INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 2

Pressurized Water Reactor Rod Control Simulator for Inherent Feedback and Velocity Actuators

Pages 24–52

Edited by John D. Hedengren

How to cite this chapter. "Pressurized Water Reactor Rod Control Simulator for Inherent Feedback and Velocity Actuators," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 24–52.

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

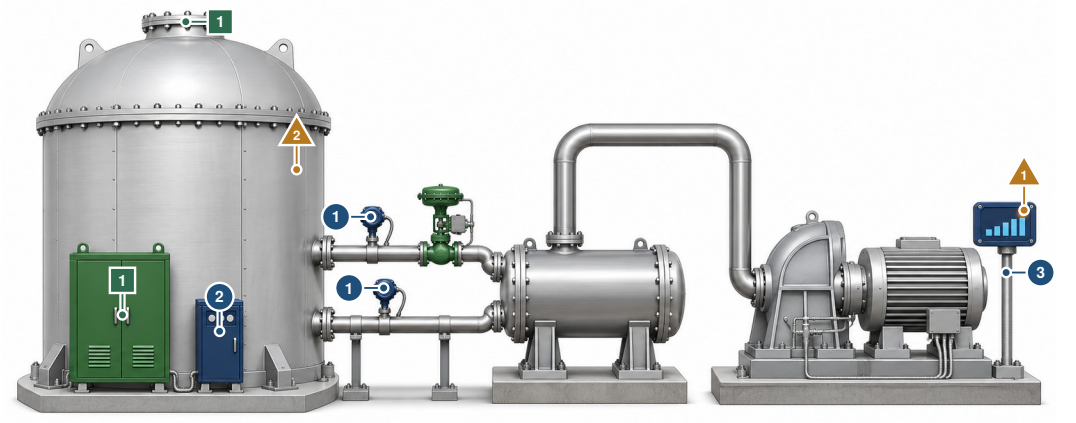
The app for this chapter
https://apmonitor.com/apps/pwr-reactor/

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

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

CHAPTER 2

Pressurized Water Reactor Rod Control Simulator for Inherent Feedback and Velocity Actuators



The controller moves 1 rod control cabinet and the rod drives on the vessel head

The controller reads 1 coolant temperature on the hot and cold legs, averaged
2 reactor power from detectors outside the vessel 3 turbine load

What disturbs it 1 turbine load change 2 a dropped rod

Notice that the rods, not a valve, trim the averaged leg temperature after a load change, so the valve drawn on the hot leg stands for nothing in the chapter's model. Illustration, not to scale.

IN PLAIN TERMS **A Reactor That Follows the Turbine**

The main idea. When the turbine takes less steam, the water in a nuclear reactor warms up, and warmer water slows the chain reaction on its own. The reactor's power then falls to match the turbine with no control rod moving.

The challenge. The reactor does not bring its water back to the right temperature, so that job is left to the control rods. The rod motor is told a speed, not a place to stop, and the slow temperature reading limits how hard the controller can push.

Why it matters. Getting this loop right lets a reactor follow a grid that has more and more wind and solar power.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the two history boxes of Section 2.2, the fact boxes of Sections 2.4.3 and 2.5.3, the try-it boxes of Sections 2.5.1 and 2.5.3, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 2.5, Section 2.6, Section 2.7 and the open challenges

The control question A pressurized water reactor follows a change in steam demand with no rod moving, because its negative temperature coefficients bring fission to the new load. What is left for the rod control system to do, and why does a drive that takes a rod speed make that task harder than the reactor does?

Learning objectives

- Derive the rod gain from bank position to coolant temperature and explain how it moves with the bank position.
- Explain why a proportional-only controller on a speed-commanded rod drive leaves no offset, and price reset in phase margin.
- Identify the measurement lag as the element that sets the achievable gain of the rod loop.
- Predict the limit cycle that a deadband narrower than the smallest rod move produces.
- Choose between feedforward and cascade from where a disturbance enters, and plan boron against the rod insertion limit and xenon.

Prerequisites Sections 1.3 (loop action), 1.4.2 (point kinetics and the rod gain), 1.4.3 (xenon), 1.4.5 (limits) and 1.6.2 (velocity-form PID). Section 1.6.4 describes how a scenario is run and credited.

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

See also Volume 1, Chapter 8 for an inventory whose own feedback arrests an offset that only integral action removes, and Volume 1, Chapter 2 for a gain that moves with the operating point. Volume 4, Chapter 2 for a protective override, and Volume 4, Chapter 9 for a slow resource that takes the load back from a fast one.

2.1 Introduction

The ordinary arrangement in a process plant is that nothing happens until the controller makes it happen. A furnace holds its outlet temperature because a controller moves fuel, a column holds a composition because a controller moves reflux, and a tank holds a level because a controller moves a valve. With the controller removed the plant leaves its operating point. A pressurized water reactor is built the other way around. Reactivity falls as the fuel and the moderator get hotter, and both coefficients are required to be negative at power, so the chain of events that follows a reduction in steam demand is complete without any command at all. Less steam leaves the secondary side, the steam generator warms, the primary loop warms behind it, the negative coefficients remove reactivity, and fission settles at the heat the turbine is now taking. In the simulator of this chapter a reduction of turbine load from 100 to 90 % with every rod stationary brings reactor power to 89.9 % within eight minutes.

What the plant does not do is put the temperature back. The reactivity balance closes at a hotter core, 13.9 °C above the programmed average coolant temperature in that same run, and closing that gap is the entire task of the rod control system. The rod control system is commonly taken to be what moves the power. It moves the temperature instead, and it does so through an actuator more difficult than the plant it acts on. Reactor self-regulation is a documented design property rather than an incidental one, and recent work on small modular plants treats the size of the inherent load-following capability as a quantity to be designed for [1, 2].

Two general classes of problem are visible in this one plant. The first is the plant with strong inherent negative feedback, which includes an exothermic reactor with a negative temperature coefficient of activity, a fuel cell whose voltage falls as current is drawn, and a compressor whose head falls with flow. In each case the controller's task is smaller than the block diagram suggests, and the useful question is what is left over. The governed grid of Volume 1, Chapter 8 has the same shape, where droop arrests a frequency offset and only integral action returns it to nominal. The second is the loop whose final element is a velocity actuator. A gripper-coil rod drive accepts a speed and not a position, and so does a motor-operated valve, a damper drive, an on-load tap changer and a screw feeder. Each of them supplies an integrator to the loop, saturates at a fixed speed, has a deadband, and cannot make a correction smaller than one step. The limit cycle that a deadband and a minimum step produce together is the mechanism that also appears in servo mechanisms with stiction [3], and

a loop that hunts because a proportional-plus-integral controller was placed in front of an actuator which already integrates is a recurring error of practice.

The physical fundamentals behind the model are the ones an introductory course already carries. Energy is conserved across three lumped capacities in series, the driving force for every heat transfer term is a temperature difference, and the reactivity balance is an algebraic constraint that the temperatures themselves close, which is the same conservation-plus-driving-force reasoning applied in the standard texts [4]. Three time scales are separated by orders of magnitude, from seconds in the thermal nodes through minutes in the delayed neutron precursors to hours in the xenon, and that separation is what makes a low-order model defensible at all [5]. The scenarios built on it show integrating dynamics contributed by an actuator rather than by an inventory, offset-free regulation under proportional-only control, the way a measurement lag rather than the process sets the achievable gain, limit cycles produced by a deadband against a minimum step, feedforward from a measured load, cascade structure on an inner flux measurement, and gain scheduling against a process gain that moves with bank position. Load-following capability now carries a direct economic value in systems with a large share of variable renewable generation, so the tuning question and the maneuver planning question are both industrial ones.

A transfer function from rod position to power conceals the fact that the power arrives with no rod motion at all, so neither class is visible to a student given only that model. The demonstration that separates the two needs nothing but the ability to park the rods and move the turbine load. This chapter separates what a pressurized water reactor does for itself from what its rod drive makes difficult, on a simulator whose manipulated variable is rod speed. Section 2.2 traces how reactor control became temperature control. Section 2.3 sets out a plant of thirteen states with its instrument chain, in enough detail to rebuild it. Section 2.4 derives the process gain, describes three controller architectures and a drive carrying three separate nonlinearities, and gives a frequency-domain account of why the achievable proportional gain of this loop is fixed by the resistance temperature detector rather than by the reactor. It also measures that reset costs exactly the phase margin a proportional-plus-integral controller predicts it will. Section 2.5 runs thirteen scenarios, each credited by its own completion rule. The first of them shows the difference between regulation and self-regulation by switching the controller off and watching the process arrive at the new load anyway. Section 2.6 sets the loop in its plant hardware, Section 2.7 states the limitations, Section 2.8 takes two consequences of the loop past the control room, and Section 2.9 summarizes the measured results.

2.2 Background

The first reactor control problem was not power regulation at all. It was approach to criticality, where the quantity of interest is the reactor period and the regulating rod is driven from a period or a flux signal by a servo whose purpose is to keep the reactor from running away. The early literature is organized around exactly that question, around the kinetics of the neutron population, the role of delayed neutrons in making the period controllable at all, and the design of servo loops on flux and period [6, 7]. A reactor treated this way is an unstable plant that needs a controller, which is the opposite of the plant studied in this chapter, and the difference is entirely the presence of temperature.

Power reactors changed the problem by making the core hot enough for its reactivity coefficients to matter. Once the Doppler broadening of the resonance absorption in the fuel and the density change in the moderator both remove reactivity as the core heats, the neutron population is no longer free, and the plant acquires the negative feedback that makes it self-regulating at power [8, 9]. Two consequences followed for the control system. Power regulation ceased to be the point, because the power arrives on its own, and something else had to be chosen as the controlled variable. The choice that pressurized water plants settled on is a programmed average coolant temperature, scheduled against turbine load so that steam pressure stays in a usable range across the operating map, and the rod control system is a temperature controller from that point onward [2, 6]. The architecture that carries the name reactor-follows-turbine is a direct expression of this. The turbine takes what it

needs, the reactor answers through its own coefficients, and the rods trim the temperature back onto its program.

> ⧖ HISTORY **A Fleet Built to Follow the Load (1970s to 1980s)**
>
> After the first oil crisis, France planned a fleet of 34 pressurized water reactors of 900 MWe and 20 of 1300 MWe under agreements of 1974 and 1976, and by 1985 nuclear units made 60 % of its electricity. Its planners foresaw in the early 1970s that the fleet would have to help balance supply and demand [10]. In the first units the control rods regulated the average coolant temperature, while boron carried the larger changes, a method too slow for daily cycling. A new mode at the end of the 1970s added grey banks, rods several times weaker than the black ones. Grey rod assemblies were first tested in a plant in 1981, and with an average coolant temperature allowed to vary they made the plants far more maneuverable [11].

Xenon-135 is what stopped the arrangement from being simple. It is a fission product with an enormous absorption cross section, produced mostly by the decay of iodine-135 with a half life of 6.7 h and destroyed by neutron capture, so a reduction in power continues to produce it while removing what destroys it. The concentration therefore rises for hours after a power reduction and falls for hours after a power increase, and it can be worth more reactivity than the entire control bank [8]. In a large core the same mechanism supports spatial oscillations in which the flux tilts from one end of the core to the other with a period near a day, which is the reason axial offset is monitored and controlled at all and the reason reduced-order models of spatial xenon dynamics remain an active subject [12]. Xenon is also why load following was for a long time treated as an operational specialty rather than a routine capability. A plant that has reduced power has a bounded window in which it can come back up, and the reactivity to do so has to come from chemistry planned in advance rather than from rods commanded on the minute.

Control research on the load-following reactor followed the tools available at the time. Optimal control formulations appeared once state-space methods were routine, posing the maneuver as a trajectory to be tracked with rod and boron together [13]. Model predictive formulations followed, with the xenon concentration written into the constraint set so that the maneuver is planned against the poison rather than against the temperature alone [14], and more recent work applies sliding-mode and observer-based designs to the same plant [15]. Alongside that, the conventional arrangement has been re-examined for the smaller plants now being designed, written out as the cascade and feedforward structure it actually is and tuned against measured plant response [16]. That structure persists because the manipulated variable has not changed. Whatever the outer algorithm computes still arrives at a gripper coil that steps at a fixed speed and refuses to move inside a band.

Grid decarbonization has turned this from a settled subject back into an open one. Fleets built to run at constant full power are being asked to maneuver because variable renewable generation displaces them at some hours and not others, and least-cost studies of deeply decarbonized systems find the value of nuclear capacity to depend strongly on whether it can follow [17]. The technical and economic constraints on doing so have been collected by the OECD Nuclear Energy Agency, and they are the ones the simulator contains, which are xenon, boron chemistry, rod worth and the insertion limits that protect shutdown margin, rather than anything about the turbine or the controller [11].

> ⧖ HISTORY **Built to Follow, Asked to Follow (1973 to 2013)**
>
> The oil price shock of 1973 sped up the German nuclear program, which aimed at 18 GWe of nuclear power by 1980 and 40 GWe by 1985. That planned share made load following a built-in feature. From the 1970s, German designs could change output over a large range at up to 5 % of rated electrical output per minute, with a part-load program that held the coolant temperature constant in the upper load range. The plants showed that capability at commissioning and then ran mainly at base load, because nuclear reached a 31 % share against the 60 % planned [10]. By 2013 more than 37 GWe of wind and 35 GWe of solar power had to be taken with priority, and flexible operation became continuous at nearly all German nuclear plants [10].

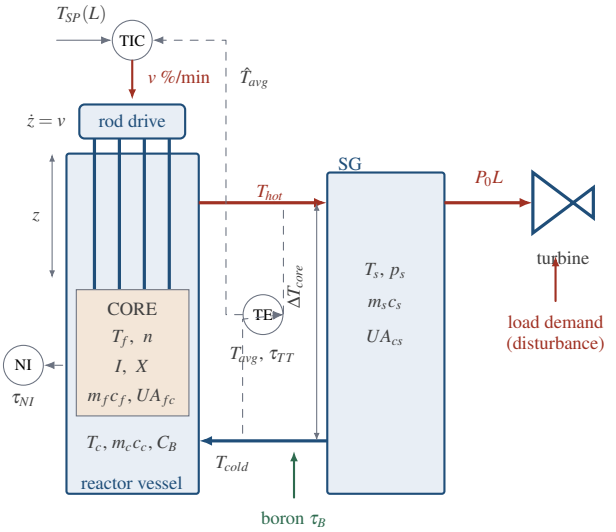


Figure 2.1 Primary circuit of the modeled four-loop plant with control bank D, one lumped steam generator and the turbine. Fission heats the fuel node T_f, which passes heat to the coolant node T_c through conductance UA_{fc}, and the coolant passes it to the secondary node T_s through UA_{cs}. Hot and cold leg temperatures straddle T_c by half the core rise ΔT_{core}, which scales with power. The rod drive integrates the commanded speed v into bank position z, the excore detector NI reports nuclear power n through filter τ_{NI}, and element TE averages the two legs and reports $\hat{T}_{avg}$ through the thermowell lag τ_{TT}. Controller TIC compares that measurement against a setpoint programmed from turbine load L. Boron C_B enters the cold leg through mixing lag τ_B, and iodine I and xenon X live in the core.

2.3 Reactor Model

The plant is a four-loop pressurized water reactor rated at 3411 MW thermal. Heat made in the fuel is carried by pressurized water at an average temperature of 310 °C to steam generators that boil a secondary side at about 270 °C and 5.5 MPa, and the turbine takes whatever steam it wants. Control bank D enters through the vessel head on a gripper-coil drive that accepts a speed rather than a position. Figure 2.1 draws the circuit with its instrument tags and every symbol the equations below use. Working units are the ones a control room reads, with power in MW and as a fraction of rated, temperatures in °C, reactivity in pcm where $1\,\text{pcm} = 10^{-5}\,\Delta k/k$, bank position in percent withdrawn, rod speed in percent per minute, boron in ppm, and time in minutes.

Table 2.1 classifies what the loop interacts with. Thirteen states are integrated. Three annunciations bound normal operation, a high-flux limit at 118 % of rated, a hot-leg limit at 340 °C and the rod insertion limit at 30 % withdrawn. All three annunciate and post an explanatory note, the first two latch, and the third clears once the bank is 3 % above the limit. None of them halts the simulation, which keeps a transient legible beyond the moment a real plant would have shut itself down.

2.3.1 Point Kinetics Under the Prompt Jump Approximation

Full point kinetics carries a prompt neutron mode whose time constant near criticality is $\Lambda/(\beta - \rho)$, about 3 ms for a generation time Λ of order 2×10^{-5} s. Integrating that mode beside a xenon transient lasting ten hours is nine orders of magnitude of stiffness in one state vector. The prompt term is therefore set to zero, which is the prompt jump approximation of Section 1.4.2, here with six delayed

Table 2.1 Classification of Variables. Nominal Values Are Those of the Design Point at Full Power, Where Every Scenario Starts.

Symbol	Role	Nominal	Range
T_{avg} (TE)	CV, measured	310.0 °C	291.7–310 °C SP
v (rod drive)	MV	0 %/min	±30 %/min
C_B	MV, slow	900 ppm	400–1400 ppm
turbine load L	DV, measured	100 %	20–100 %
RTD noise	DV, unmeasured	0 °C	0–1 °C 1σ
RTD bias	DV, unmeasured	0 °C	±3 °C
flux noise	DV, unmeasured	0 %	0–2 % 1σ
dropped rod	DV, unmeasured	0 pcm	−150 pcm step
$W_1 \dots W_6$	states	β_i/λ_i	≥ 0
T_f	state	609.0 °C	—
T_c	state	310.0 °C	—
T_s	state	269.9 °C	—
z	state	60.0 %	0–100 %
I, X	states	1, 1	≥ 0
C_B	state	900 ppm	≥ 0

groups in place of one, and what survives is the algebraic relation of Equation (2.1).

$$\frac{\mathrm{d}W_i}{\mathrm{d}t} = \beta_i n - \lambda_i W_i, \qquad n = \frac{\sum_{i=1}^{6} \lambda_i W_i}{\beta - \rho}, \qquad W_i = \Lambda c_i \tag{2.1}$$

where n is nuclear power as a fraction of rated, c_i is the concentration of delayed neutron precursor group i, β_i is its delayed fraction, λ_i its decay constant, $\beta = \sum_i \beta_i$ and ρ is reactivity in $\Delta k/k$. Scaling each precursor by the generation time, $W_i = \Lambda c_i$, removes Λ from both expressions in (2.1) completely, so the model has no neutron generation time in it at all and cannot be sensitive to a quantity that is difficult to measure. Both sides of the differential form carry units of inverse minutes, since W_i is dimensionless in the same sense that n is.

The approximation is a standard one for slow transients and its error is understood, being smallest well below prompt critical and growing as ρ approaches β [18]. Equation (2.1) makes the failure visible rather than subtle, because the denominator vanishes at $\rho = \beta$. The application does not integrate through that point. It tests $\rho \geq \beta$ every step, annunciates PROMPT CRITICAL, and states that nothing past that instant is a physical prediction. The power expression is floored at $\beta - \rho \geq 0.05\beta$ so that a student driving the model toward prompt critical gets a finite number and a warning rather than an overflow. Delayed neutron data are for thermal fission of uranium-235, $\beta_i = \{215,\ 1424,\ 1274,\ 2568,\ 748,\ 273\} \times 10^{-6}$ summing to $\beta = 0.006\,502$, with decay constants $\lambda_i = \{0.0124,\ 0.0305,\ 0.111,\ 0.301,\ 1.14,\ 3.01\}\,1/\mathrm{s}$ converted to inverse minutes [7].

2.3.2 Reactivity Balance and Control Bank Worth

Reactivity is measured from the design critical condition, so every term is zero at full power and the plant starts exactly critical, as shown in Equation (2.2).

$$\rho = W(z) - W(z_0) + \alpha_F (T_f - T_{f,0}) + \alpha_M (T_c - T_{c,0}) - \rho_{Xe,0}(X - 1) + \alpha_B (C_B - C_{B,0}) + \rho_{ext} \tag{2.2}$$

where $W(z)$ is the integral worth of the bank at position z, α_F is the Doppler coefficient acting on average fuel temperature, α_M is the moderator temperature coefficient acting on average coolant temperature, $\rho_{Xe,0}$ is the equilibrium xenon worth at full power, X is the xenon concentration normalized to that equilibrium, α_B is the differential boron worth and ρ_{ext} carries a dropped rod. The

design values are $z_0 = 60\,\%$, $T_{f,0} = 609.0\,^\circ\mathrm{C}$, $T_{c,0} = 310.0\,^\circ\mathrm{C}$ and $C_{B,0} = 900\,\mathrm{ppm}$. Equation (2.2) is the reactivity ledger of Equation (1.3) with boron and an external term added, reported term by term. Power moves only when the ledger stops summing to zero, and most of the entries in it are temperatures, which is what makes the self-regulation visible.

Integral bank worth follows the classic S curve produced by the axial flux shape, given with its slope in Equation (2.3).

$$\begin{aligned} W(z) &= W_{tot}\left[\frac{z}{100} - \frac{\sin(2\pi z/100)}{2\pi}\right] \\ \frac{\mathrm{d}W}{\mathrm{d}z} &= \frac{W_{tot}}{100}\left[1 - \cos\frac{2\pi z}{100}\right] \end{aligned} \tag{2.3}$$

where W_{tot} is the total worth of the bank in $\Delta k/k$ and z is in percent withdrawn. Differential worth is zero at both ends of travel and peaks at half insertion, at 20 pcm per percent for the 1000 pcm bank used here, and 18.09 pcm per percent at the design position of 60 %. The process gain of the rod loop therefore depends on where the bank happens to be, which is a gain schedule the plant supplies whether or not the controller knows about it, and it is why a fully withdrawn bank has no authority at all.

2.3.3 Fuel, Coolant and Steam Generator Energy Balances

Three lumped thermal capacities carry the heat from the fuel to the turbine, as shown in Equations (2.4) to (2.6).

$$m_f c_f \frac{\mathrm{d}T_f}{\mathrm{d}t} = f_P P_0 n - UA_{fc}(T_f - T_c) \tag{2.4}$$

$$\begin{aligned} m_c c_c \frac{\mathrm{d}T_c}{\mathrm{d}t} &= (1 - f_P)P_0 n + UA_{fc}(T_f - T_c) \\ &\quad - UA_{cs}(T_c - T_s) \end{aligned} \tag{2.5}$$

$$m_s c_s \frac{\mathrm{d}T_s}{\mathrm{d}t} = UA_{cs}(T_c - T_s) - P_0 L \tag{2.6}$$

where P_0 is rated thermal power, f_P is the fraction of fission power deposited in the fuel rather than directly in the moderator, and L is the turbine load fraction after the ramp limiter. Every term carries units of MW, so the capacities are in MW min/°C and the conductances in MW/°C. The three time constants are $m_f c_f/UA_{fc} = 5.0\,\mathrm{s}$, $m_c c_c/(UA_{fc} + UA_{cs}) = 11.4\,\mathrm{s}$ and $m_s c_s/UA_{cs} = 10.6\,\mathrm{s}$. The disturbance L enters at the far end of that chain in (2.6). Temperature error is therefore the last quantity anywhere in the plant to register that something happened, which is the argument for a feedforward taken from an earlier link.

Hot and cold leg temperatures and the steam pressure are algebraic, as shown in Equations (2.7) and (2.8).

$$T_{hot} = T_c + \tfrac{1}{2}\Delta T_{core}\, n, \qquad T_{cold} = T_c - \tfrac{1}{2}\Delta T_{core}\, n \tag{2.7}$$

$$p_s = 0.1037\,(T_s/100)^4 \tag{2.8}$$

where ΔT_{core} is the core rise at full power and p_s is in MPa with T_s in °C. The rise scales with power in (2.7) because primary flow is constant. Equation (2.8) is an empirical saturation fit good to about 2 % between 240 and 300 °C, and it is a readout rather than a controlled variable.

2.3.4 Iodine, Xenon and Boron

Iodine-135 and xenon-135 are carried normalized so that both equal one at full-power equilibrium, as shown in Equations (2.9) and (2.10).

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \lambda_I (n - I) \tag{2.9}$$

$$\frac{\mathrm{d}X}{\mathrm{d}t} = (\lambda_X + a)\left[(1-\gamma)I + \gamma n\right] - (\lambda_X + a n)X \tag{2.10}$$

where λ_I and λ_X are the decay constants for half lives of 6.71 h and 9.21 h, a is the burnup rate $\sigma_X \Phi_0$ at rated flux, and γ is the share of xenon produced directly in fission rather than through iodine decay, taken from yields of 6.386 % for iodine and 0.228 % for xenon, so $\gamma = 0.03447$. They are Equation (1.6) with the direct yield included, and a is the burnup rate that Chapter 1 writes r_X. Setting $n = I = X = 1$ makes both derivatives vanish identically, which is the eighth of the ten checks the application runs on itself. The steady xenon concentration for a sustained power fraction is Equation (2.11).

$$X_\infty(n) = n\,\frac{\lambda_X + a}{\lambda_X + a n} \tag{2.11}$$

which returns 0.878 at 60 % power and 0.828 at 50 %. The burnup rate $a = 4.77 \times 10^{-3}$/min is 3.8 times the xenon decay constant, so at power the concentration is set mainly by neutron capture rather than by decay, and it is that asymmetry that makes xenon rise after a power reduction. Boron reaches the core through the charging line as a single mixing lag, as shown in Equation (2.12).

$$\tau_B \frac{\mathrm{d}C_B}{\mathrm{d}t} = C_{B,dem} - C_B \tag{2.12}$$

where $C_{B,dem}$ is the operator's demand in ppm and τ_B is 20 min. Chemistry is therefore a plan rather than a response, because a correction started now arrives over the following half hour.

2.3.5 Temperature Program, Rod Drive and Instruments

The setpoint is not constant. Average coolant temperature is scheduled against turbine load so that steam pressure stays in a usable range, and the rod drive integrates a commanded speed, as shown in Equation (2.13).

$$T_{SP} = T_{nl} + (T_{fl} - T_{nl})\,L + \Delta T_{SP}, \qquad \frac{\mathrm{d}z}{\mathrm{d}t} = v \tag{2.13}$$

where T_{nl} and T_{fl} are the programmed temperatures at no load and full load, ΔT_{SP} is an operator trim, and v is the applied rod speed in percent per minute. A load change therefore moves the setpoint and disturbs the plant at the same instant. The load itself reaches the machine as a ramp at the configured rate rather than as a step. Bank position is clipped to 0–100 % and the derivative in (2.13) is forced to zero at either stop whenever the commanded speed points out of travel, which stops the Runge-Kutta stages from stepping across a mechanical limit and back inside one interval.

The instrument chain is the only path by which the controller learns anything. Average coolant temperature is measured through a resistance temperature detector in a thermowell and the excore detector is filtered, as shown in Equation (2.14).

$$\tau_{TT} \frac{\mathrm{d}\hat{T}_{avg}}{\mathrm{d}t} = T_c + b + \eta(t) - \hat{T}_{avg}, \qquad \tau_{NI} \frac{\mathrm{d}\hat{n}}{\mathrm{d}t} = n + \eta_n(t) - \hat{n} \tag{2.14}$$

where b is a transmitter bias, η is colored noise on the temperature channel with a correlation time of 0.05 min and η_n is colored noise on the flux channel with a correlation time of 0.02 min, each scaled so that its stationary standard deviation is the selected noise level. Each lag is stepped with the closed-form zero-order-hold recursion $x_k = \alpha x_{k-1} + (1-\alpha)u_k$ at $\alpha = \exp(-\Delta t/\tau)$, so no sample period can destabilize it. Equation (2.14) contains the single most consequential number in the application, because τ_{TT} and not the reactor is what limits the gain.

Table 2.2 Model Parameters. All Are Displayed and All Are Adjustable, Because No Identification Experiment Depends on Concealing Them.

Symbol	Meaning	Value	Unit
P_0	rated thermal power	3411	MW
f_P	power deposited in the fuel	0.974	—
$m_f c_f$	fuel heat capacity	0.925	MW min/°C
UA_{fc}	fuel-to-coolant conductance	11.11	MW/°C
$m_c c_c$	primary heat capacity	18.33	MW min/°C
UA_{cs}	steam generator conductance	85.0	MW/°C
$m_s c_s$	secondary heat capacity	15.0	MW min/°C
ΔT_{core}	core rise at full power	33.0	°C
α_F	Doppler coefficient	−3.0	pcm/°C
α_M	moderator coefficient	−4.0	pcm/°C
α_B	differential boron worth	−8.0	pcm/ppm
W_{tot}	control bank total worth	1000	pcm
$\rho_{Xe,0}$	equilibrium xenon worth	2600	pcm
a	xenon burnup rate	4.77×10^{-3}	1/min
v_{max}	maximum rod speed	30	%/min
z_{lim}	rod insertion limit	30	%
τ_B	boron mixing time	20	min
—	turbine load ramp rate	5	%/min
T_{nl}	programmed T_{avg}, no load	291.7	°C
T_{fl}	programmed T_{avg}, full load	310.0	°C
τ_{TT}	T_{avg} RTD time constant	0.05	min
τ_{NI}	excore detector filter	0.01	min
—	high-flux annunciation	118	%
—	hot-leg temperature limit	340.0	°C

2.3.6 Parameters, Initialization and Numerics

Constants are listed in Table 2.2. None is concealed, since no identification experiment depends on hiding one, and every one is adjustable. The experiments establish what the constants do rather than identify their values.

No scenario begins with a transient already in progress. The coolant temperature starts at the programmed value, the fuel and secondary temperatures follow from Equations (2.4) and (2.6) at zero derivative, the precursors are set to $W_i = \beta_i L/\lambda_i$, iodine and xenon to their equilibrium from (2.11), boron to 900 ppm, and the bank is placed by bisection on the monotone worth curve of (2.3) at the position that cancels the temperature and xenon feedback exactly. A load chosen before Run sets that start only where such a position exists. Below about 62 % load the xenon that the lower power burns off is worth more than the bank has left, so the plant starts at the case's own load and the turbine ramps to the chosen one, as it would after Run. When a parameter edit leaves no position at either load, the routine returns 0 or 100 %, which reports a control bank of insufficient worth rather than failing silently. The first of the ten self-checks confirms that every derivative at the design point is smaller than 10^{-9} in magnitude.

States advance under fourth-order Runge-Kutta on a fixed 0.0005 min step. That is 0.03 s, comfortably shorter than the 0.33 s mean life of the sixth precursor group, which is the fastest mode left in the model once the prompt term is gone. Rod control fires on its own 0.002 min period, or 0.12 s, and twenty logged samples cover each simulated minute. The default seed is 20260829 and the five speed multipliers run from 0.5× to 20× simulated minutes per wall second, with the meaning Section 1.6.1 gives them. The xenon maneuver needs the fastest of them to cover sixteen hours.

2.3.7 Assumptions

Detail is carried where a controller has to be designed around it and omitted elsewhere. Every simplification below has a place where it stops being true.

Point kinetics with the prompt jump. There is one neutron population and no spatial dependence, and the prompt mode has been removed algebraically. The model is silent about anything approaching prompt critical, including a rod ejection, and it says so rather than integrating through it [18].

No spatial flux distribution. One node represents the whole core, so axial offset, quadrant power tilt and spatial xenon oscillation cannot appear. This is the largest omission. Those phenomena are the reason large cores need axial flux control and part-length rods, and reproducing them requires at least a two-node axial model [12].

Reactivity coefficients are constants. In a real core the Doppler and moderator coefficients both move with power, burnup and boron concentration, and the moderator coefficient changes most over a fuel cycle [9]. Rather than model that dependence, the application exposes α_M as an adjustable parameter, which shows what the variation costs and predicts nothing about any particular core.

Three lumped capacities and no transport delay. A real primary loop has an 8 to 12 s circulation time that appears as dead time in the temperature measurement, and dead time is harder on a loop than the lag that stands in for it here. Primary flow is constant, so pump trips, natural circulation and flow-induced reactivity are outside the model.

The secondary side is one saturated node. There is no drum level, no feedwater control, no steam pressure controller, no steam dump and no turbine bypass. The steam dump in particular is what a real plant would use for a large load rejection, as Chapter 5 simulates, and its absence is why the high-flux and hot-leg limits in this application annunciate rather than trip.

Boron is a first-order lag. Real boration and dilution are flow-limited operations through the chemical and volume control system, with a rate limit rather than a lag, a finite makeup inventory, and a concentration that is not uniform during the transient.

The numerics are more repeatable than a control system. Freezing the step and the execution period reproduces one trajectory on every machine, and a real rod control system is analog in parts, moves the bank in discrete steps rather than at a continuous speed, and has a deadband that is not perfectly repeatable.

2.4 Dynamic Analysis and Controller Design

2.4.1 Process Gain From Bank Position to Temperature

Every steady state of Equation (2.1) has $\rho = 0$ exactly, by the argument of Section 1.4.2 applied to each of the six groups. Steady state in this plant is therefore a statement about the reactivity ledger and not about the power, and the power itself is fixed by the turbine through (2.6), which requires $n = L$ when T_s is stationary.

Those two facts give the process gain of the rod loop in one step. Equation (2.4) at steady state fixes the fuel-to-coolant difference at $f_P P_0 n / U A_{fc}$, which depends on power alone, so at fixed load a change in coolant temperature moves the fuel temperature by exactly the same amount. Differentiating (2.2) at $\rho = 0$ with load, xenon and boron held gives Equation (2.15), the rod gain of Equation (1.4).

$$K_p \;=\; \frac{\mathrm{d}T_c}{\mathrm{d}z} \;=\; \frac{-\,\mathrm{d}W/\mathrm{d}z}{\alpha_F + \alpha_M} \tag{2.15}$$

where K_p is in °C per percent withdrawn. At the design bank position the differential worth of (2.3) is 18.09 pcm per percent and the coefficients sum to −7 pcm/°C, so $K_p = 2.58\,^{\circ}\mathrm{C}$ per percent. Withdrawing the bank makes the plant hotter and does not, at steady state, make it more powerful. That result separates this loop from one in which the manipulated variable sets the power.

The gain is scheduled by the bank itself. Equation (2.15) inherits the shape of (2.3), giving 2.86 °C per percent at half insertion, 0.99 °C per percent at 80 % withdrawn, 0.27 °C per percent at 90 %, and zero at either stop. A controller tuned at the design position is working with more than nine times the

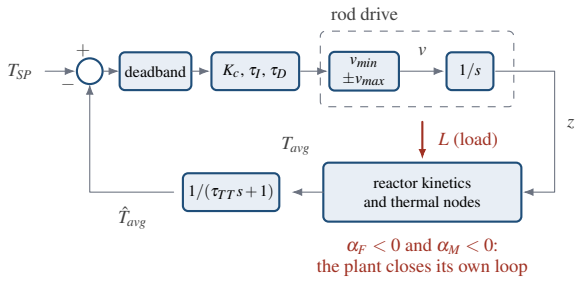


Figure 2.2 The rod control loop as implemented. The manipulated variable is rod speed, so the drive contains the integrator, and the controller ships proportional-only. The deadband acts on the error rather than on the output, which makes it a gap controller and keeps it free of windup, while the minimum speed and the maximum speed act on the output. Turbine load L enters at the steam generator and moves the setpoint through the temperature program at the same time, and the negative reactivity coefficients close a loop inside the plant block that no controller setting can alter.

gain it will have when the bank is at 90 %, which is what makes the end of the xenon maneuver in Section 2.5 difficult rather than merely slow.

2.4.2 Controller Implementation

Figure 2.2 shows the loop as coded. Inside the controller is the velocity-form PID of Section 1.6.2, differentiating the process value so that a moving setpoint cannot kick the drive, and the three drive nonlinearities sit between that output and the plant.

The deadband is applied to the error, as shown in Equation (2.16), so that the controller cannot accumulate anything inside its own band.

$$e_k = T_{SP,k} - \hat{T}_{avg,k}, \qquad e_{db,k} = \operatorname{sgn}(e_k)\max(0, |e_k| - DB) \tag{2.16}$$

Gating the output instead would leave the integral free to grow while the actuator was held still, and the loop would then arrive at the band edge already wound up. Each execution of the controller computes an increment to the rod speed rather than the speed itself [19]. The three increments are Equations (2.17) to (2.20), and the recursion that assembles them is Equation (2.21).

$$\hat{y}_k = \alpha \hat{y}_{k-1} + (1-\alpha)\hat{T}_{avg,k}, \qquad \alpha = \mathrm{e}^{-\Delta t_c/\tau_f} \tag{2.17}$$

$$\Delta v_k^P = K_c\,(e_{db,k} - e_{db,k-1}) \tag{2.18}$$

$$\Delta v_k^I = \frac{K_c}{\tau_I}\, e_{db,k}\,\Delta t_c \tag{2.19}$$

$$\Delta v_k^D = -\frac{K_c \tau_D\,(\hat{y}_k - 2\hat{y}_{k-1} + \hat{y}_{k-2})}{\Delta t_c} \tag{2.20}$$

$$\begin{aligned} s_k &= s_{k-1} + \Delta v_k^P + \Delta v_k^D \\ c_k &= u_{ff,k} + s_k + I_{k-1} \\ v_{cmd,k} &= \operatorname{clamp}(c_k,\ -v_{max},\ v_{max}) \\ I_k &= I_{k-1} + \Delta v_k^I + \frac{\Delta t_c}{T_t}\,(v_{cmd,k} - c_k) \end{aligned} \tag{2.21}$$

where $\hat{y}$ is the measurement filtered with $\tau_f = 0.05$ min for the derivative alone, I is the integral state, s_k accumulates the proportional and derivative increments without limit, c_k is the speed the algorithm demands, and T_t is the tracking time. In the symbols of Equations (1.13) to (1.16), s_k is S_k, c_k is v_k, $v_{cmd,k}$ is u_k and the feedforward u_{ff} takes the place of the bias u_b. The integral used at scan k was

accumulated through scan $k-1$. The loop is reverse acting, $a=+1$ in the convention of Section 1.3, because a temperature below its program calls for withdrawal and withdrawal makes the coolant hotter. Windup protection is the back-calculation term in the last line of (2.21), which drives I toward the value that reproduces the clipped speed, so the integral alone is corrected and the running sum s_k is never truncated. The application passes $T_t=\tau_I$ for every finite integral time, so each scan removes $\Delta t_c/\tau_I$ of the clipped excess, one part in 500 at $\tau_I=1$ min. The integral update runs only when the integral gain is nonzero. The highest available integral time, 60 min, is read as no reset at all, which removes Δv_k^I and the tracking term together. Dropping the second with the first is not cosmetic. With no integral to unwind, a live tracking term would leave a standing offset in I after each excursion against the speed limit, and the loop would come back from a large transient holding an error that nothing erases. Away from the speed limit the recursion equals the positional law with the same three gains, as Section 1.6.2 explains. The tracking time matters only in runs that reach the limit. At $\tau_I=1$ min in scenario 5 the demand is clipped for 59 % of the hour, and a counterfactual integral that absorbed the whole clipped excess at each scan still hunts, but peaks at 12.8 °C rather than 23.2 °C and spends 1137 % of rod travel rather than 1352 %. The deadband gate, minimum speed and travel stops of Equation (2.22) act after the clamp and outside the tracking, so with a finite τ_I the integral keeps accumulating while the bank stands at a stop.

In AUTO the drive applies its three limits in the order of Equation (2.22).

$$v=\begin{cases}0, & |e|\le DB,\ DB>0,\ |u_{ff}|<v_{min}\\ \operatorname{sgn}(v_{cmd})\,v_{min}, & 0<|v_{cmd}|<v_{min},\ v_{cmd}\,d>0\\ 0, & 0<|v_{cmd}|<v_{min},\ v_{cmd}\,d\le 0\\ v_{cmd}, & \text{otherwise}\end{cases} \tag{2.22}$$

where the direction d is e_{db} outside the band and u_{ff} inside it, with a further block at either stop when the commanded speed points out of travel, and $v=0$ when the drive-failure event is armed. In MANUAL the first and third cases do not apply: operator demand bypasses the temperature gate, and a nonzero demand below v_{min} is rounded up to it. On transfer from AUTO to MANUAL, a sub-minimum automatic demand is instead zeroed immediately, so a parked drive stays parked. An explicit manual demand chosen before Run is preserved. Explicit manual demands still obey maximum speed, minimum moving speed, travel stops and drive failure. The feedforward test in the first case matters only in MISMATCH, where it is what keeps the deadband from being applied twice. A load ramp produces its largest feedforward demand while the temperature error is still small, so a drive held still on the temperature error alone would discard the feedforward at the moment it is worth most. In scenario 7 with the drive of scenario 6 restored, the exception moves the rods in 103 of the 17166 in-band executions of the hour (0.6 %) at the shipped $K_{mm}=1.2$. At the normal 5 %/min ramp it never acts at $K_{mm}=1.2$ and acts in 38 of 15201 at $K_{mm}=2$. In the plant the two signals are summed ahead of the speed program, so what the drive stands still for is a small demand rather than a small error. The rod control system runs on a 0.002 min period unrelated to the 0.0005 min physics step, and it always visits six steps in one order, which is the turbine ramp limiter, the instruments of (2.14), the setpoint program, the boron demand, the controllers and finally (2.22). No controller signal ever comes from a state. In MANUAL the integral tracks the operator's demand, $I_k=v_{man}-u_{ff,k}-s_k$, and in automatic the demand follows the controller output, so transfer is bumpless both ways as Section 1.6.2 describes.

Three architectures share that machinery and are selected by an operator switch. In TEMP the temperature controller drives the rods and $u_{ff}=0$. In MISMATCH the controller bias becomes $u_{ff}=K_{mm}(L-\hat{n})\times 100$, clamped to $\pm v_{max}$, which is the difference between what the turbine is taking and what the reactor is making. Entering it as the bias rather than adding it afterward means that saturation and bumpless transfer both account for it. In CASCADE the temperature controller output is reinterpreted as a nuclear power setpoint bounded to 0–115 %, and an inner proportional controller on the excore measurement drives the rods. Changing architecture while running changes

Table 2.3 Ultimate Gain and Ultimate Period of the Proportional Rod Loop Against the T_{avg} Measurement Lag, From the Linearized Model at the Design Point. The Last Two Columns Are Measured in the Application, as the Smallest Gain on the 0.25 Grid at Which the Response to the 10 % Load Reduction of Scenario 4 Stops Decaying, and the Period of That Oscillation.

τ_{TT} (min)	τ_{TT} (s)	ω_{180} (rad/min)	P_u (min)	K_u	meas. K_u	meas. P_u (min)
0	0.0	9.34	0.67	37.3	29.75	0.75
0.02	1.2	6.40	0.98	19.2	15.75	1.07
0.05	3.0	4.17	1.51	9.01	7.5	1.62
0.10	6.0	2.75	2.29	4.18	3.75	2.36
0.20	12.0	1.88	3.34	2.09	2.0	3.35

Gains in %/(min °C).

what the outer output means, so for one execution the rod speed is held, the inner controller tracks it, and the outer controller tracks it or, in CASCADE, the measured power.

2.4.3 Frequency Response and the Case for Proportional-Only Control

Combining Equation (2.15) with the rod drive gives the loop the form of Equation (2.23).

$$\frac{\hat{T}_{avg}(s)}{v(s)} = \frac{K_p}{s} \cdot \frac{H(s)}{\tau_{TT} s + 1}, \qquad H(0) = 1 \tag{2.23}$$

where $H(s)$ collects the delayed neutron dynamics and the three thermal nodes with their own reactivity feedback already closed. The pole at the origin belongs to the actuator. It is there because z is the integral of v and because reactivity is a function of z, not because any inventory in the plant fails to self-regulate, and that is the distinction the simulator exists to make. A proportional controller on (2.23) therefore leaves no offset at all, since the loop transfer from setpoint to temperature is unity at $s = 0$ for any positive K_c.

🕮 **Pause.** Try Exercise 2.1 before reading on.

Evaluating (2.23) numerically from a Jacobian of Equations (2.1) to (2.13) taken at the design point confirms the low-frequency asymptote, since $sG(s)$ evaluates to 2.58 at $\omega = 0.001\,\text{rad/min}$ against the 2.584 that (2.15) predicts. The same evaluation gives the ultimate gain of the proportional loop, listed in Table 2.3.

Table 2.3 says that the reactor is not what limits this loop. Removing the thermowell lag entirely raises the ultimate gain from 9.01 to 37.3 %/(min °C), a factor of 4.14, and doubling the lag to 6 s lowers it to 4.18, a further factor of 2.16. The application's own runs follow the same scaling. On the load reduction of scenario 4 the oscillation stops decaying at $K_c = 7.5$ with the thermowell in place, 29.75 with it removed and 3.75 with it doubled, ratios of 3.97 and 2.0 against the 4.14 and 2.16 of the linearization. Each measured onset sits 4 to 20 % below the linearized ultimate gain, partly because the bank has settled near 56.5 % withdrawn where differential worth is 6 % higher than at the design position, and each measured period is within 12 % of P_u. At $\tau_{TT} = 0.05\,\text{min}$ the shipped $K_c = 3$ sits at a gain margin of 3.00 and a phase margin of 10.4°, and with the lag removed $K_c = 12$ sits at a gain margin of 3.11 and a phase margin of 10.2°. Those are the same loop with the same robustness and four times the gain. Raising K_c at the real thermowell walks the margin down in the way the trends show, to 1.50 at $K_c = 6$ where the response begins to ring and 1.13 at $K_c = 8$ where it hunts continuously, and at $\tau_{TT} = 0.1\,\text{min}$ a gain of 5 is already above the ultimate gain of 4.18 and cannot be stable at all.

🕮 **Pause.** Try Exercise 2.3 before reading on.

Table 2.4 Phase Margin Against Integral Time at $K_c = 3$ and $\tau_{TT} = 0.05$ min. The Phase a PI Controller Removes at Crossover Accounts for the Entire Change.

τ_I (min)	∞	6	3	2	1
ω_c (rad/min)	2.344	2.347	2.356	2.371	2.440
$\arctan(1/\omega_c\tau_I)$ (°)	0	4.1	8.1	11.9	22.3
Phase margin (°)	10.4	6.4	2.3	−1.7	−12.6
Rod travel (%)	18	28	92	828	1352

The same calculation prices reset. With $K_c = 3$ the loop crosses over at 2.34 rad/min, and a proportional-plus-integral controller subtracts $\arctan(1/\omega_c\tau_I)$ of phase at that frequency while barely moving the crossover. Table 2.4 compares that prediction against the margin computed for the full loop.

The loop has 10.4° of phase margin to spend and an integral time of 2 min spends 11.9° of it, which is why Table 2.4 changes sign between $\tau_I = 3$ and $\tau_I = 2$ min and why the measured rod travel changes character at the same place. Nothing about the gain caused this. The rod drive is an integrator and the controller added a second one, so the open loop starts at −180° and the only phase lead available comes from the reset zero arriving below crossover. Every longer integral time converges back on the proportional-only answer while spending more rod to get there, which makes the correct integral time for this loop infinite.

> ✦ **DID YOU KNOW?** **Reset Buys Rod Travel**
>
> From the simulator: in scenario 5 at $K_c = 3$ the bank travels 18 % with no reset, 92 % with an integral time of 3 min, 828 % at 2 min and 1352 % at 1 min (Table 2.4). The loop without reset already returns the temperature to its program.

2.5 Results

Table 2.5 lists thirteen scenarios in the sequence the ideas arrive, with the mechanism each exercises. Each preloads a mode, an architecture, a set of tuning constants, the drive nonlinearities, the disturbance levels and a schedule of timed events. Each case is credited when the learner performs the actions listed in the table and the simulation reaches the listed state, and Section 1.6.4 describes the briefing, the completion rules and the challenge that every case shares. The challenge is judged on the case as given, so changing a model parameter or a disturbance setting voids the attempt. Numbers quoted below are read at the end of the run length each case names, 40 min for scenarios 4, 6, 8, 9 and 12 and 60 min for 5, 7 and 13. Integrated error and rod travel come off the scorecard, and the peak deviation is the largest logged error. An integrated absolute error therefore means something as a comparison against another entry in the same table rather than on its own, while peak deviations and rod travel are not sensitive to where the run was stopped.

2.5.1 Inherent Feedback With the Bank Parked

Scenario 1 answers the question the simulator is built around before any controller is involved. The bank is parked, the turbine load demand goes from 100 to 90 % at $t = 2$ min, and nothing else is touched. Reactor power is at 91.0 % at $t = 6$ min and 89.9 % at $t = 10$ min, having moved 346 MW of thermal power with every rod stationary. What moved it is the primary warming behind a steam generator that is taking less duty, which by (2.2) removes reactivity until fission matches the new heat sink. What the plant has not done is put the temperature back. The balance closes 13.9 °C above the programmed value at $t = 10$ min, and by $t = 30$ min that deviation has drifted to 10.8 °C as xenon begins to build at the reduced power and takes over part of the job.

Table 2.5 Scenarios and Completion Requirements

#	Scenario	What it exercises	Kind	Required to complete
1	Rods parked, load to 90 %	inherent feedback with no controller	Experiment	Keep the bank parked in Manual through the load reduction at 2 min until $t = 10$ min, and then on until $t = 30$ min.
2	Drive the bank by hand	a velocity actuator under an operator	Experiment	After the load change at 2 min, move the bank with the speed demand, then from 18 min on hold it stopped for 2 min with T_{avg} within 3 °C of program.
3	First automatic loop, $K_c = 1$	no offset without integral action	Experiment	In Auto at $K_c = 1$, cut turbine load to 85–95 % after the first minute, and after a deviation of at least 1 °C hold T_{avg} within 0.2 °C of program for 2 min.
4	A proper gain, $K_c = 3$	gain limited by the measurement	Experiment	Three runs of the load reduction to 40 min, at $K_c = 3$, at $K_c \geq 6$ with the RTD lag in place, and at $K_c \geq 6$ with $\tau_{TT} = 0$, each setting held from 2.1 min.
5	Reset added, $\tau_I = 1$ min	two integrators in series	Demonstration	Three runs of the load reduction to 60 min, at $\tau_I = 1$ min, at τ_I between 2 and 6 min, and with reset removed ($\tau_I = 60$ min), each held from 2.1 min.
6	Deadband 0.8 °C, $v_{min} = 3.5\,\%/\text{min}$	offset against actuator wear	Experiment	Three runs of the load reduction to 40 min, at $DB = 0.8\,°\text{C}$ with $v_{min} = 3.5\,\%/\text{min}$, at $DB = 0.8\,°\text{C}$ with $v_{min} = 0$, and at $DB = 0.2\,°\text{C}$ with $v_{min} = 3.5\,\%/\text{min}$, each held from 2.1 min.
7	30 % ramp at 20 %/min	power-mismatch feedforward	Experiment	Three runs of the ramp to 60 min, at $K_c = 3$ with the mismatch feedforward, at $K_c = 3$ on T_{avg} alone, and at $K_c \geq 5$ on T_{avg} alone, each held from 2.1 min.
8	Dropped rod, −150 pcm	cascade on a reactivity disturbance	Experiment	Three runs through the drop at 5 min to 40 min, on T_{avg} alone at $K_c = 3$, with the mismatch feedforward, and in cascade at $K_c = 5$, $\tau_I = 4$ min and inner gain 2.5, each held from 5.1 min.
9	RTD noise $\sigma = 0.3\,°\text{C}$, steady load	derivative and deadband against noise	Experiment	Four 40 min runs at steady load, at $\tau_D = 0.4$ min with no deadband, at $\tau_D = 0$, at $\tau_D = 0.4$ min with the derivative filter at 0.2 min or more, and at $\tau_D = 0.4$ min with $DB = 0.8\,°\text{C}$, each held from 0.5 min.
10	60 % load for ten hours	the xenon transient	Experiment	Run the whole 960 min maneuver through both scheduled load changes.
11	Load to 40 %	the rod insertion limit and boron	Demonstration	Run until the bank reaches the insertion limit, then change the boron demand and stay in Auto for 60 min after the change.
12	$\alpha_M = +0.5\,\text{pcm}/°\text{C}$	process gain over a fuel cycle	Demonstration	Three runs of the load reduction to 40 min, in Auto at $\alpha_M = +0.5\,\text{pcm}/°\text{C}$, with the bank parked in Manual at the same coefficient, and in Auto at $\alpha_M \leq -4\,\text{pcm}/°\text{C}$, each held from 2.1 min.
13	Graded challenge, three events	the whole workflow under time	Challenge	Over the 60 min run hold $\lvert T_{avg} - T_{SP} \rvert \leq 2\,°\text{C}$ throughout, the bank above the 30 % insertion limit and total rod travel under 60 %, with the model parameters and disturbances as given.[a]

[a] With the shipped settings (P-only at $K_c = 1$ on T_{avg} alone with the real drive) the worst T_{avg} error is 8.10 °C and rod travel 71 %, so the run misses the error and travel targets by 6.1 °C and 11 % while the bank stays above 49.0 %. With the suggested settings (cascade at $K_c = 8$, $\tau_I = 3$ min, inner gain 4, deadband and minimum speed off) the worst error is 1.74 °C and rod travel 51.6 %, a margin of 0.26 °C and 8.4 %, with the bank never below 49.6 %. The suggested settings carried through a Reset give the same run.

Table 2.6 Proportional Gain Against the 10 % Load Reduction of Scenario 4, With the Deadband and Minimum Speed Disabled and $\tau_{TT} = 0.05$ min, Read at $t = 40$ min. Gain Margin Is From Table 2.3.

K_c (%/(min °C))	1	3	5	6	8
Peak deviation (°C)	4.08	1.58	0.99	0.83	1.26
IAE (°C min)	15.3	5.7	1.6	2.8	19.5
Rod travel (%)	15.4	17.4	8.7	17.2	153
Swing, 30–40 min (°C)	0.004	0.014	0.009	0.080	2.51
Gain margin	9.01	3.00	1.80	1.50	1.13
Behavior	slow	settles	settles	rings	hunts

TRY IT IN THE APP **1 · The Plant With Nobody Driving (Manual)**

Load scenario 1 and press Run. Leave the rod speed at zero and watch reactor power after the turbine load drops to 90 % at 2 min. Power reaches 91.0 % at 6 min and 89.9 % at 10 min with every rod still, while the average coolant temperature settles 13.9 °C above its program.

Removing that residual is the entire task of the rod control system, and Equation (2.2) prices it. Once the temperature is back on its program the core is cooler than the reference condition everywhere, by 1.83 °C in the coolant and 31.7 °C in the fuel, so the two temperature terms hand back $\alpha_F(-31.7\,^\circ\mathrm{C}) + \alpha_M(-1.83\,^\circ\mathrm{C}) = +102.4$ pcm that the bank has to absorb. At 18.09 pcm per percent that is under six percent of travel. Scenario 2 hands the job to an operator. The manual demand is a speed, so holding a temperature means integrating by hand while the plant corrects itself underneath, and inserting too far takes power below the load and calls for the rods to be withdrawn again.

Pause. Try Exercise 2.2 before reading on.

2.5.2 Proportional Gain and the Measurement Lag

Scenario 3 puts the loop in automatic at $K_c = 1\,\%/\mathrm{min}$ per degree with no reset at all and the drive nonlinearities switched off. The error returns to within 0.05 °C of zero by $t = 40$ min, which is the result (2.23) predicts and the one students expect a proportional controller not to produce, and the small residue is the slow xenon drift that the loop follows with a lag. What the low gain costs is speed. Table 2.6 sweeps the gain against the 10 % load reduction of scenario 4.

Peak deviation falls from 4.08 to 0.99 °C and the integrated error from 15.3 to 1.6 °C min between $K_c = 1$ and $K_c = 5$, and then the improvement stops. At $K_c = 6$ the trend rings and is still swinging 0.08 °C peak to peak in the last ten minutes, and at $K_c = 8$ the loop hunts with a growing swing, 2.5 °C by the end of the run, and spends 153 % of rod travel on a disturbance that needs 17.4 %. The shipped gain is lightly damped, as its 10.4° phase margin predicts. At $K_c = 3$ the error rebounds to 0.91 °C below program 1.7 min after its 1.58 °C peak and rings down over the next twenty minutes, while $K_c = 5$ rebounds by 0.10 °C and spends 8.7 % of rod travel instead of 17.4 %. Table 2.6 sets the measured behavior beside the computed gain margin, and the two agree without any fitting. Setting τ_{TT} to zero and running the sweep again shifts the whole pattern by about four, as Table 2.3 measures, so the same loop settles cleanly at $K_c = 12$, and doubling the lag to 6 s makes $K_c = 5$ hunt. Three seconds of thermowell lag decides how much gain this loop can carry, and the experiment supplies that number before any further filtering of a noisy signal is proposed.

The bank finishes at 56.5 % at $t = 40$ min for every gain that settles, because the end state is fixed by (2.2) and only the path there changes. That is 3.5 % of travel rather than the 5.7 % the load change alone would need, since 64.4 pcm comes from the rods and the balance from the xenon that has already started to build.

Table 2.7 Deadband and Minimum Speed on the Same 10 % Load Reduction at $K_c = 3$ (Scenario 6, Read at 40 min), and on Forty Idle Minutes With $\sigma = 0.3\,°C$ of RTD Noise and No Derivative (Scenario 9).

DB (°C)	0	0.8	0.8	0.2
v_{min} (%/min)	0	3.5	0	3.5
IAE, load step (°C min)	5.7	18.4	32.1	14.7
Rod travel, load step (%)	17.4	17.2	12.0	105
Rod travel, idle and noisy (%)	30	0	0	101
Steady deviation (°C)	0	≤ 0.8	0.8	cycles

2.5.3 Deadband and Minimum Rod Speed

Scenario 6 restores the drive a plant actually has, a 0.8 °C deadband and a 3.5 %/min minimum speed, and Table 2.7 separates what each of the two settings does.

The deadband turns a loop with no offset into one that parks up to 0.8 °C off program, and the integrated error on the load step rises from 5.7 to 18.4 °C min. What it buys is invisible on a clean load change, where rod travel is 17.2 % against 17.4 % without it, because the minimum speed gives back what the band saves. Dropping the minimum speed to zero with the band still at 0.8 °C recovers a 31 % saving, from 17.4 to 12.0 % of travel, at the price of 32.1 °C min of integrated error, because the drive then creeps to the edge of the band and stops there. The band earns its keep in the third row of Table 2.7, where forty minutes of a noisy measurement at steady load cost 30 % of rod travel with the band off and none with it on, because no excursion of the noise is large enough to leave the band. Over seven further noise seeds the first figure ranges from 27 to 36 % and the second never exceeds 0.3 %.

> **DID YOU KNOW? The Same Band in a Real Fleet**
>
> In the French 900 MWe pressurized water reactors the control rods move in only when the average coolant temperature stands more than 0.8 °C above its reference, and out only when it stands more than 0.8 °C below [11]. That is the deadband of scenario 6.

The fourth column is the failure the two settings produce together. A 0.2 °C band with a 3.5 %/min minimum speed cannot land inside itself. Once the drive starts it cannot creep, so it crosses the band, the error changes sign, the drive comes back, and the loop hunts for the rest of the run, with a 1.20 °C peak-to-peak cycle of 1.7 min period and 105 % of rod travel by forty minutes on the load step, and 101 % in forty idle minutes on the noisy measurement. The rule the two settings enforce is that the full band width $2\,DB$ must exceed the smallest temperature excursion the drive can produce, and the measurement supplies that number directly. A 1.2 °C cycle corresponds through (2.15) to 0.13 min of motion at the minimum speed, which is the same order as the 0.05 min thermowell and the 0.19 min coolant node together, so the excursion is what the loop cannot stop before it happens. At $2\,DB = 1.6\,°C$ the cycle fits inside the band and dies, and the last rod motion is at 33.2 min. At $2\,DB = 0.4\,°C$ it does not, and no tuning removes it. The boundary sits where the rule puts it, since a 0.5 °C band still cycles at forty minutes and a 0.7 °C band stops moving at 34.8 min.

> **TRY IT IN THE APP 6 · The Real Rod Drive: Deadband and Minimum Speed**
>
> Load scenario 6 and run it to 40 min with its 0.8 °C deadband and 3.5 %/min minimum speed. The temperature parks inside the band for 17.2 % of rod travel. Set the deadband to 0.2 °C and run again: the loop hunts with a 1.20 °C peak-to-peak cycle of 1.7 min period and spends 105 % of rod travel by 40 min.

The direction test in the second and third cases of (2.22) exists because of the mismatch feedforward and the cascade's inner loop. Just outside the band the demand can be a few thousandths of a percent per minute whose sign alternates between executions, since each rod step moves the flux and

Table 2.8 Three Architectures Against Two Disturbances. The Ramp Is 30 % of Load at 20 %/min and the Reactivity Step Is a Dropped Rod Worth −150 pcm. Cascade Uses $K_c = 5$, $\tau_I = 4$ min and an Inner Gain of 2.5. Ramp Runs Are Read at 60 min and Dropped-Rod Runs at 40 min.

	T_{avg} only	+ mismatch	cascade
Turbine ramp, 30 % at 20 %/min			
Peak deviation (°C)	6.29	4.64	5.37
IAE (°C min)	41.1	10.1	23.7
Rod travel (%)	123	30.7	24.7
Dropped rod, −150 pcm			
Peak deviation (°C)	5.93	3.58	2.23
IAE (°C min)	17.9	4.3	2.1
Rod travel (%)	53.5	14.2	19.9

both signals answer it within one execution. Rounding every such demand up to $\pm 3.5\,\%/\text{min}$ made the drive step back and forth, and over the sixteen hours of scenario 10 the rod travel count reached 2667 % against 225 % of logged bank motion. With the direction test the same run reverses the drive four times and counts 154 %, equal to the bank's path, and scenario 11 counts 87 % by 245 min.

2.5.4 Feedforward and Cascade Against Two Disturbances

Two disturbances enter this plant in completely different places. A turbine ramp arrives at the far end of the thermal chain in (2.6) and reaches the temperature last, while a dropped rod arrives directly in (2.2) and appears in the flux immediately. Table 2.8 runs all three architectures against both.

On the ramp the mismatch feedforward takes the peak deviation from 6.29 to 4.64 °C, the integrated error from 41.1 to 10.1 °C min and the rod travel from 123 to 30.7 %, which is three quarters off both the error and the actuator wear. Cascade uses the least rod on the ramp, 24.7 %, but returns the temperature slowly, at 5.37 °C and 23.7 °C min, because its inner loop holds power close to the load and the excess heat leaves only through the small power offset the outer loop allows. Raising the gain instead does not substitute for the feedforward. At $K_c = 5$ on temperature alone the peak deviation is better at 3.91 °C, but the loop has spent 383 % of rod travel by 60 min and is still swinging 0.8 °C peak to peak two hours after the ramp finished, between 120 and 130 min, because feedback speed is bounded by Table 2.3 while feedforward does not go round the loop at all. The same run also shows a gain comfortable at full power becoming uncomfortable at 70 %, which is Equation (2.15) evaluated at a different bank position.

On the reactivity step the ordering changes. Cascade gives 2.23 °C against 3.58 for the feedforward and 5.93 for temperature alone, because the excore detector sees the disturbance before the primary has changed temperature at all. Neither structure wins both columns, and the architecture question is always where the disturbance enters rather than which arrangement is more advanced. Volume 3, Chapter 2 reaches the same verdict on a tank whose two disturbances reach the level by different routes. The cascade also reverses the lesson of Table 2.4. In cascade the rod drive integrator is enclosed by the inner flux loop, so the inner loop is proportional-only and the outer loop needs an integral time near 3 to 4 min to reposition the power setpoint. The same plant gives opposite answers about integral action depending on which loop encloses the actuator, and that is a structural question rather than a tuning one. Figure 2.3 draws Table 2.8 as bars.

Scenario 9 separates what the deadband and the derivative each contribute. With 0.3 °C of RTD noise at steady load, derivative action at $\tau_D = 0.4$ min with the band off costs 108 % of rod travel in forty idle minutes, removing the derivative drops that to 30 %, raising the derivative filter to 0.2 min instead gives 65 %, and switching the 0.8 °C band back on gives zero. Over eight noise seeds the four figures span 98–111, 27–36, 57–67 and 0–0.3 %, so the ordering does not depend on the draw. The two defenses differ in kind, because a filter attenuates what the controller sees and adds the

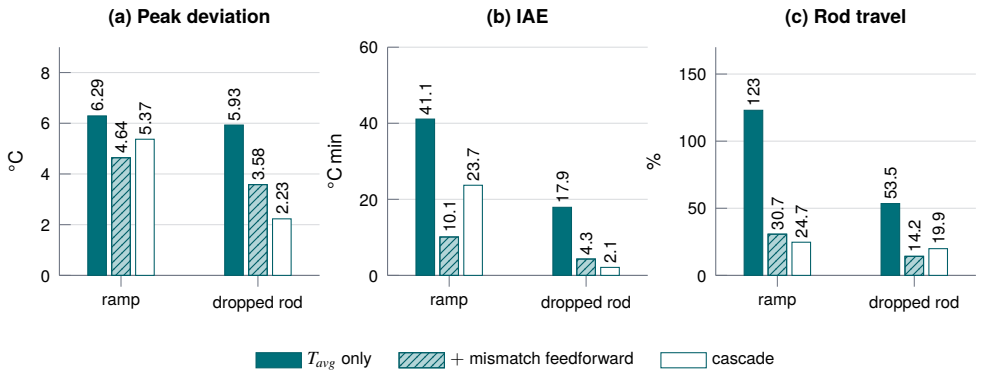


Figure 2.3 The three architectures of Table 2.8 against the 30 % turbine ramp at 20 %/min and the −150 pcm dropped rod: (a) peak deviation, (b) integrated absolute error and (c) rod travel. Solid: T_{avg} only. Hatched: with the power-mismatch feedforward. White: cascade. Notice that the feedforward gives the smallest peak and error on the ramp, which enters at the far end of the thermal chain, and the cascade the smallest on the dropped rod, which the ex-core detector sees first.

measurement lag that Table 2.3 prices, while a deadband stops the actuator responding to anything small whatever its source and pays for it in offset.

2.5.5 Xenon Transient and the Rod Insertion Limit

Scenario 10 reduces the load to 60 % for ten hours and returns it, and the bank traces 60, 39, 68, 60, 100 and finally 44 % over sixteen hours. The first move is the load change, which inserts the bank to 39.5 % eight minutes after the reduction. Then it reverses, because Equation (2.10) keeps producing xenon from the iodine already in the core while removing the flux that destroys it, the mechanism of Section 1.4.3. The concentration peaks at 121 % of its full-power equilibrium 5.2 hours after the reduction, worth −548 pcm, and the bank has to stand at 68 % to cancel it. Over the next five hours xenon decays toward the 87.8 % equilibrium that (2.11) gives for 60 % power and the bank returns to 59.7 % by the time the load comes back.

The return to full power is the more demanding half of the maneuver. Reactivity for the load increase and reactivity for the xenon that is now burning out are needed together, and eight minutes after the return the bank reaches 100 %, fully withdrawn, where Equation (2.15) gives a process gain of exactly zero and the controller has no authority left at all. It stays there for thirteen minutes while the average coolant temperature falls 7.9 °C below its program. By the end of the run, five hours after the return, xenon is down to 87.5 % of equilibrium, its ledger entry reads +325 pcm that nothing asked for, and the bank has come all the way down to 43.6 %. A plant would borate before the increase so that the bank has somewhere to go, which is why a return to full power is planned hours ahead rather than commanded on the minute, and it is the operational fact behind the load-following constraints collected in the OECD review [11].

Scenario 11 finds the constraint that the controller cannot see. A reduction to 40 % asks for about 615 pcm, while the bank holds 545 pcm between its starting position and the insertion limit at 30 % withdrawn. The bank alarms at $t = 16.2$ min and bottoms out at 26.8 % a minute later, three percent inside a margin the safety analysis assumes is there, and no tuning prevents it because nothing in the controller's objective mentions shutdown margin. What rescues it in the simulation is not the operator but xenon, which adds negative reactivity over the following hours and withdraws the bank on its own, to 33.8 % by $t = 30$ min, 43.9 % by 60 min and 57.7 % by 120 min. The deliberate route is chemistry, and its sign is the reverse of the one usually expected. Boration removes reactivity, so the bank must withdraw to replace it, and boration therefore raises the bank while dilution lowers it.

Raising the boron demand from 900 to 920 ppm at the alarm puts the bank at 41.9 % twenty minutes later, against 36.2 % without it. Setting the demand to the concentration that would make the plant critical at the present bank position is no help in automatic, because the bank already holds the reactor critical, and at the alarm that concentration is 902 ppm. Room in the bank is a target the operator has to choose.

2.5.6 Sensitivity to the Moderator Temperature Coefficient

Scenario 12 sets the moderator coefficient to $+0.5\,\mathrm{pcm}/^{\circ}\mathrm{C}$, which is where a core sits at the beginning of a fuel cycle when the boron concentration is high [9]. The plant still follows the turbine, settling at 89.6 % with the rods parked, because the Doppler coefficient acts on fuel temperature and fuel temperature is proportional to power. What changes is the price. At $t = 10\,\mathrm{min}$ the parked plant stands 32.1 °C above program instead of 13.9 °C, because the coolant term in (2.2) opposes the correction instead of helping it, and the deviation is still climbing, to 34.1 °C at 14.4 min. On the way the hot leg passes its 340 °C trip setting at 4.9 min and peaks at 357 °C, and only xenon brings the deviation back, to 17.5 °C at an hour. The loop notices too. Equation (2.15) has $\alpha_F + \alpha_M = -2.5\,\mathrm{pcm}/^{\circ}\mathrm{C}$ instead of −7 pcm/°C, so the process gain rises by a factor of 2.8, and the tuning that was comfortable in Table 2.6 leaves the loop swinging 12.6 °C peak to peak in the last ten minutes of a forty-minute run, with a peak deviation of 6.4 °C, 76 °C min of integrated error and 225 % of rod travel. At $\alpha_M = -4\,\mathrm{pcm}/^{\circ}\mathrm{C}$ the same tuning costs 5.7 °C min and 17 % of travel, and at −8 pcm/°C it costs 4.0 °C min and 12 %. The requirement that the coefficient be negative at power is a control requirement as much as a safety one, and the loop gets monotonically easier as it becomes more negative.

2.5.7 Graded Challenge

The challenge applies a 20 % load reduction at $t = 4\,\mathrm{min}$, an unannounced dropped rod at $t = 15\,\mathrm{min}$ and a return to full load at $t = 28\,\mathrm{min}$, and asks for T_{avg} held within $\pm 2\,^{\circ}\mathrm{C}$ with the bank never below 30 % and total rod travel under 60 %, judged over a 60 min run. Left as shipped, on the cautious gain of scenario 3 and the real drive, the schedule costs 8.1 °C of peak deviation and 71 % of travel, so it meets only the bank target, with the bank never below 49.0 %. With the deadband and minimum speed switched off, cascade at $K_c = 8$, $\tau_I = 3\,\mathrm{min}$ and an inner gain of 4 gives the smallest peak deviation, 1.74 °C for 51.6 % of travel and 15.7 °C min, and passes. The mismatch feedforward at $K_c = 5$ and $K_{mm} = 2$ on the same drive peaks higher at 2.41 °C, which misses the target, and wins on the other two at 8.5 °C min and 41.1 %. No architecture wins every number, which is why the scorecard shows the competing quantities rather than one composite score. The outer loop of the cascade never reaches its bounds, so the tracking time plays no part in this result.

The scorecard reports quantities that compete. Alongside integrated error, the present temperature error and time at the speed limit it carries bank margin above the insertion limit, total rod travel, the power mismatch between turbine and reactor, and the xenon concentration. Rod travel and the total variation of the rod speed demand are separate quantities, the first being $\int |v|\,\mathrm{d}t$ in percent of full travel and the second $\sum_k |v_k - v_{k-1}|$ in %/min. Ten verification tests run inside the application itself, covering the steady state at the design point, exact criticality at rated power, the sign of the rod, fuel and moderator reactivity terms, the inherent response of the secondary to a reduced steam demand, the monotonicity and peak slope of the worth curve of (2.3), the sign of the controller action, the xenon stationarity of (2.10), the reduced equilibrium of (2.11), and the sign of the feedforward gain. Figure 2.4 shows scenario 4 part way through its run.

2.6 Technology in Practice

The tags of Figure 2.1 are, in a plant, resistance thermometers on the hot and cold legs, ex-core neutron detectors beside the vessel, a measurement of turbine load and the drives of the control bank. Figure 2.5 draws the rod loop with that hardware in place of the schematic's symbols.

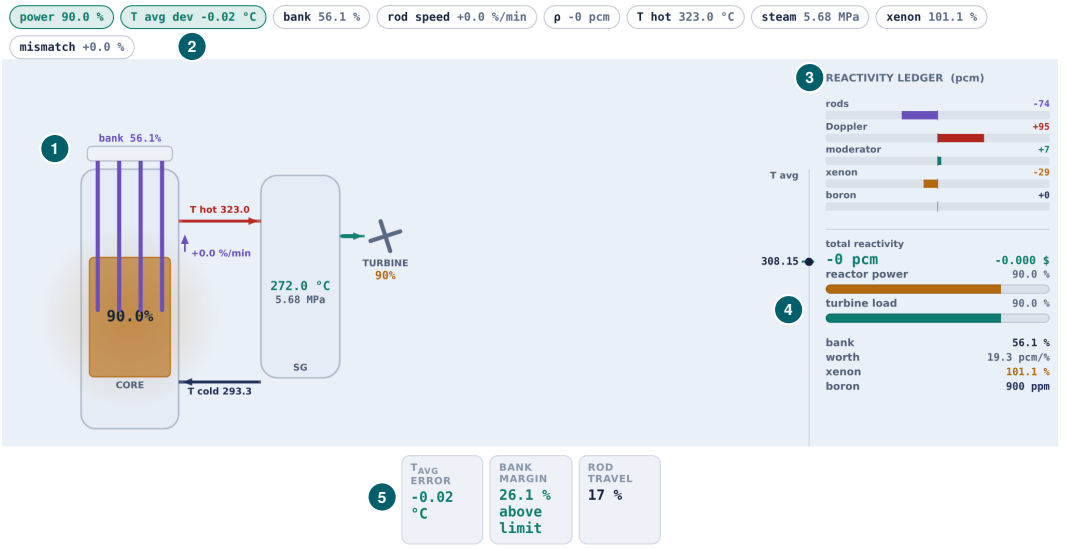


Figure 2.4 Scenario 4 at $t = 30$ min, twenty-eight minutes after the turbine load demand went from 100 to 90 % with $K_c = 3$ and no reset. The cut keeps the plant view with its badge row and three tiles of the control performance panel. Callouts mark (1) the control bank, (2) the deviation of T_{avg} from its program, (3) the reactivity ledger, (4) reactor power against turbine load and (5) rod travel and bank margin. The reactivity ledger reads rods −74 pcm, Doppler +95 pcm, moderator +7 pcm and xenon −29 pcm for a total of zero. With the temperature back on its program the core is cooler than the reference condition, so the two temperature terms hand back 102 pcm for the bank to absorb, and the rods carry less than three quarters of it because xenon has already begun to build. Nuclear power and turbine load have both arrived at 90.0 %, the average coolant temperature is 0.02 °C below its programmed 308.17 °C, and the bank is at 56.1 % and still drifting outward as xenon builds toward its new equilibrium. Rod travel is 17 % and bank margin above the insertion limit is 26.1 %.

What Measures It

The controlled variable is the average of two leg temperatures, which the element TE of Equation (2.14) reports behind the thermowell lag τ_{TT}. Thermocouples and resistance temperature detectors are the most common temperature sensors in nuclear plants, mounted in contact with the fluid or in thermowells [20]. The second measurement of the loop, nuclear power for the mismatch feedforward and the inner loop of the cascade, comes from ex-core neutron detectors, which give a global or regional flux as a representation of reactor power. In the power range they are typically uncompensated ion chambers [20]. The chapter filters that signal through τ_{NI}.

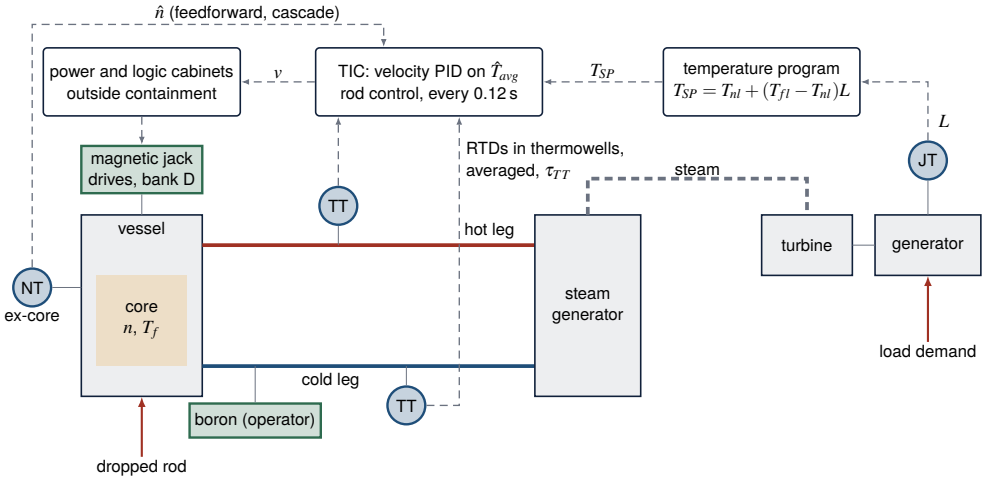


Figure 2.5 The rod control loop of Figure 2.1 with the hardware of a plant. Resistance thermometers (TT) in thermowells on the hot and cold legs give $\hat{T}_{avg}$, an ex-core detector (NT) gives nuclear power, and a load measurement (JT) schedules the setpoint. The velocity PID writes a rod speed every 0.12 s to the power and logic cabinets, which step the magnetic jack drives on the vessel head. Boron is the operator's slow second element. Thick solid lines carry primary coolant, the dashed line steam, and thin dashed arrows signals.

> **TECHNOLOGY** **Sensor: Resistance Thermometers in Thermowells**
>
> The coolant temperature is read by resistance temperature detectors, fine wires of pure metal, usually platinum, whose resistance rises with temperature. They are relatively insensitive but very linear [21]. In a plant they sit in thermowells that protrude into the pipe [20], so an element can be replaced without depressurizing the circuit. Every thermowell adds a first-order lag by its mass, and a lag large enough can compromise the stability of the loop. An element not seated at the bottom of the well leaves an air gap that adds a second lag and more phase shift [21], so a loop that settled may begin to oscillate at the same gain. The chapter's element averages the two legs behind the 3 s lag τ_{TT}.

What Moves It

Neutron-absorbing control rods regulate power by controlling the neutron flux [20], and the chapter moves one bank of them. The rod control system has a power part that converts three-phase alternating current to the direct current of the drive coils, and a control part that sequences the groups so that reactivity changes smoothly and continuously [22]. The chapter's continuous rod speed, with its minimum moving speed and its deadband, stands in for that stepping sequence, as Section 2.3.7 states. Boron is the second and slower final element, set by the operator through the charging line.

TECHNOLOGY Final Element: The Magnetic Jack Rod Drive

The rod drives are flanged on top of the reactor vessel head. A common design is the magnetic jack: stator coils outside the pressure housing make magnetic fields that work latches inside it, and energizing the coils in sequence steps the rod into or out of the core [22]. Position is known twice, from a detector on the housing and from a count of the pulses sent to the drive, and after a dropped or stuck rod the two disagree. Failing or drifting power supplies, and gripper coils fed at the wrong voltage, have caused dropped and slipped rods [22]. A dropped rod reaches the loop as the reactivity step of scenario 8.

Where the Controller Runs

The power and control cabinets of a rod drive system sit outside containment, in a more controlled environment than the drives on the vessel head [22]. Control loops of this kind run with cycle times from milliseconds to minutes [23]. The chapter's rod control fires every 0.12 s above a 0.03 s physics step, and the temperature program, the mismatch feedforward and the velocity PID of Equation (2.21) all run in that one execution, ahead of the drive limits of Equation (2.22). In the first French units the same structure was a rod program in which the rods regulated the average coolant temperature [11].

IN PRACTICE Count Every Step

Following load with the rods is likely to shorten the life of the drive mechanisms, and a worn drive can leave rods inoperable or slow to insert on a trip [10]. Plants that follow load therefore keep a monitoring system that counts and records rod motion in steps against the design life of the drives, and they maintain the drives more often. One national practice rotates the duty among the rods of the banks, and another replaces the drives before they reach the end of their life [10]. The rod travel on the chapter's scorecard is that count.

What Is Modernizing It

Three developments touch this loop. The first is the move from analog to digital instrumentation and control, which new plants make whole and operating plants make one system upgrade at a time [20]. The second is condition monitoring. An aging study of rod drive systems recommended monitoring techniques with trendable results that can detect aging [22], and plants that follow load count rod steps against the design life of the drives [10]. The third is control of the power distribution together with the temperature. In a later French mode all banks regulate the temperature while two overlapping groups hold the axial offset, selected from on-line measurements of the core [11]. That spatial problem is the one the chapter's single core node leaves out.

2.7 Limitations

No claim about a real core can rest on this model. The largest omission is spatial, because one node represents the whole reactor, so axial offset, quadrant power tilt and spatial xenon oscillation cannot appear at all and would need at least an axially divided model [12]. The application can show a xenon transient and cannot show a xenon oscillation, and the distinction matters because the second is what axial flux control exists for. The prompt jump approximation of (2.1) puts every fast reactivity event out of reach, including a rod ejection. The secondary side is one saturated node with no feedwater control and no steam dump, and there is no protection system, so the high-flux and hot-leg limits annunciate and the transient continues, which is a teaching convenience and not a statement about what a plant would do. The primary loop has no transport delay, and the 8 to 12 s circulation time of a real plant would appear as dead time in (2.14), which is harder on a loop than the lag standing in for it and would lower every entry in Table 2.3. Only PID with two feedforward structures is implemented, so the model predictive and sliding-mode formulations applied to this same plant in the literature [14, 15] are outside its scope. Two of the disturbances are stylized. A gripper coil failure does not deliver an exact −150 pcm step, but that step is the cleanest available disturbance entering as reactivity rather than as heat, which is what Table 2.8 needs it for. A real machine would not be

asked for a 20 %/min turbine ramp either, although nothing smaller makes the mismatch signal large enough to read off a trend. The six features kept through the simplification are the integrator inside the actuator, the deadband, the minimum step, the speed limit, the finite bank travel and the two time scales of temperature and xenon. Each one is there because a controller cannot be designed without accounting for it. Section 2.3.7 states where each assumption stops holding.

> **OPEN CHALLENGE** **Following Load Late in the Fuel Cycle**
>
> Late in a fuel cycle the boron in the coolant is highly diluted, so each change of power needs large amounts of dilution and makes large volumes of borated wastewater. Tank and evaporator capacity can then limit flexible operation, and operators may curtail it as the cycle ends [10]. Grey rods change power with less disturbance to the power distribution, but many plants have none [10]. Progress would be a rod program that carries the daily cycle to the end of the fuel cycle without large boron changes. In the chapter, boron is a lag with no inventory, and scenario 11 calls on it freely.

2.8 Discussion

Two consequences of this loop reach past the control room. The first is that a constraint can exist for reasons the controller has no representation of. The rod insertion limit is not a process limit and is not there to protect the temperature. It is there so that the rods still withdrawn can shut the reactor down, and the number follows from a safety analysis carried out outside the control room. A well-tuned controller holding the programmed temperature perfectly will pass straight through it, as scenario 11 does at 26.8 % withdrawn, because nothing in the objective mentions shutdown margin. Putting bank margin on the scorecard beside temperature error is the cheapest available fix, and the general shape, which is a controller approaching a limit that comes from outside the loop with a slower and less convenient manipulated variable available behind it, is the ordinary form of override control in refining, power and mining. Volume 4, Chapter 2 treats a protective override of that kind, and Volume 4, Chapter 9 a valve-position controller that returns the load to the slower and cheaper resource, as boron relieves the bank here.

The second is that automation built on a physical property inherits the obligation to keep verifying that the property still holds. The negative moderator coefficient is a licensing requirement and it is also what makes this loop tractable, and it moves measurably over every fuel cycle [9]. Scenario 12 shows the same controller oscillating on a plant whose coefficient has gone slightly positive. The honest answers are gain scheduling, periodic retuning, or a design robust across the range, and choosing among them requires knowing the range, which is a measurement rather than an assumption.

Both questions are becoming more practical. Least-cost studies of deeply decarbonized systems find the value of nuclear capacity to depend on whether it can follow [17], and whether a particular plant can do so is a question about xenon, boron chemistry, rod worth and insertion limits rather than about the controller [11]. The application contains exactly those four constraints, and at 20× speed, twenty simulated minutes per second, the sixteen-hour maneuver runs in under a minute.

> **OPEN CHALLENGE** **Replacing Analog Rod Control**
>
> Analog instrumentation and control is going obsolete, and new plants will use fully digital systems. Operating plants move to digital one system upgrade at a time, constrained by licensing concerns, and the experience base in nuclear plants is still limited [20]. Safety systems and high-consequence control systems generally use only mature, well-proven technology, while a plant outlives several generations of its instrumentation [20]. Progress would be a digital rod control that can be shown, before it is licensed, to keep the behavior of the system it replaces, including its deadband and minimum speed.

2.9 Conclusion

This chapter studied a pressurized water reactor, a plant that performs its own load change. The simulator integrates thirteen states, six delayed neutron precursor groups under the prompt jump

approximation, three lumped thermal capacities, the control bank position, iodine, xenon and boron, with the neutron generation time cancelled out of the model entirely and the region of validity annunciated rather than extrapolated. Reactivity feedback takes reactor power from 100 to 89.9 % within eight minutes with every rod stationary and leaves the average coolant temperature 13.9 °C above its program, which is the whole of what the rod control system has to remove.

Because rod speed and not rod position is the manipulated variable, the process gain from bank position to temperature is 2.58 °C per percent in series with a pole at the origin that belongs to the actuator, the correct integral time is infinite, and an integral time of 2 min spends 11.9° of a phase margin that is only 10.4° wide. The ultimate gain of the loop is 9.01 with the three-second thermowell in place and 37.3 with it removed, so the measurement and not the reactor sets the achievable gain, and the application’s own runs stop decaying at 7.5 and 29.75 in the same two cases.

A 0.2 °C deadband against a 3.5 %/min minimum speed hunts for as long as the run lasts, at 105 % of rod travel in forty minutes, while 0.8 °C contains the same excursion and removes all noise-driven motion. A sixteen-hour maneuver to 60 % load and back walks the bank through 60, 39, 68, 60, 100 and 44 % and exhausts its withdrawal travel, and a reduction to 40 % takes the bank through a licensing limit that no tuning respects.

A fuller model would add an axial division so that xenon oscillation and axial offset control become reachable, a transport delay around the primary loop, a steam dump on the secondary side, and coefficients that move with burnup. The simulator separates what a plant does for itself from what its actuator makes difficult.

KEY POINTS

- A pressurized water reactor follows a cut in steam demand by itself: with every rod still, power goes from 100 to 89.9 % within eight minutes.
- The plant is left 13.9 °C above its temperature program, and removing that is the whole task of the rod control system.
- Because the drive takes a speed, a proportional controller leaves no offset, and the correct integral time for this loop is infinite.
- The measurement and not the reactor sets the achievable gain: the ultimate gain is 9.01 with the thermowell and 37.3 without it.
- A deadband narrower than the smallest excursion of the drive hunts for as long as the run lasts.

Where else this applies. A drive that moves in fixed steps inside a band also sets the tap changer of Volume 17, Chapter 8, and the offset that inherent feedback leaves returns in the governed grid of Volume 1, Chapter 8.

References

[1] M. Zarei, “On the inherent load following features of a small modular nuclear reactor,” *Nuclear Engineering and Design*, vol. 370, p. 110896, 2020.

[2] T. W. Kerlin and B. R. Upadhyaya, “Pressurized water reactors,” in *Dynamics and Control of Nuclear Reactors*. London: Academic Press, 2019, pp. 139–166.

[3] S. Jeon, “Integrator leakage for limit cycle suppression in servo mechanisms with stiction,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 134, no. 3, p. 034502, 2012.

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

[5] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[6] M. A. Schultz, *Control of Nuclear Reactors and Power Plants*, 2nd ed. New York: McGraw-Hill, 1961.

[7] M. Ash, *Nuclear Reactor Kinetics*, 2nd ed. New York: McGraw-Hill, 1979.

[8] J. J. Duderstadt and L. J. Hamilton, *Nuclear Reactor Analysis*. New York: John Wiley & Sons, 1976.

[9] M. Antonopoulos-Domis and C. Housiadas, “Moderator temperature coefficient of reactivity in pressurized water reactors: Theoretical investigation and numerical simulations,” *Nuclear Science and Engineering*, vol. 132, no. 3, pp. 337–345, 1999.

[10] 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.

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

[12] A. Chakraborty, S. Singh, and M. P. S. Fernando, “An improved reduced order model for nonlinear stability analysis of spatial xenon oscillations,” *Progress in Nuclear Energy*, vol. 116, pp. 62–75, 2019.

[13] M.-S. Yim and J. M. Christenson, “Application of optimal control theory to a load-following pressurized water reactor,” *Nuclear Technology*, vol. 100, no. 3, pp. 361–377, 1992.

[14] H. Eliasi, M. B. Menhaj, and H. Davilu, “Robust nonlinear model predictive control for nuclear power plants in load following operations with bounded xenon oscillations,” *Nuclear Engineering and Design*, vol. 241, no. 2, pp. 533–543, 2011.

[15] J. Hui and J. Yuan, “Load following control of a pressurized water reactor via finite-time super-twisting sliding mode and extended state observer techniques,” *Energy*, vol. 241, p. 122836, 2022.

[16] J. Wan, P. Wang, S. Wu, and F. Zhao, “Conventional controller design for the reactor power control system of the advanced small pressurized water reactor,” *Nuclear Technology*, vol. 198, no. 1, pp. 26–42, 2017.

[17] L. Duan, R. Petroski, L. Wood, and K. Caldeira, “Stylized least-cost analysis of flexible nuclear power in deeply decarbonized electricity systems considering wind and solar resources worldwide,” *Nature Energy*, vol. 7, no. 3, pp. 260–269, 2022.

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

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

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

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

[22] E. Grove, W. Gunther, and K. Sullivan, “Effect of component aging on PWR control rod drive systems,” Brookhaven National Laboratory, Upton, NY, Tech. Rep. BNL-NUREG-47470, 1992, aging Research Information Conference, Rockville, MD, March 1992. OSTI 10147987.

[23] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, “Guide to operational technology (OT) security,” National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

2.10 Exercises

2.1 ○ Concept check · 2 min

The output of the rod controller is a rod speed demand in percent per minute, and the drive integrates it into bank position. A proportional-only controller is in Auto when the turbine load is cut from 100 to 90 %. What happens to T_{avg} once the transient has died out?

(a) It settles off its program by an offset that shrinks as the gain K_c is raised

(b) It settles off its program until xenon builds and brings it back over several hours

(c) It returns to its program only if reset is added, because the controller has no integral

(d) It returns to its program, because the drive supplies the integrator the loop needs

2.2 ○ Concept check · 1 min

After a 10 % cut in turbine load, which statement describes the task of the rod control system?

(a) Return the average coolant temperature to its load program by moving the rods
(b) Bring reactor power down to the new turbine load by inserting the bank
(c) Hold the bank at its position so that the turbine load stays where it was set
(d) Insert the bank far enough to keep the reactor subcritical during the change

2.3 Concept check · 2 min

The shipped rod loop rings at $K_c = 6$ and hunts at $K_c = 8$ on the load cut of scenario 4. Which element of the loop sets that gain ceiling?

(a) The delayed neutron precursors, whose longest group lives about a minute
(b) The 3 s lag of the resistance temperature detector in its thermowell
(c) The maximum rod speed of 30 %/min, which the drive cannot exceed
(d) The xenon concentration, which changes the reactivity balance at reduced power

2.4 Theory and calculation · 5 min

The control bank is worth $W_{tot} = 1000\,\text{pcm}$ and follows the S curve of Equation (2.3). The coefficients are $\alpha_F = -3\,\text{pcm}/^\circ\text{C}$ and $\alpha_M = -4\,\text{pcm}/^\circ\text{C}$.

(a) Find the differential worth and the rod gain K_p at the design position, 60 % withdrawn. *Checkpoint:* 18.09 pcm per percent and 2.58 °C per percent.

(b) Find K_p at 80 and 90 % withdrawn, and the ratio of the gain at 60 % to the gain at 90 %. *Checkpoint:* 0.99 and 0.27 °C per percent, a ratio of 9.47.

(c) Find the controller gain at 80 % withdrawn that gives the same loop gain $K_c K_p$ as $K_c = 3$ at the design position. *Checkpoint:* 7.85.

2.5 Theory and calculation · 8 min

The plant of Table 2.2 runs at full power with $T_{c,0} = 310.0\,^\circ\text{C}$, and the turbine load is cut to 90 %. The rod loop returns the coolant to its program, and xenon and boron are held.

(a) Find the programmed T_{avg} at 90 % load and the change in coolant temperature from the reference. *Checkpoint:* 308.17 °C and −1.83 °C.

(b) Find the change in average fuel temperature from the reference. *Checkpoint:* −31.7 °C.

(c) Find the reactivity the two temperature terms hand back, and the bank insertion that absorbs it at the design differential worth of 18.09 pcm per percent. *Checkpoint:* About +102.5 pcm and 5.7 % of travel.

2.6 Theory and calculation · 5 min

The real rod drive has a deadband DB on the temperature error and a minimum speed $v_{min} = 3.5\,\%/\text{min}$. The loop cannot stop a move in less than about 0.13 min once it starts, and the rod gain is $K_p = 2.58\,^\circ\text{C}$ per percent.

(a) Find the coolant temperature change produced by one such move. *Checkpoint:* About 1.2 °C.

(b) The drive parks only if the full band $2\,DB$ is wider than that change. Find the smallest deadband that parks, and say what the 0.2 °C band and the 0.8 °C band each do. *Checkpoint:* About 0.6 °C. The 0.2 °C band hunts and the 0.8 °C band parks.

2.7 App challenge · 6 min

Open scenario 1, *The plant with nobody driving (MANUAL)*, as shipped, at the default seed. The bank stays parked in Manual, and the turbine load demand goes to 90 % at 2 min. Run to 30 min at speed 5× (6 s of wall time) and export the CSV.

(a) From the rows at 10 min, read the reactor power (`power_pct`) and the coolant error (`tavg_error_C`), each within 0.1. How many megawatts of thermal power did the plant shed with no rod moving? *Checkpoint:* 89.9 %, +13.9 °C, about 346 MW.

(b) From the row at 30 min, read the coolant error and the xenon in percent of equilibrium (the Xenon tile, or 100 times `xenon_rel`), each within 0.1. Explain why the error has fallen with no controller acting. *Checkpoint:* +10.8 °C and 101.1 % of equilibrium.

2.8 App challenge · 9 min

Open scenario 4, *A proper gain on a load change*, at the default seed. The turbine load demand goes to 90 % at 2 min, with the deadband and minimum speed off. For each run below, set the values before Run, run to 40 min at speed 5× (8 s), export the CSV and read the last row: the integrated error `tavg_iae` (the IAE tile), the peak error `tavg_max_error_C`, the rod travel `rod_travel_pct` (the Rod travel tile) and the bank `rod_pos_pct`.

(a) Run at Proportional gain 3, 5 and 8. Give the four numbers for each, within 0.1 in the unit read and within 1 % for rod travel. *Checkpoint:* IAE 5.7, 1.6 and 19.5 °C min. Peak 1.58, 0.98 and 1.26 °C. Travel 17.4, 8.7 and 153 %. Bank 56.5, 56.5 and 57.2 %.

(b) Reload the scenario, set T_{avg} *RTD time constant* to 0 in Show Model, Parameters, and Proportional gain to 12, and run again. Give the integrated error and the rod travel, and say what sets the gain ceiling of part (a). *Checkpoint:* About 1.1 °C min and 13 %.

2.9 App challenge · 10 min

Open the challenge, *Challenge, run the load-follow schedule*, at the default seed. It cuts the load to 80 % at 4 min, drops a rod at 15 min and returns to full load at 28 min. The target is a worst T_{avg} error within 2 °C, a bank above 30 % and under 60 % of rod travel over 60 min. Run each case to 60 min at speed 5× (12 s) and export the CSV. Read the worst error as the largest $|$`tavg_error_C`$|$, the lowest bank as the smallest `rod_pos_pct`, and the rod travel from the Rod travel tile, each within 0.1 in the unit read and within 1 % for travel.

(a) Run the case as shipped. Give the three numbers and the targets missed. *Checkpoint:* 8.10 °C, 49.0 % and 71 %: the error and travel targets are missed.

(b) Reload, and before Run set Control architecture to Cascade, Proportional gain 8, Integral time 3 min, Inner flux gain 4, Rod deadband 0 and Minimum rod speed 0. Give the three numbers. *Checkpoint:* 1.74 °C, 49.6 % and 52 %: a pass.

(c) Reload, and set + Mismatch, Proportional gain 5, Power mismatch gain 2, Rod deadband 0 and Minimum rod speed 0. Give the three numbers, and explain why the cascade has the smaller worst error. *Checkpoint:* 2.41 °C, 49.2 % and 41 %: the error target is missed.

2.10 Application area · 4 min

In scenario 11 the load is cut to 40 %. The cut asks the bank for about 615 pcm, while the bank holds 545 pcm between its starting position and the insertion limit at 30 % withdrawn. The differential boron worth is −8 pcm/ppm.

(a) Find the shortfall, and the boron change, with its sign, that would supply it so that the bank stays above the limit. *Checkpoint:* 70 pcm, a boration of about 9 ppm.

(b) The boron demand reaches the core through a 20 min mixing lag. Say when the operator must act, and why the rod controller cannot protect the limit on its own. *Checkpoint:* Before the cut, since the controller's objective holds no shutdown margin.

2.11 Application area · 3 min

Early in a fuel cycle, with boron concentration high, the moderator coefficient can be slightly positive. Scenario 12 sets it to $\alpha_M = +0.5\,\text{pcm}/^\circ\text{C}$ with α_F unchanged at −3 pcm/°C.

(a) By what factor does the rod gain K_p rise against the design coefficient of -4 pcm/°C, and what does the shipped tuning of $K_c = 3$ do? *Checkpoint:* A factor of 2.8. The loop swings about 12.6 °C peak to peak.

(b) Name two ways a plant engineer could keep the loop acceptable over a whole fuel cycle, and the measurement each one needs. *Checkpoint:* Gain scheduling, periodic retuning or a robust design, each needing the coefficient's range.

2.12 Review the agent · 5 min

An AI assistant reviews the shipped rod loop and writes: "The temperature controller is proportional-only, and a proportional-only loop always leaves an offset after a load change. Add integral action with $\tau_I = 2$ min. At $K_c = 3$ the loop has a gain margin of 3.00, so it has ample robustness to absorb the extra term." Find the flaw in the proposal, and say how you would check it in the chapter and in the app.