147 s and the customer is lost.

available. Master-slave on the identical events fails in both directions. The step down has nothing fast to absorb it, so the header reaches 3.874 MPa and the relief lifts for 147 s, and the ramp up then drives it to 2.647 MPa and the customer is lost. Same reactors, same rate limits, same disturbance, and the only difference is whether a fast actuator is in the loop. Figure 9.3 draws the table as ranges and bars.

Scenario 7 is milder and shows the same ordering, a nadir of 3.308 MPa under coordinated control against 3.241 MPa under master-slave and an overshoot on the return of 3.416 MPa against 3.567 MPa. The process flow controller is part of the mechanism rather than a bystander, because it has to open its valve further than it planned to hold the flow it was asked for while the header is falling, which pulls the header down again.

The store in scenario 8 improves the nadir from 3.124 to 3.336 MPa and reduces the energy thrown to the condenser from 3.97 to 3.37 MWh, from a tank that ends the run at 44.4 % charge against the 45.0 % it started with. Two limits appear at once. The delivery pressure of (9.13) falls with the charge, from 5.40 MPa full to 3.05 MPa empty, so a nearly flat tank cannot push steam into a healthy header, and storage steam is barely superheated, so a hard discharge cools the header through (9.11). Both are real properties of a sensible-heat store, and both are why it is a fourth participant rather than an answer [16].

9.6.5 The Allocation Rule Matters as Much as the Architecture

Scenario 9 fouls Unit 3's steam generator from 0.86 to 0.70 at 120 s and then brings the evening peak. Header pressure dips to 3.303 MPa at 157 s on the fouling step and is back inside its band by 194 s, because the total goes on being delivered, and the only lasting indication is a module temperature warning at 245 s. With the equal split the run ends at 900 s with the three modules at 198, 198 and 166 MW carrying margins of 8.5, 4.9 and 0.0 °C, a spread of 8.43 °C, so Unit 3 is pinned on its 315 °C limit while Unit 1 sits with eight and a half degrees spare (9.2, 4.9 and 0.0 °C once Unit 1 has finished unloading at 1200 s). Switching to the margin-proportional rule redistributes the same total to 208.9, 196.6 and 155.4 MW with margins of 5.41, 5.66 and 4.59 °C, a spread of 1.07 °C. The plant sells the same product from nearly the same steam, 242.2 against 243.7 kg/s, and has simply stopped choosing which module to break. The architecture is identical in both runs, so the 7.36 °C reduction

in margin spread belongs entirely to the splitting rule, which is the result unit-specific allocation across a multi-unit plant is built to obtain [8].

The MPC of Section 9.5.5 still believes Unit 3 is at 0.86, so at its worst instant its model allows Unit 3 203.4 MW where the true ceiling of (9.22) is 165.5 MW, and only the ceiling clamp every architecture shares holds the module. Unit 3 ends 0.10 °C past its limit, against 0.09 °C under Coordinated with the equal split, and the clamp has a cost. The plan keeps asking for a ramp the clamp refuses, the header IAE is 20.04 against 18.04 and 15.16 MPa s for the equal and margin splits, and the header settles 0.004 MPa under setpoint, 3.396 MPa at 2400 s. In a diagnostic copy without the clamp the MPC runs Unit 3 5.12 °C past its limit at 900 s. The comparison is therefore not between a controller that knows the fouling and one that does not, because the clamp and Coordinated's allocator both estimate the ceiling from the measured coolant temperatures, which follow the fouling.

9.6.6 A Module Trip Is a Dispatch Decision

Scenario 10 scrams Unit 1 at 90 s. Core power collapses toward decay heat with the 1.6 s scram time constant, 39 MW three seconds after the trip, and the isolation valve is shut in six through (9.9), so 82.0 kg/s of the header's 228.2 kg/s stops arriving into a pipe holding 7.55 s of steam. The two survivors have ceilings of 218 and 204 MW against a plant that was using 530 MW, so the arithmetic does not close and no architecture makes it close. Three runback choices were measured over 1500 s with the coordinated architecture and the margin split throughout. With no runback the nadir is 1.936 MPa, the header spends 474 s below P_p^{min} and the customer is lost. With the turbine armed the nadir is 2.943 MPa, no time at all is spent below P_p^{min}, the customer is kept and 41.7 MW is shed. With the process user armed the nadir is 2.082 MPa, 473 s are spent below P_p^{min}, the customer is lost anyway and 18.7 kg/s is shed. All three runs end with header pressure back on 3.400 MPa.

The no-runback case deserves attention for how it ends rather than for how deep it goes. The customer's 56 kg/s of demand leaves with the customer, the two surviving reactors then have enough capability for what remains, and header pressure recovers neatly to 3.400 MPa. The plant balanced itself by losing the customer, the header tile is green again by 900 s, and an operator reading the pressure trend at that instant sees a plant that handled a reactor trip. The rest of the scorecard does not agree, with the industrial user, worst margin, reserve and time off band all red.

> **TRY IT IN THE APP** **10 · Unit 1 Scrams**
>
> Load scenario 10 and press Run with the runback on None. Unit 1 trips at 90 s, the header falls to 1.936 MPa and the industrial customer is lost, after which the pressure recovers to 3.400 MPa. Reset, arm the runback on the turbine and run again: the lowest pressure is 2.943 MPa, the customer stays, and 41.7 MW of generation is shed.

Arming the runback on the turbine changes the outcome without changing anything about the reactors. The fast stage of (9.32) takes 79 MW off the generator setpoint within the fifteen seconds after the trip that would otherwise not have existed, the nadir is 2.943 MPa, and the starvation timer never starts. The fast stage then hands load back as the header recovers until the latched trim holds 41.7 MW, and the plant finishes on 3.400 MPa with the two survivors carrying 407 MW. The trade is 41.7 MW of electricity not sold against a customer whose interruption costs 40 000 $ plus 9000 $/h. Arming it on the process user cannot work, for reasons of arithmetic rather than tuning, because the runback may take the customer down only to a little above its 34 kg/s shutdown flow, which is 18.7 kg/s of relief against a shortfall of about 47 kg/s. Figure 9.4 shows the no-runback case 540 s after the trip.

Under the MPC, which after the scram plans the two survivors with Unit 1's ramp held at zero, the armed turbine runback picks up at 92.30 s against 92.05 s. The plan shuts the bypass from 3.02 % in its first decision, at 90.25 s, where Coordinated has it shut at 91.75 s, so the header reaches the pick-up threshold a scan later. The nadir is 2.941 against 2.943 MPa, 42.2 against 41.7 MW is shed,

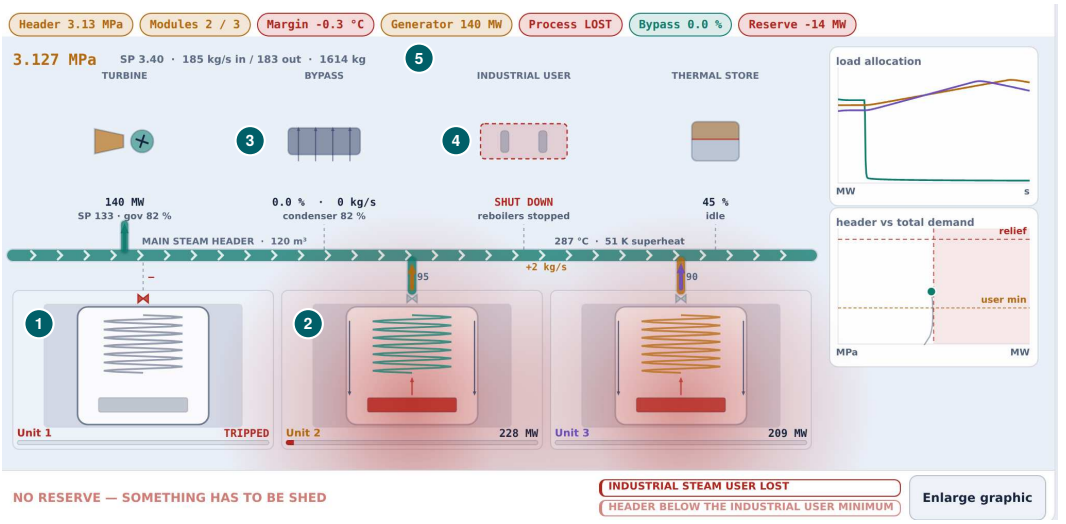


Figure 9.4 Scenario 10 at $t = 10.50$ min, 540 s after the Unit 1 scram, with the runback set to None and the coordinated architecture in service, in the plant view of the app. Header pressure reads 3.127 MPa (5) and is climbing back toward its 3.40 MPa setpoint for the wrong reason, because the industrial user (4) was lost and its 56 kg/s of demand went with it. Unit 1 (1) is tripped and pale, Units 2 and 3 (2) carry 228 and 209 MW, above ceilings of 222 and 207 MW, the worst module margin is –0.3 °C, thermal reserve is –14 MW, and the bypass (3) is hard shut.

and the IAE is 78.02 against 77.42 MPa s. The high-temperature trip, its limit lowered to 302.5 °C below Unit 1's coolant temperature at 10 s of scenario 5, fires at 10.05 s under both.

9.6.7 The Graded Challenge

Scenario 12 loads the plant weakly tuned, on master-slave with an equal split, a bypass hard shut at zero percent and Unit 3 already fouled to 0.78, then scrams Unit 1 at 120 s and asks the customer for 62 kg/s at 450 s, with a target of ±0.08 MPa of 3.40 MPa, no lost customer and no relief lift. Run exactly as loaded, the header reaches 2.091 MPa, the customer is lost 56 s after the scram, and the worst residual in the final third of the run, from $t = 594$ s, is 0.725 MPa, about nine times the band. One change is both necessary and sufficient for this target. A runback armed on the turbine, with everything else as loaded, keeps the customer, lifts no relief and leaves a worst residual of 0.042 MPa, while every run without it loses the customer, including the full suggested set with the runback switched off. The other choices add margin rather than feasibility. Coordinated allocation and the margin-proportional split with the runback, on the weak tuning and a shut bypass, give 0.033 MPa, and the suggested settings add a 6 % bypass resting position, storage on automatic and stronger tuning. They give a nadir of 3.087 MPa, no relief lift, 31.4 MW of generation shed, and a worst residual in the final third of 0.046 MPa against the 0.08 MPa tolerance, inside the band. The split rule cannot matter once the trip has happened, because both survivors then sit on their temperature limits whatever the rule, with a worst margin of 0.0 °C at 900 s in all three runs, so the target rewards the dispatch decision rather than the structure. The MPC confirms it, and the challenge itself is unchanged. With only the architecture switched to MPC the runback stays off, the customer is lost 54.4 s after the scram and the worst residual is 0.736 MPa, while the MPC with the suggested settings meets the target at 0.038 MPa, with Unit 3 ending 0.02 °C past its limit on the shared clamp where the suggested coordinated run ends it on the limit.

Table 9.7 Scenario 11 and the Other Cases, Coordinated Against the MPC at Scenario 11's Settings With the Same Seed, Events and Noise. Header IAE in MPa s From the 4 Hz Historian, to 900 s Except Scenarios 4 and 7, to 600 and 780 s, Unless Stated

Case and measure	Coordinated	MPC
Scenario 11, 0–900 s		
IAE, whole run	72.56	71.61
IAE, load drop, 0–300 s	3.36	1.04
IAE, process step, 300–540 s	1.22	0.35
IAE, trip, 540–900 s	67.98	70.22
At the rate limits (module-s)	1262	383
Module (MW), bypass (%) travel	237, 60.2	126, 58.3
To the condenser (MWh)	3.15	4.72
Worst coolant margin (°C)	1.31	4.08
Plans, rejected, mean solve (ms)	—	900, 0, 3.80
IAE, and to the condenser (MWh)		
4, +8 MW load step	5.58, 0.84	7.95, 1.02
6, evening ramp	47.15, 4.35	111.73, 5.82
7, process step	10.82, 2.09	12.76, 2.97
8, store in service	19.98, 3.37	19.03, 5.65
8, store in service, to 2400 s	22.63	33.18
9, equal split	18.04, 2.32	20.04, 1.84
9, margin split	15.16, 2.21	20.04, 1.84
10, no runback	504.80, 0.87	513.19, 0.79
10, turbine runback	77.42, 1.20	78.02, 1.26
Scenario 12, suggested settings		
Worst residual (MPa), Unit 3 margin	0.046, 0.00 °C	0.038, –0.02 °C

9.6.8 Model Predictive Control Against Coordinated Allocation

Scenario 11 runs one day on the tuned loops with the margin split, the store out of service and the turbine runback armed. The grid takes 23 MW off the generator at 60 s, the process user steps from 56 to 68 kg/s at 300 s, and Unit 1 scrams at 540 s. Table 9.7 sets Coordinated against the MPC on it and on the earlier cases, with Coordinated on each case's own tuning, the MPC at scenario 11's settings throughout and no preview given to either structure, and Figure 9.5 shows the day.

Over the whole run the header IAE is a tie, 71.61 against 72.56 MPa s, 1.3 % apart, and the difference lies in how the steam is found. On the load drop Coordinated opens the bypass to 24.9 % at 125.5 s, then walks the reactors down behind it, 93 MW of module travel by 300 s with the bypass back at 8.0 %. The MPC opens it to 27.6 % at 148.3 s and still holds 20.1 % at 300 s, having moved the modules 40 MW. Over the day the modules spend 383 against 1262 module-seconds ramping at their rate limits and travel 126 against 237 MW, and the price is 4.72 against 3.15 MWh to the condenser, because a ramp is priced in (9.33) and an open bypass only through q_b. Both structures bottom out at 2.97 MPa after the trip. The worst coolant margin of 4.08 against 1.31 °C comes from loading the survivors less, 389.6 against 405.6 MW of core power at 900 s, and not from the plan's temperature limit, which carries no slack on any of the 900 plans. It holds for scenario 11 only, where every steam generator is in its declared condition. The bypass rest weight was lowered once before tuning, from 0.444 to 0.02, and it decides this ranking, since at 0.444 the IAE is 84.24 MPa s with 2.82 MWh to the condenser.

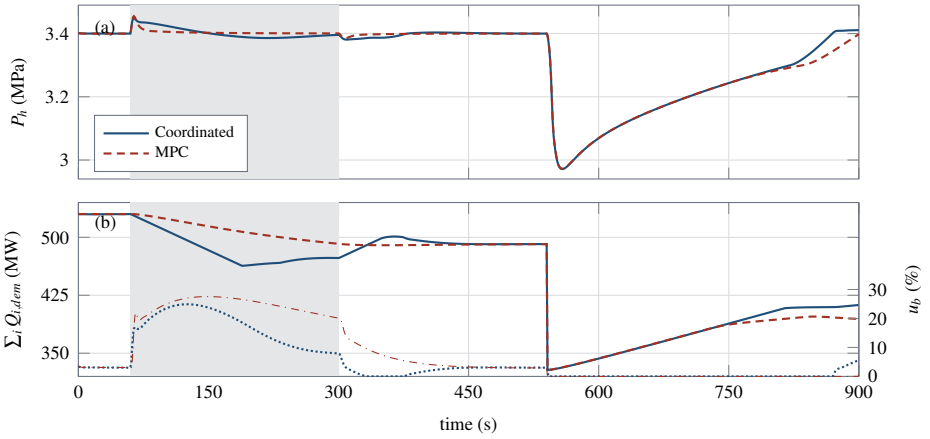


Figure 9.5 Scenario 11 under Coordinated (blue, solid) and the MPC (red, dashed), the decisive load drop shaded. (a) Header pressure. (b) Total module demand (left axis, Coordinated solid and MPC dashed) and bypass command (right axis, Coordinated dotted and MPC dash-dotted), with the bypass held open longer and the reactors moved less by the MPC.

> **+ DID YOU KNOW? An Optimizer That Only Ties**
>
> From the simulator: over the day of scenario 11 the model predictive controller and coordinated allocation hold the header equally well, 71.61 against 72.56 MPa s of integrated error. The plan keeps the modules at their ramp limits for 383 module-seconds instead of 1262, and pays with 4.72 instead of 3.15 MWh sent to the condenser (Table 9.7).

The other cases show where pricing the ramp costs pressure. The MPC spends less time at the rate limits in every case but scenario 10 without the runback, 1059 against 1047 module-seconds, by factors of 1.4 to 9.0, and moves the modules less in all of them. It wins the downward halves, 1.11 against 3.59 MPa s over the first 480 s of scenario 6 and 0.42 against 1.03 on scenario 7's return. It loses the upward events, scenarios 7 and 4, where λ delays the first ramp, and loses scenario 6's evening ramp by more than a factor of two, the header falling to 2.926 MPa against 3.166 MPa. The plan holds w_u at its estimate for two minutes while the ramp keeps raising the draw, so after 480 s its header prediction 30 s ahead is 0.052 MPa high on average. It sends more steam to the condenser in most cases but less in scenario 9 and in scenario 10 without the runback. Scenario 8 depends on the run's length, 19.03 against 19.98 MPa s to 900 s and 33.18 against 22.63 to 2400 s. The store's discharge controller, offered the setpoint 0.04 MPa low, holds the header near 3.368 MPa from 900 to 1300 s while the plan walks the total demand up from 532.5 to 567.9 MW at its priced rate. No plan was rejected in any case.

9.7 Technology in Practice

The symbols of Figure 9.1 are, in a plant, pressure transmitters on the main steam header, power and flow transmitters on the generator and the process line, the reactors' own power control, a set of air-operated valves, and the controllers of a plant control system. Figure 9.6 draws the coordinated architecture with that hardware in place of the symbols.

What Measures It

Header pressure comes from pressure transmitters, which convert the pressure to a standard signal calibrated to a span, most often an analog current loop and, in newer smart transmitters, a digital signal that can also report the transmitter's own health [23]. The chapter carries each of the three as

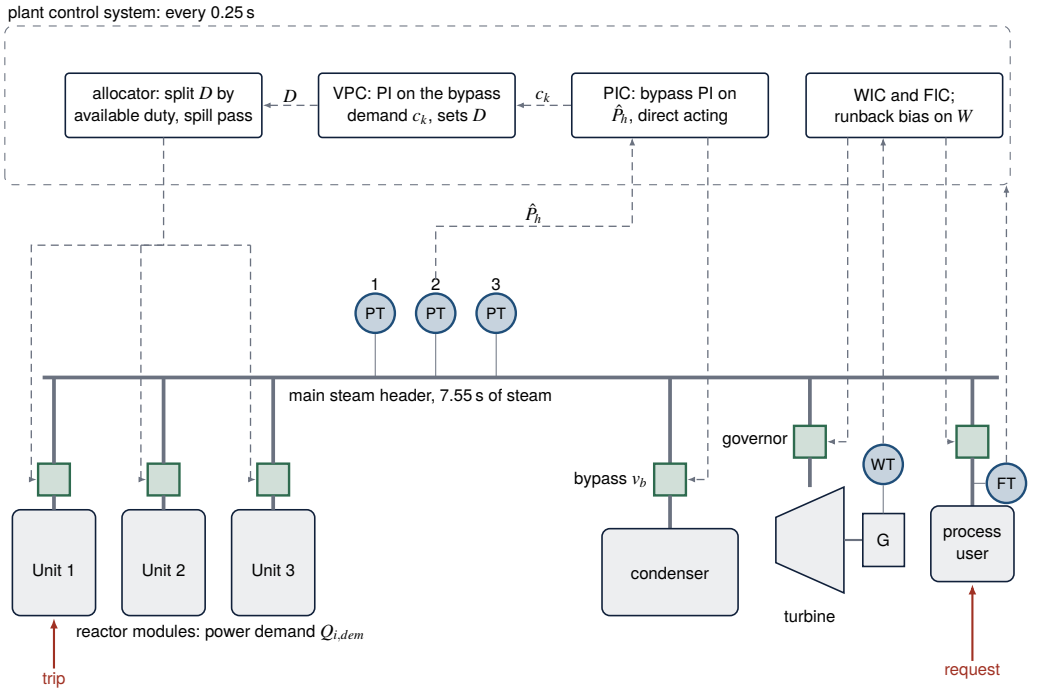


Figure 9.6 The coordinated architecture of Section 9.5 with plant hardware in place of the symbols of Figure 9.1. Sensors are circles, final elements are green squares, and dashed arrows are signals. The bypass PI block reads the header transmitters' unbiased value $\hat{P}_h$, equal to the reading of PT-2, and moves the bypass in seconds. The valve position block reads that block's unclamped demand and sets the total D, which the allocator divides into the three module power demands. The generator output (WT) and process flow (FT) loops move the governor and process valves. The storage valves and the relief are not drawn. Notice that one reading of the header feeds one pressure controller here, where the independent architecture gives each module a controller on its own transmitter.

a 0.4 s filter, Equation (9.19), and gives them the calibration spread of Equation (9.20). Generator output and process steam flow have their own transmitters, carried as 0.8 s and 1.2 s lags. None of the three measures what scenario 3 needs to see, the margin of each module, which comes from the coolant temperatures.

> **TECHNOLOGY** **Sensor: The Header Pressure Transmitter**
>
> A common electronic pressure sensor is the differential capacitance cell. A taut metal diaphragm sits between two fixed metal plates, and pressure flexes it toward one plate and away from the other, changing the two capacitances [24]. The transmitter turns the reading into a standard signal, most often 4 to 20 mA over a calibrated span, and the live zero at 4 mA reveals an open circuit [23]. The error that matters here is a zero shift, which moves every point of the range by the same amount [24]. On a trend it shows as nothing at all, only a steady and believable number. The chapter's three transmitters carry such shifts of $+0.015$, 0 and -0.015 MPa behind a 0.4 s lag.

What Moves It

The slow final element is each reactor's power control, which the chapter reduces to a rate-limited demand and a 14 s lag. In an island-mode dispatch optimization of small modular reactor units, most of the variation in demand was met by moving core thermal output with the control rods, and the turbine bypass was used below minimum power and for short periods when xenon limited the rods [9].

The fast final elements are valves: the governor, the process valve, the bypass and the two storage valves, which the chapter strokes as first-order lags of 0.6 to 5.0 s with no dead band and no stiction.

> **TECHNOLOGY** **Final Element: The Turbine Bypass Valve**
>
> The bypass to the condenser is an air-operated valve. Air pressure on a diaphragm pushes the stem against a spring, and a positioner adjusts the air until the stem reaches the position the controller asks for [25]. Valves that dump main steam to the condenser must stroke very quickly on a turbine load rejection, yet hold stable pressure control during startup and shutdown. Wear shows as air leakage, friction and dead band, which a digital positioner can find before it causes an upset. Safety analyses also consider the opposite failure, one turbine bypass valve opening by mistake, which raises steam flow and cools the reactor coolant [25]. In the chapter the bypass strokes as a 2.0 s lag and passes 1.00 kg/s per percent.

Where the Controller Runs

Plant control systems run supervisory and lower-level loops with cycle times from milliseconds to minutes, often nested so that one loop sets the setpoint of another [26]. The chapter's loops execute every 0.25 s and the predictive controller every 1 s. The coordinated structure is a chain of function blocks: a direct-acting PI block on header pressure that drives the bypass, a valve position PI block that reads that block's unclamped output and sets the total demand, and an allocation block that divides the demand by available duty with its spill pass. The runback adds a bias to the generator load setpoint.

> **IN PRACTICE** **Checking Three Transmitters Against Each Other**
>
> Where several instruments read the same quantity, operators perform a channel check: they compare the readings of the redundant channels on a regular basis and confirm that they agree within pre-defined criteria. On-line monitoring does the same from recorded data and can show that a sensor has not drifted since its last calibration. That matters because mis-calibration during maintenance is one of the most common modes of failure [23]. The three header transmitters of Equation (9.20) would pass such a check, since a spread of 0.015 MPa is a good calibration. Agreement within tolerance is not a license for three integrators on three readings.

What Is Modernizing It

Three developments reach this plant. Smart transmitters allow a calibration check from a remote location [23], and on-line monitoring of instrument channel calibration status uses data the control system already records [23], which is how a growing transmitter spread could be caught before a second integrator finds it. Digital positioners diagnose air leakage, friction and dead band on the valves, so that degradation is found before an upset [25]. And dispatch across units is becoming an optimization over days and years, with outage schedules and markets for grid services in the same problem [9], the long-horizon relative of the predictive controller of Section 9.5.5.

9.8 Limitations

The application is not a reactor model and cannot support a claim about a real plant's transient capability, safety margins or licensability. With no neutronics, nothing about reactivity, xenon or a rod-worth limit can be studied, and with no primary flow state the flow-limited power that constrains a natural-circulation module is absent. The steam generator is one lumped conductance where a real helical-coil once-through unit has moving section boundaries [12], and feedwater flow is not independent, so the coupled feedwater problem across several once-through generators is out of reach [11]. The header is one lumped volume with no pressure drop, which understates the disagreement between transmitters at opposite ends of a real header and therefore understates the failure of Figure 9.2. Repeatability is itself unrealistic, because a fixed 0.05 s step and a fixed 0.25 s scan reproduce one trajectory on every machine while a distributed control system does no such thing.

The four features that survived the simplification, a small shared inventory between rate-limited producers, a ceiling that moves with the shared variable, a fast expensive actuator mid-ranged against a slow free one, and a consumer that can be lost, all survived because a control system has to be designed around each of them. Where each assumption breaks is set out in Section 9.3.6.

The MPC of Section 9.5.5 holds the users' draw constant over its two-minute horizon and leaves out their loops, which is why its model overstates the header's open-loop response by about a third and why it loses the evening ramp by 111.73 against 47.15 MPa s, and it plans every module at its declared fouling, so a fouled module is planned against a ceiling it no longer has. The shared clamp and Coordinated's allocator, by contrast, estimate each ceiling from the measured coolant temperatures and the declared conductance, so for a fouled module their estimate can differ from the physical ceiling during a transient and approaches it as the measured temperature approaches its limit.

9.9 Discussion

The result that generalizes furthest is not the divergence of Figure 9.2 but the reason an operator cannot see it. Header pressure is the variable that appears on the control room display, carries the alarm limits and enters the scorecard, and through the whole of scenario 3 it holds 0.0162 MPa of a 0.16 MPa band, while the evidence that the plant is coming apart is on per-module power and margin trends and a margin-spread score, none of which is a controlled variable. The same asymmetry appears in the trip results, where the run that loses the customer ends with header pressure on setpoint just as the run that keeps it does. Any plantwide design that selects its displays and alarms from the list of controlled variables reproduces both blind spots, which is an argument for choosing the monitored set from the constraint list rather than from the loop list [1, 2].

The economics of the fast path are worth stating with numbers because they are frequently argued without them. A bypass resting at 3 % costs 7.0 MW thermal continuously, which is 99 $/h of electricity not sold at the default price, while losing the industrial customer costs 40 000 $ at once plus 9000 $/h until it returns. The one-off loss alone equals about 17 days of continuous bypass loss, and each hour the customer stays off adds almost four days more. On this plant, however, neither a 3 % resting position in scenario 10 nor a 6 % one in the challenge keeps the customer through the trip without the runback, so the money buys transient margin rather than the customer. The application exposes both the resting position and the four prices so that the trade can be computed rather than asserted.

The model also supports methods beyond the implemented loops. The ceiling of (9.22), the ramp limits of (9.3) and the customer's two limits are inequality constraints that a receding-horizon formulation enforces explicitly rather than through a spill pass and a runback latch, and the storage state of (9.17) is what makes the horizon matter. The controller of Section 9.5.5 enforces the ramp limits as hard bounds and the temperature limit as a priced one, and it leaves the store, the customer's two limits and the runback outside its plan. Coupling one heat source to several products is the arrangement that nuclear-renewable integrated energy system studies address with dynamic optimization over a day rather than with regulatory loops [15, 16], and allocation is then answered by the optimizer's constraint set rather than by a participation shape [7, 8].

> **OPEN CHALLENGE** **Dispatching Many Reactors Over Many Years**
>
> A multi-unit site can share load between its units, but it also shares services such as refueling equipment, and its units may follow a fixed sequence of fuel reloads. In one study, joining the markets for grid services lengthened the fuel cycles, and over a decade the refueling outages drifted by three months, into the summer. The cost of component wear was left for future work [9]. Progress would be an allocator that divides load across units under their fuel, outage and wear constraints as well as their thermal margin, which is all this chapter's allocator weighs.

9.10 Conclusion

This chapter studied three small modular reactors on one steam header as an interactive laboratory for load allocation on a shared variable. The simulator integrates eighteen states, three core powers on rate-limited demands, three average primary coolant temperatures, three once-through steam generator pressures, the header mass and temperature, a storage charge, five valve positions and a generator output lag. The header holds 1724 kg of steam and passes 228.2 kg/s for 7.55 s of residence against a primary loop time constant of 52.5 s and an aggregate ramp capability of 31.3 MW/min. Independent per-module pressure control separates the three modules by 91.7 MW in twenty minutes with no disturbance applied while header pressure stays inside 0.0162 MPa, master-slave holds a 23.2 MW spread and a 196 s recovery from an 8 MW load step, and coordinated allocation recovers from the same step in 110 s because a 2 s bypass valve is in the loop ahead of the reactors. An equal megawatt split leaves margins of 8.5, 4.9 and 0.0 °C where the margin-proportional rule leaves 5.41, 5.66 and 4.59 °C on the same total demand. A Unit 1 scram with no runback takes the header to 1.936 MPa and loses the industrial customer, after which the plant recovers to 3.400 MPa precisely because the customer is gone, while a turbine runback holds the nadir at 2.943 MPa and sheds 41.7 MW, and in the graded challenge that runback is the one change both necessary and sufficient. A model predictive controller planning the three ramps and the bypass ties coordinated allocation on header error, 71.61 against 72.56 MPa s over a day with a load drop, a process step and a trip, while spending 383 rather than 1262 module-seconds at the rate limits and 4.72 rather than 3.15 MWh on the condenser. A fuller model would add a sectioned steam generator with independent feedwater control, a primary flow state to bound module power, a tank thermocline, and a receding-horizon allocator that infers each module’s fouling from its coolant temperature rather than assuming the declared value. The plant shows, in a form a learner can watch, that a controlled variable on setpoint is not the same thing as a plant under control.

KEY POINTS

- A pressure on its setpoint is not proof of a plant under control: with no disturbance, three controllers on one pressure walk the reactors 91.7 MW apart while the header stays inside its band.
- One controlled variable gets one integrator, and the split is a separate decision made by a separate mechanism.
- A 2 s bypass valve ahead of the reactors recovers an 8 MW load step in 110 s, where master-slave needs 196 s.
- Splitting by margin leaves 5.41, 5.66 and 4.59 °C of margin where an equal split leaves 8.5, 4.9 and 0.0 °C.

Where else this applies. Plantwide structure decides the answer in Volume 6, Chapter 7 and Volume 6, Chapter 8, and Volume 4, Chapter 10 mid-ranges a fast, costly bypass against slower actuators in the same way.

References

[1] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, “Plantwide control design procedure,” *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

[2] S. Skogestad, “Control structure design for complete chemical plants,” *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[3] F. G. Shinskey, *Process Control Systems*, 4th ed. New York, NY: McGraw-Hill, 1996.

[4] F. Paparella, L. Dominguez, A. Cortinovis, M. Mercangoz, D. Pareschi, and S. Bittanti, “Load sharing optimization of parallel compressors,” in *2013 European Control Conference (ECC)*, 2013, pp. 4059–4064.

[5] A. Torzhkov, P. Sharma, C. Li, R. Toso, and A. Chakraborty, “Chiller plant optimization – an integrated optimization approach for chiller sequencing and control,” in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 2741–2746.

[6] J. M. Guerrero, M. Chandorkar, T.-L. Lee, and P. C. Loh, "Advanced control architectures for intelligent microgrids—part I: Decentralized and hierarchical control," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 4, pp. 1254–1262, 2013.

[7] A. Nuerlan, P. Wang, J. Wan, and F. Zhao, "Decoupling header steam pressure control strategy in Multi-Reactor and Multi-Load Nuclear Power Plant," *Progress in Nuclear Energy*, vol. 118, p. 103073, 2020.

[8] H.-H. Byun and M.-S. Yim, "Genetic algorithm-based non-synchronous unit-specific optimal load-following control of multi-unit small modular reactor," *Energy*, vol. 314, p. 134091, 2025.

[9] S. Alhadhrami, G. J. Soto, and B. Lindley, "Dispatch analysis of flexible power operation with multi-unit small modular reactors," *Energy*, vol. 280, p. 128107, 2023.

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

[11] P. Sun and X. Zhang, "Modeling and simulation of feedwater control system with multiple once-through steam generators in a sodium-cooled fast reactor (I): Dynamic model development," *Annals of Nuclear Energy*, vol. 138, p. 107208, 2020.

[12] F. U. Zaman, K. Qureshi, I. Haq, and W. Siddique, "Thermal hydraulics analysis of a helical coil steam generator of a small modular reactor," *Annals of Nuclear Energy*, vol. 109, pp. 705–711, 2017.

[13] J. Carter, "Small modular reactor deployment: Learning from the past and the present," *The Electricity Journal*, vol. 29, no. 8, pp. 1–8, 2016.

[14] D. T. Ingersoll, Z. J. Houghton, R. Bromm, and C. Desportes, "NuScale small modular reactor for Co-generation of electricity and water," *Desalination*, vol. 340, pp. 84–93, 2014.

[15] R. A. Jacob and J. Zhang, "Modeling and control of nuclear–renewable integrated energy systems: Dynamic system model for green electricity and hydrogen production," *Journal of Renewable and Sustainable Energy*, vol. 15, no. 4, p. 046302, 2023.

[16] A. Rigby, M. J. Wagner, and B. Lindley, "Dynamic modelling of flexible dispatch in a novel nuclear-solar integrated energy system with thermal energy storage," *Annals of Nuclear Energy*, vol. 204, p. 110534, 2024.

[17] 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.

[18] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

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

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

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

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

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

[24] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[25] 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.

[26] 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.

9.11 Exercises

9.1 Concept check · 2 min

Three reactors feed one main steam header, which supplies a turbine, an industrial steam user, a condenser bypass and a thermal store. Which quantity is the single controlled variable that every module and every user shares?

(a) The fouling multiplier $m_{f,i}$ of each module's steam generator
(b) The calibration spread σ_b between the three pressure transmitters
(c) The main steam header pressure P_h, in MPa
(d) The mass of steam m_h held in the header, in kg

9.2 Concept check · 1 min

In scenario 3 each module runs its own well-tuned PI controller on header pressure, the transmitter spread is at its default of 0.015 MPa, and no disturbance is applied. The module powers walk 91.7 MW apart while the header stays inside its band. What keeps them walking?

(a) A slow load change on the grid that the three loops share unequally
(b) The coupling through saturation temperature, which moves every module's heat removal
(c) A gain three times too high on each loop, which makes the header oscillate
(d) Three integrators on readings that differ by a bias, so no pressure zeroes all three errors

9.3 Concept check · 2 min

The coordinated architecture has three layers: a bypass pressure loop, a valve position controller and an allocator. Which statement describes what the three layers achieve together?

(a) They hold the header on setpoint while dividing the reactor load within each module's thermal limit
(b) They send as much steam as the condenser accepts through the bypass to hold the header
(c) They hold each steam generator pressure at a separate setpoint of its own
(d) They keep the thermal store full whatever the header pressure does

9.4 Theory and calculation · 4 min

At the design point the header holds 1724 kg of steam and passes 228.2 kg/s. A module's steam side holds about 470 kg against an average 76.1 kg/s. Use Tables 9.2 and 9.1.

(a) Compute the residence time of the header and of one module's steam side. *Checkpoint:* 7.55 s and 6.18 s.

(b) Compute the primary loop time constant C_{p1}/UA_1 of Equation (9.4) for Unit 1. *Checkpoint:* 52.5 s.

(c) Compute the aggregate ramp capability of the three modules in MW/min, and the shortest time in which they can add 24 MW of core power together. *Checkpoint:* 31.25 MW/min and 46 s.

9.5 Theory and calculation · 8 min

At the design demand $D = 530$ MW the three ceilings are 232, 218 and 204 MW, and the conductances are $UA_i = 3.43\, m_{f,i}$ with fouling multipliers 1.00, 0.93 and 0.86. Treat the ceilings as fixed and use the margin of Equation (9.24).

(a) Divide D in the ratio of the ceilings and compute the three margins. *Checkpoint:* 12.8, 13.0 and 13.1 °C.

(b) Divide D equally instead and compute the three margins and their spread. *Checkpoint:* 16.1, 13.0 and 9.3 °C, a spread of 6.9 °C.

(c) Under the equal split, at what total demand does Unit 3 reach its limit, and at what total demand does the proportional split use up every margin at once? *Checkpoint:* 612 MW and 654 MW.

9.6 Theory and calculation · 6 min

The bypass rests at 3 % and passes 3.00 kg/s of header steam at 3.40 MPa and 296.8 °C to the condenser. Feedwater is at 149 °C, the turbine makes 783 kJ/kg, electricity sells for 42 $/MWh, and losing the industrial user costs 40 000 $ plus 9000 $/h.

(a) From Equation (9.1) with $c_p = 2.65\,\mathrm{kJ/kg/K}$, compute T_{sat}, the header enthalpy and $h_f(T_{fw})$, and the thermal power the bypass throws away. *Checkpoint:* 241.0 °C, 2951.9 kJ/kg, 631.9 kJ/kg, 6.96 MW.

(b) Compute the electricity that steam would have made and its value per hour. *Checkpoint:* 2.35 MW, 98.7 $/h.

(c) How many days of bypass loss equal the one-off penalty, and how many days does each hour of interruption add? *Checkpoint:* 16.9 days and 3.8 days.

9.7 App challenge · 8 min

Load scenario 3, Independent pressure control, fighting itself, and change no setting. Run to 1200 s at a speed of 60× at the default seed and export the CSV file. The file's time column is in minutes.

(a) From columns `Q1_MW`, `Q2_MW` and `Q3_MW`, read the three module powers and their spread at 10 and 20 minutes. *Checkpoint:* 131.5, 202.0 and 198.1 MW, spread 70.5 MW, then 120.1, 211.8 and 198.4 MW, spread 91.7 MW, each within 0.2 MW.

(b) From column `Ph_MPa`, find the lowest and highest header pressure over the run, and say what an operator watching only that trend would conclude. *Checkpoint:* 3.385 and 3.401 MPa, each within 0.001 MPa.

(c) Press Reset, switch to Advanced, set Pressure transmitter calibration spread to 0 on the Modules tab, and run to 1200 s again. Find the lowest and highest header pressure. *Checkpoint:* 2.693 and 3.917 MPa, each within 0.005 MPa.

9.8 App challenge · 8 min

Load scenario 4, Master–slave load sharing, and change no setting. The grid asks for 8 MW more at 90 s. Run to 600 s at a speed of 20× at the default seed and export the CSV file, whose time column is in minutes.

(a) From column `Ph_MPa`, find the lowest header pressure after the step, and the time after the step of the last row more than 0.02 MPa from the 3.40 MPa setpoint. *Checkpoint:* 3.277 MPa, within 0.002 MPa, and 196 s, within 2 s.

(b) Press Reset, set Pressure control architecture to Coordinated, and repeat the run and the two readings. *Checkpoint:* 3.339 MPa, within 0.002 MPa, and 110 s, within 2 s.

(c) Explain why the master-slave recovery cannot be made much faster by tuning, and what makes the coordinated run faster.

9.9 App challenge · 9 min

Load scenario 10, Unit 1 scrams, and change no setting: Coordinated, margin split, Runback target None. Unit 1 trips at 90 s. Use a speed of 60× and the default seed for every run, and export the CSV file of each.

(a) Run to 900 s. Find the lowest header pressure in column `Ph_MPa`, and read the Industrial user and Header tiles at the end. *Checkpoint:* 1.936 MPa, within 0.005 MPa, LOST, and 3.400 MPa.

(b) Press Reset, switch to Advanced, set Runback target to Turbine on the Dispatch tab before Run, and run to 900 s. Find the lowest header pressure and the generator load shed in the last row of column `shed_bias_MW`, and read the Industrial user tile. *Checkpoint:* 2.943 MPa, within 0.005 MPa, −41.7 MW, within 0.2 MW, and kept, 0 s low.

(c) Load scenario 12, Challenge: Ride through the trip. Change only Runback target to Turbine, run the full 900 s and read the verdict. *Checkpoint:* worst header error 0.042 MPa over the last third, within 0.002 MPa, user kept, relief 0 s.

9.10 Application area · 4 min

In scenario 10 the turbine runback keeps the industrial user by shedding 41.7 MW of generation, at the default electricity price of 42 $/MWh. Losing the user costs 40 000 $ plus 9000 $/h until it returns.

(a) What does the shed generation cost per hour, and after how many hours of shedding would it equal the one-off penalty alone? *Checkpoint:* 1751 $/h and 22.8 h.

(b) The process runback can cut the user only to a little above its 34 kg/s shutdown flow, 18.7 kg/s of relief against a shortfall of about 47 kg/s. What does that mean for the choice of product to shed?

9.11 Application area · 3 min

An operator display for this plant shows header pressure with alarms at ±0.08 MPa and nothing else. Using the Discussion and the Limitations, name the readings that would have revealed the divergence of scenario 3 and the lost user of scenario 10, and say whether a real header with a pressure drop along it makes the scenario 3 failure better or worse.

9.12 Review the agent · 5 min

An AI assistant is asked how to keep a separate pressure controller on each module and replies as follows. “Scenario 3 fails because each module loop has $K_c = 950$ against a master of 320, so three loops act at three times the master’s gain. Set Per-module pressure gain to 110 MW/MPa, a third of the master, and the three loops together give the master’s loop gain of 2.2. I ran scenario 3 for ten minutes with that gain: the spread is 7.2 MW and the header stays between 3.399 and 3.401 MPa. Independent control now works as well as master-slave.” Find the flaw and say how you would check it.