INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 7

Molten Salt Reactor Simulator Where the Pump Sets the Stability Margin

Pages 174–203

Edited by John D. Hedengren

How to cite this chapter. "Molten Salt Reactor Simulator Where the Pump Sets the Stability Margin," 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. 7, pp. 174–203.

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

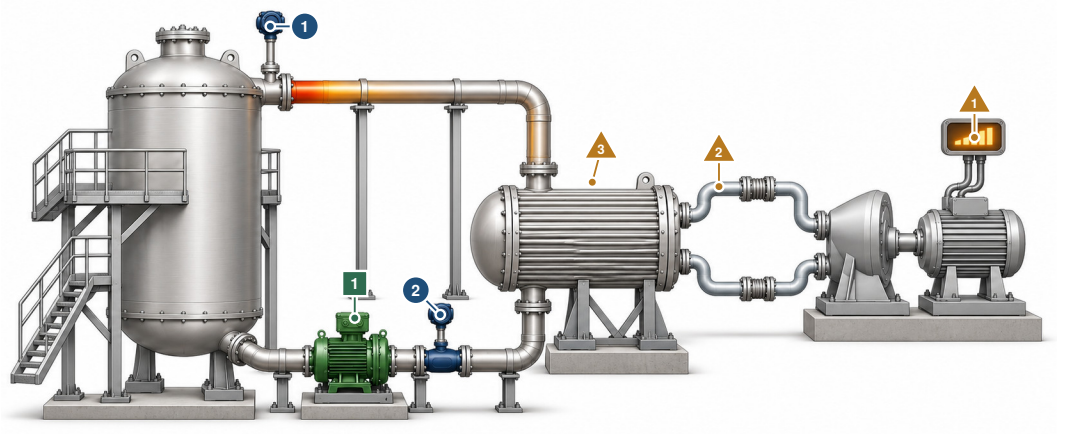
The app for this chapter
https://apmonitor.com/apps/molten-salt-reactor/

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

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

CHAPTER 7

Molten Salt Reactor Simulator Where the Pump Sets the Stability Margin



The controller moves 1 fuel salt pump speed

The controller reads 1 core outlet temperature 2 fuel salt flow, for the cascade

What disturbs it 1 turbine load 2 secondary salt flow 3 heat exchanger fouling

Notice that one pump carries both the heat and the atoms that release late neutrons round the loop, so its speed sets the outlet temperature and the stability margin together. Illustration, not to scale.

IN PLAIN TERMS A Pump That Also Sets the Safety Margin

The main idea. In this reactor the fuel is a hot liquid salt. A pump drives it through the core and out to a heat exchanger that passes the heat on. The pump's speed is how the plant holds the temperature of the salt leaving the core.

The challenge. A few neutrons, the tiny particles that keep the chain reaction going, are set free late, after the salt has left the core. A faster pump loses more of them, so it also shrinks the margin against a sudden jump in power.

Why it matters. Plain temperature control here is also a choice about safety margin.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	The first page, the history boxes of Sections 7.2 and 7.4, the fact box of Section 7.3, the try-it and fact boxes of Section 7.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 7.6, Section 7.8, Section 7.9 and the open challenge

The control question The fuel pump of a molten salt reactor carries delayed neutron precursors out of the core, so the actuator that holds the core outlet temperature also sets the effective delayed neutron fraction. Which actuator should hold that temperature, which loop structure does the plant call for, and how much stability margin does each tuning spend?

Learning objectives

- Evaluate the effective delayed neutron fraction at a given pump speed, and express a rod insertion in dollars.
- Select the manipulated variable from steady-state gains and step responses, and explain why the secondary flow fails as a pairing.
- Explain the offset a proportional loop leaves on a position actuator and the limit transport places on the loop gain.
- Compare feedback alone, static and lagged feedforward and a flow cascade on load ramps and delivery faults.
- Read temperature error and delayed neutron margin together when judging a tuning or a predictive controller.

Prerequisites Sections 1.4.2 and 1.4.3 (point kinetics, circulating precursors), 1.3 (loop action) and 1.6.2 (velocity form), and Chapter 2 for rod control and inherent feedback. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/molten-salt-reactor/

See also Volume 3, Chapter 2 for an inner loop against a disturbance in the actuator, Volume 3, Chapter 3 for how fast that inner loop must be and Volume 3, Chapter 6 for feedforward ahead of a long transport. Volume 5, Chapter 2 for control through a fuel transport delay.

7.1 Introduction

A control loop is ordinarily analyzed as an arrangement in which the manipulated variable moves the operating point and leaves the transfer function between them unchanged [1]. Gain scheduling is introduced as the remedy for the cases where that assumption fails, and the standard examples are a nonlinear valve characteristic or an aircraft flying at a new dynamic pressure [2]. A liquid-fueled reactor is a stronger instance of the same effect, because the actuator does not only move the process gain. It moves the quantity that sets how fast the plant can respond to any reactivity disturbance at all.

The mechanism is the fuel. Fissioning uranium dissolved in a fluoride salt releases about 99.35 % of its neutrons within microseconds and the remaining 0.65 % seconds to minutes later, from the beta decay of fission fragments called delayed neutron precursors [3]. That small late fraction is the entire reason a reactor can be controlled by mechanical equipment rather than by microsecond electronics. In a solid-fuel core the precursor decays where it was born. In a molten salt reactor the precursor is dissolved in the fuel salt, the pump carries it out of the core, and it emits its neutron in a pipe or in a heat exchanger where there is no fuel geometry to multiply it. The effective delayed neutron fraction of the plant of this chapter is 419 pcm against the 650 pcm the fuel itself produces, so 35.5 % of the delayed neutrons are discarded by the act of pumping, and the control rod carries 231 pcm of permanent excess reactivity for no purpose other than paying for that loss. Slowing the pump reduces the loss. At the natural circulation floor of 6 % of design flow the effective fraction returns to 604 pcm. Section 1.4.3 sets out the circulating precursor balance behind these numbers.

Reactivity expressed in dollars is not a fixed currency on this plant. Reactivity divided by the delayed fraction is the standard measure of how close an insertion is to prompt criticality, and in a solid-fuel plant that denominator cannot move. Here the denominator is set by the fuel pump. The same 100 pcm rod is worth 0.247 dollars at 120 % pump speed and 0.203 dollars at 40 %, with the rod and the core unchanged. A block diagram of the loop does not display that dependence, because the block and the input to the block share a variable.

The class of problem generalizes in two directions. The first is any loop whose manipulated variable changes the dynamics rather than only the operating point, such as a recycle valve that sets residence time in the reactor it recycles to, a feed pump that sets the conversion of the reaction it feeds, or an antisurge valve that moves the compressor along its own map. One fixed tuning on such a plant is a claim about an operating point, and the theory of scheduling controllers across an envelope exists because that claim is usually false [1, 2]. The second is any process whose transport time is comparable to its reaction time, which covers long pipelines, blend headers, conveyor-fed dryers and every recycle loop acting on a measurement that is one circulation out of date. This plant is the extreme member of that family, because the same 16 s of pipe is simultaneously the dead time in the temperature loop and the mechanism that strips a third of the delayed neutrons.

The experiments of this chapter exercise a specific set of process dynamics and control principles: pairing selection from measured steady-state gains and step responses, transport delay as the limit on achievable closed-loop bandwidth, the structural reason a position actuator leaves proportional offset, feedforward with dynamic compensation, cascade control against a disturbance inside the inner loop, gain scheduling where the process gain follows the manipulated variable, and derivative action on a noisy measurement [1, 4]. The physical fundamentals behind the model are conservation of energy on two core nodes and a heat exchanger node, a neutron and precursor population balance closed by the criticality condition, an effectiveness relation that bounds the duty a finite secondary stream can carry, and a separation of time scales that spans two decades between the fuel salt node and the graphite. The behavior matters industrially because liquid-fueled designs are being engineered and taken through regulatory review again after a gap of four decades, and circulating precursor loss is the first item on the neutronics side of each of them [5, 6].

The subject is normally first met analytically, with no means of observing a plant property change under an operator input, of applying a pump trip and reading the reactivity that follows, or of measuring that the apparent pairing between the heat sink and the core outlet has the wrong sign. This chapter works with a browser simulation of a circulating-fuel reactor, set out completely enough to be rebuilt from the text, and reports what happens when the fuel pump is asked to hold a temperature. Section 7.2 traces the plant from the Aircraft Reactor Experiment to the present designs. Section 7.3 gives the balances and the constants that reproduce the simulator: twelve states, a delay line on the return leg and an annunciator that declares when its own approximation has expired. Section 7.4 measures the three candidate manipulated variables on one plant and finds that the heat sink pairing carries a sign inversion and 16 s of pure dead time while the fuel pump carries neither. Section 7.5 describes the implemented loop, with load feedforward, an inner flow cascade and a predictive option. Section 7.6 runs the scenario sequence, which reports the stability margin beside the controlled variable so that a tuning decision and a safety margin are read together. Sections 7.7 to 7.10 discuss the passive shutdown, name the hardware a real loop of this kind runs on, state the limitations and summarize.

7.2 Background

Liquid fuel arrived as an answer to an aviation problem, in the Aircraft Reactor Experiment at Oak Ridge in 1954 [7], whose story the History box below tells.

> ⧖ HISTORY **A Reactor That Followed Its Load (1954)**
>
> The Aircraft Reactor Experiment was built at Oak Ridge for work on nuclear propulsion of aircraft. Its fuel, a molten mixture of the fluorides of sodium, zirconium and uranium, was pumped round a closed loop through a core moderated by beryllium oxide and on to a heat exchanger. Uranium went into the fuel on October 30, 1954, and the reactor was shut down for the last time on November 12. It ran at up to 2.5 megawatts. Its strongly negative temperature coefficient made it, in the words of its operating report, a slave to the load: the power followed the rate at which heat was taken out. No more than 5 percent of the xenon stayed in the circulating fuel [8].

Chapter 8 follows that aviation line with a notional flight reactor. The behavior this chapter is about was present from the beginning, because a circulating core cannot be analyzed with the delayed neutron data of the fuel alone. The program that followed was aimed at civil power. The Molten Salt Reactor Experiment was a graphite-moderated loop with the fuel dissolved in a lithium and beryllium fluoride salt, and its design report fixed the loop geometry, the transit times, the drain tanks and the frozen salt plug that the simulator of this chapter still models [9]. It reached criticality in 1965 and operated until 1969 [10]. It remains the only reactor of its kind that has operated, which bounds what any circulating-fuel model can be validated against.

The pump experiments at that reactor are often paraphrased loosely, so what they measured is stated here precisely. The reactivity difference between fuel circulating and fuel stationary was obtained at zero power, where no temperature feedback exists to mask it, by starting and stopping the fuel pump and observing the compensating rod motion and the reactor period. The theory of the period during circulation was worked out and then checked against the measurement on the operating loop [11], and the wider zero-power program established the rod worths, the temperature coefficients and the circulating delayed neutron fraction the later dynamics work depended on [12]. What those experiments produced is a reactivity, not a delayed fraction, and the delayed fraction follows from it only through a model of the loop. The reactivity was substantial against a total delayed fraction that is itself only a fraction of a percent, and it had to be held in the shim rods for as long as the pump was running. The reactor was also the subject of a full dynamics program, in which a lumped model of the core, the loop and the heat exchanger was built and then validated against pulse and frequency response tests performed on the operating plant [13, 14]. The model of this chapter follows the structure of that lumped model rather than inventing one.

Two other results from that program shape what this simulator does and does not contain. Xenon-135 is the dominant fission product poison in a solid-fuel reactor and the reason a shut-down plant cannot always be restarted, but a molten salt reactor sparges noble gases continuously out of the fuel salt in the pump bowl, and the xenon behavior of the operating reactor was measured with that stripping in place [15]. Xenon is accordingly absent from the model here and the omission is declared in the assumptions rather than left silent. The second result concerns operator training. Simulators for the plant were written and used at the laboratory in 1966 on analog computing equipment, for the same reason an interactive model is written today, which is that no prior operating experience existed for a reactor whose delayed neutron fraction changes when the pump starts [16].

Work on the subject then stopped. The molten salt program was ended in the early 1970s in favor of the liquid metal fast breeder, and for roughly forty years the literature consisted of a small number of design studies revisiting the original reports. Commercial and national interest returned in the 2000s, when the molten salt reactor was named one of the six systems of the Generation IV program and a coordinated modeling and materials effort restarted around it [5]. The designs now under development range from graphite-moderated thermal converters close in layout to the original loop to fast-spectrum chloride and fluoride machines with no moderator at all [6]. Each of them has to establish what the pump does to the delayed neutrons and what follows when it stops. Control work specific to these plants is correspondingly recent, and the core outlet temperature is the variable it takes as the controlled one [17], for reasons the steady-state analysis in Section 7.4 makes concrete. Calculation of the effective delayed fraction with circulation, including the precursor distribution the point model here does not resolve, is still an active subject [18].

7.3 Reactor Model

The plant is a graphite-moderated liquid-fueled reactor rated at 250 MW thermal. Fuel salt is pumped upward through channels cut in a graphite stack, leaves the core at 665 °C, gives its heat to a secondary salt stream in one heat exchanger, and returns to the core inlet at 565 °C. That 100 °C core temperature rise is the quantity the control system manipulates. The mean fuel salt temperature is 615 °C, the graphite sits 25 °C above it at 640 °C, and a single control rod is 55 % withdrawn out of a total worth

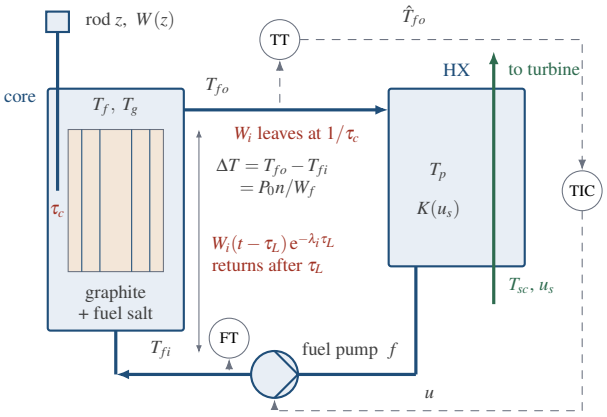


Figure 7.1 The fuel salt loop with the core, the hot leg, the heat exchanger and the fuel pump. Precursor group W_i leaves the core at a rate $1/\tau_c$, decays along the external loop by $e^{-\lambda_i \tau_L}$ and re-enters one loop transit later, and the same salt carries the heat, so the cold leg temperature T_{fi} is the heat exchanger outlet delayed by the same τ_L. TT measures the core outlet, TIC drives the fuel pump command u, and FT closes the optional inner flow loop around the pump.

Table 7.1 Classification of Variables. Nominal Values Are Those of the Design Point.

Symbol	Role	Nominal	Range
T_{fo} (TT)	CV, measured	665 °C	640–690 °C SP
u (fuel pump)	MV	100 %	20–120 %
L turbine load	DV, measured	100 %	30–110 %
$u_{s,\mathrm{cmd}}$	DV, operator	100 %	10–150 %
Δz rod trim	DV, operator	0 %	±10 %
g pump delivery	DV, unmeasured	1.00×	0.60–1.15×
b TT bias	DV, unmeasured	0 °C	±5 °C
$W_1 \ldots W_6$	states	see Table 7.2	≥ 0
T_f	state	615.0 °C	—
T_g	state	640.0 °C	—
T_p	state	615.0 °C	—
z	state	55.0 %	0–100 %
f	state	100 %	0–138 %
u_s	state	100 %	0–150 %

of 800 pcm. Figure 7.1 carries the layout and every symbol that appears in the equations. Time is in seconds, reactivity is carried in $\Delta k/k$ and displayed in pcm where 1 pcm is 10^{-5}, and both flows are in percent of their design values.

Table 7.1 classifies what the loop interacts with. Twelve states are integrated. Five annunciations bound normal operation, an outlet deviation alarm at 5 °C, a high flux alarm at 130 % of rated power, a core inlet alarm at 500 °C where the salt approaches freezing, a prompt critical declaration at $\rho \geq \beta$, and the freeze plug at a core outlet of 750 °C. Only the last of the five changes the model, and the other four annunciate while the run continues so that the transient that raised them stays readable.

7.3.1 Prompt Jump and Circulating Precursor Kinetics

The neutron population and the six precursor groups are written in point kinetics, and the prompt mode is removed algebraically because its time constant at criticality is a few milliseconds while the graphite in the same model has a time constant of 277 s. Setting the prompt derivative to zero

and scaling each precursor concentration by the neutron generation time, $W_i = \Lambda c_i$, gives the neutron power of Equation (7.1).

$$n = \frac{\sum_{i=1}^{6} \lambda_i W_i}{\beta - \rho} \tag{7.1}$$

where n is the nuclear power as a fraction of rated, λ_i is the decay constant of group i in 1/s, ρ is the total reactivity in $\Delta k/k$ and β is the *static* delayed fraction of 650.2 pcm. The generation time cancels out of the model entirely [3], as in the one-group form of Equation (1.1). The fraction in the denominator of (7.1) is the static one and not the effective one, because it arrives from the production term and production does not care where the precursor ends up. What circulation changes is the size of the numerator that a given power can sustain.

The precursor balance is the ordinary six-group equation with two extra terms, as shown in Equation (7.2).

$$\frac{\mathrm{d}W_i}{\mathrm{d}t} = \beta_i n - \lambda_i W_i - \frac{W_i}{\tau_c} + \frac{W_i(t-\tau_L)\,\mathrm{e}^{-\lambda_i \tau_L}}{\tau_c} \tag{7.2}$$

where β_i is the static fraction of group i, τ_c is the time the salt spends in the core and τ_L is the time it spends outside. Scaling by Λ leaves W_i in seconds, so every term in (7.2) is dimensionless. The third term removes precursors at the rate the pump sweeps the core volume, and the fourth returns whatever survives one trip round the external loop. Both transit times scale inversely with the delivered pump speed f in percent of design, as shown in Equation (7.3).

$$\tau_c = \frac{\tau_{c,0}}{\max(f_{\mathrm{nat}}, f)/100}, \qquad \tau_L = \frac{\tau_{L,0}}{\max(f_{\mathrm{nat}}, f)/100} \tag{7.3}$$

where f_{nat} is a thermosiphon floor of 6 %, present because a hot core above a cold heat exchanger does not stop circulating when the pump stops and because the transit times would otherwise be unbounded.

Pause. Try Exercise 7.2 before reading on.

Setting the derivative in (7.2) to zero and replacing $W_i(t-\tau_L)$ by W_i gives the steady result that the whole application is organized around, as shown in Equation (7.4).

$$\beta_{\mathrm{eff}} = \sum_{i=1}^{6} \beta_i \frac{\lambda_i \tau_c}{\lambda_i \tau_c + 1 - \mathrm{e}^{-\lambda_i \tau_L}} \tag{7.4}$$

Equation (7.4) is Equation (1.5) of Chapter 1. It says that a group keeps its worth in proportion to how much of its life it spends inside the core. Table 7.2 evaluates it group by group at the design flow, where the salt spends 9 s in the core and 16 s outside it.

The retained column of Table 7.2 carries the mechanism. Group 6 has a mean life of 0.33 s and keeps 96.4 % of its worth, because a precursor that decays in a third of a second never reaches the outlet nozzle of a core it takes 9 s to cross. Group 1 has a half-life of 55.9 s and keeps 38.3 %, because almost all of its life is spent somewhere else. The largest group of the six is worth 256.8 pcm in the fuel and contributes 188.0 pcm here. Adding the column gives 419.4 pcm against a static 650.2 pcm, so 35.5 % of the delayed neutrons decay outside the core. A circulating core is therefore critical at *positive* reactivity, and the excess it has to hold is $\beta - \beta_{\mathrm{eff}}$, which at the design flow is 230.8 pcm.

DID YOU KNOW? Twenty-Five Seconds Round the Loop

At its design flow of 1,200 gallons per minute, the fuel salt of the Molten Salt Reactor Experiment took 25.2 seconds to go once round its loop, 9.4 of them between the top and the bottom of the graphite core [9]. The plant of this chapter spends 9 s in the core and 16 s outside it.

Table 7.2 Six-Group Delayed Neutron Data for Thermal Fission of Uranium-235 and the Worth Each Group Retains at the Design Pump Speed, $\tau_c = 9\,\text{s}$ and $\tau_L = 16\,\text{s}$.

Group	λ_i (1/s)	Half-life (s)	β_i (pcm)	Retained	Contribution
1	0.0124	55.9	21.5	38.3 %	8.2 pcm
2	0.0305	22.7	142.4	41.6 %	59.2 pcm
3	0.111	6.24	127.4	54.6 %	69.6 pcm
4	0.301	2.30	256.8	73.2 %	188.0 pcm
5	1.14	0.61	74.8	91.1 %	68.2 pcm
6	3.01	0.23	27.3	96.4 %	26.3 pcm
Total			650.2	64.5 %	419.4 pcm

7.3.2 Reactivity and the Criticality Condition

Reactivity is measured from the design reference condition and carries a rod term, two temperature terms and whatever a fault has inserted, as shown in Equation (7.5).

$$\begin{aligned}\rho &= \rho_0 + W(z) - W(z_0) + \rho_x \\ &\quad + \alpha_f (T_f - T_{f,0}) + \alpha_g (T_g - T_{g,0})\end{aligned} \tag{7.5}$$

$$W(z) = W_{\text{tot}} \left[\frac{z}{100} - \frac{\sin(2\pi z/100)}{2\pi} \right] \tag{7.6}$$

where $\rho_0 = \beta - \beta_{\text{eff}}$ evaluated at the design flow is the 230.8 pcm the design point holds, $W(z)$ is the integral rod worth in $\Delta k/k$, α_f and α_g are the fuel salt and graphite temperature coefficients, $z_0 = 55\,\%$, $T_{f,0} = 615\,^\circ\text{C}$, $T_{g,0} = 640\,^\circ\text{C}$ and ρ_x carries transient insertions such as a dropped rod. Equation (7.6) is the usual S curve with the differential worth peaking at half insertion, which at the design position is 15.6 pcm per percent withdrawn.

Every steady state of the plant satisfies $\rho = \beta - \beta_{\text{eff}}$, because a settled reactor is critical and a circulating one has to make up what it is pumping away. That criticality condition governs every steady operating point of the plant. With the rod fixed and the pump fixed, the two temperature terms in (7.5) must sum back to the same number they started at whatever the power, so the weighted core temperature is pinned and the *mean* fuel salt temperature barely moves. Taking the turbine from 100 to 70 % with the pump parked moves the mean fuel salt by 0.8 °C in ten minutes while the core outlet falls 10.0 °C and the core inlet rises 11.6 °C. The core temperature spread is what moved.

7.3.3 Thermal Balances of the Core, Loop and Heat Exchanger

The core is two well-mixed nodes, the fuel salt and the graphite, and the heat exchanger primary side is a third. Each node carries a linear temperature profile, so its mean is the average of its inlet and outlet and the outlet follows algebraically. The thermal balances are Equations (7.7) to (7.9).

$$C_f \frac{\mathrm{d}T_f}{\mathrm{d}t} = (1 - f_G) P_0 n - 2W_f (T_f - T_{fi}) - h_{fg} (T_f - T_g) \tag{7.7}$$

$$C_g \frac{\mathrm{d}T_g}{\mathrm{d}t} = f_G P_0 n + h_{fg} (T_f - T_g) \tag{7.8}$$

$$C_{hx} \frac{\mathrm{d}T_p}{\mathrm{d}t} = 2W_f (T_{fo} - T_p) - K(u_s)(T_p - T_{sc}) \tag{7.9}$$

where C_f, C_g and C_{hx} are heat capacities in MW s/°C, P_0 is the rated thermal power, f_G is the share of fission energy deposited in the graphite as gammas and fast neutrons, h_{fg} is the salt-to-graphite conductance, and $W_f = C_f/\tau_c$ is the flow capacity rate in MW/°C. Every term is a megawatt. The salt is the fuel, so 93.5 % of the fission energy appears in the salt directly and only 6.5 % has to be conducted back out of the moderator.

The three loop temperatures follow from the linear profiles and from the transport around the external loop, as shown in Equation (7.10).

$$\begin{aligned} T_{fo} &= 2T_f - T_{fi}, \qquad T_{p,\mathrm{out}} = 2T_p - T_{fo} \\ T_{fi}(t) &= T_{p,\mathrm{out}}(t - \tau_L) \end{aligned} \tag{7.10}$$

The delay in (7.10) is the same τ_L that appears in (7.2), because the same salt carries the heat and the precursors through the same pipe. The dead time of the temperature loop and the precursor loss of the neutronics problem are therefore one number, which is the central coupling of this plant.

The heat exchanger passes its duty to a secondary salt stream of finite capacity rate, and the effectiveness result for such a stream exchanging with a well-mixed source is Equation (7.11).

$$K(u_s) = C_s \left[1 - \mathrm{e}^{-UA/C_s} \right], \qquad C_s = C_{s,0}\,\frac{u_s}{100} \tag{7.11}$$

where UA is the exchanger conductance and $C_{s,0}$ is the secondary capacity rate at design flow. Equation (7.11) is monotone in u_s and saturates at UA, so the effective conductance is 1.852 MW/°C at design flow and 2.193 MW/°C at 150 %, and no amount of secondary pumping passes 3.20 MW/°C. That ceiling belongs to the tubes rather than to the pump.

The turbine reaches the reactor only through the secondary cold leg, as shown in Equation (7.12).

$$T_{sc} = T_{sc,0} + \Delta T_{sc}\,(1 - L) \tag{7.12}$$

where L is the turbine load as a fraction of rated after a ramp limiter of 1 %/s. Closing the throttle raises the steam pressure, raises the saturation temperature on the steam side, and returns the secondary salt to the exchanger hotter, so the cold leg is 480 °C at full load, 510 °C at 70 % and 550 °C at 30 %. A hotter cold leg is less driving temperature difference and therefore less duty, so reactor power falls with nothing commanded.

7.3.4 Actuators and Instruments

The fuel pump is a position actuator with a first-order lag, the secondary pump is the same with a longer time constant, and the rod is a position servo with a hard speed limit, as shown in Equations (7.13) and (7.14).

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{g\,u - f}{\tau_p}, \qquad \frac{\mathrm{d}u_s}{\mathrm{d}t} = \frac{u_{s,\mathrm{cmd}} - u_s}{\tau_{\mathrm{sec}}} \tag{7.13}$$

$$\frac{\mathrm{d}z}{\mathrm{d}t} = \mathrm{sat}\left(\frac{z_{\mathrm{sp}} - z}{\tau_z},\ \pm v_{\max} \right) \tag{7.14}$$

where g is the delivery the pump achieves per unit of command, $v_{\max} = 0.5\,\%/\mathrm{s}$ and $\tau_z = 1\,\mathrm{s}$ is the approach time constant that lets the rod settle onto its demand instead of chattering across it. The rod moves at $v_{\max}$ whenever the position error exceeds 0.5 %. On a pump trip the target in (7.13) becomes zero and τ_p becomes the flywheel coastdown time of 10 s, which is exactly how fast the effective delayed fraction climbs back toward its static value. Nothing downstream of the controller integrates, and Section 7.5 shows what that costs a proportional loop.

Two instruments are read. The core outlet thermocouple adds a bias and a noise sequence and then filters the result with a time constant τ_{TC}, and the flow transmitter does the same with τ_{FT}, as shown in Equation (7.15).

$$\tau_{TC}\frac{\mathrm{d}\hat{T}_{fo}}{\mathrm{d}t} = T_{fo} + b + n_T(t) - \hat{T}_{fo} \tag{7.15}$$

where b is the thermocouple bias and n_T is a Gaussian sequence passed through a first-order filter and rescaled so that its stationary standard deviation matches the specified noise level. The filter is advanced by the exact zero-order-hold update, which is stable at any sample period.

Table 7.3 Model Parameters. Every Constant Is Reported and Operator Adjustable.

Symbol	Meaning	Value	Unit
P_0	rated thermal power	250	MW
f_G	fission energy into graphite	0.065	—
C_f	core fuel salt heat capacity	22.5	MW s/°C
C_g	graphite heat capacity	180	MW s/°C
h_{fg}	salt-to-graphite conductance	0.65	MW/°C
$\tau_{c,0}$	core transit at design flow	9.0	s
$\tau_{L,0}$	loop transit at design flow	16.0	s
f_{nat}	natural circulation floor	6.0	%
τ_p	fuel pump response	4.0	s
—	fuel pump coastdown	10.0	s
α_f	fuel salt coefficient	−5.0	pcm/°C
α_g	graphite coefficient	−1.0	pcm/°C
W_{tot}	total control rod worth	800	pcm
$v_{\max}$	maximum rod speed	0.5	%/s
UA	heat exchanger conductance	3.20	MW/°C
$C_{s,0}$	secondary capacity rate	2.632	MW/°C
C_{hx}	exchanger primary capacity	25	MW s/°C
τ_{sec}	secondary pump response	6.0	s
$T_{sc,0}$	cold leg at full load	480	°C
ΔT_{sc}	cold leg rise at zero load	100	°C
—	turbine ramp rate	1.0	%/s
τ_{TC}	outlet thermocouple filter	2.0	s
τ_{FT}	flow transmitter filter	0.5	s
—	outlet deviation alarm	5.0	°C
—	cold leg low alarm	500	°C
—	freeze plug drain temperature	750	°C
—	high flux alarm	130	%

7.3.5 Parameters, Initialization and Numerics

Table 7.3 lists the constants. All of them are reported and adjustable, because no quantity is withheld for identification. What is not reported is the consequence, because $\tau_{L,0} = 16\,\mathrm{s}$ is given without stating that the same number sets the gain limit of the temperature loop and the value of Equation (7.4).

No scenario begins anywhere but on an exact equilibrium. The design point is solved in closed form rather than by relaxation, because the core rise and the duty are both linear in n. Putting the core outlet on the setpoint gives the power directly from $K(T_{fo} - T_{sc}) = P_0 n \left[1 + K/(2W_f)\right]$, after which T_f, T_g and the rod position that makes that combination critical follow. Below about 25 % pump speed at full load that position lies past the bottom stop, so the plant starts at the steady state with the rod on the stop, an outlet of 700.2 °C at 20 %. The precursors are initialized from the steady form of (7.2), $W_i = \beta_i n / [\lambda_i + (1 - \mathrm{e}^{-\lambda_i \tau_L})/\tau_c]$. The delay line is seeded with the equilibrium heat exchanger outlet temperature so that the first loop transit of every run reads a settled plant rather than a fabricated inlet temperature.

States advance on a fixed 0.02 s fourth-order Runge-Kutta step, which is sized against the 0.33 s mean life of the fastest precursor group. The controller executes every 0.2 s and the historian logs five samples per simulated second. The delay line holds 8192 samples pushed every third step, or every 0.06 s, and interpolated linearly, which is 491 s of history, enough to cover the natural circulation floor where τ_L reaches 267 s. The noise seed is 20260829. The model spans two decades of thermal time constant, from 4.5 s on the fuel salt node and 3.65 s on the heat exchanger primary to 277 s on the graphite.

A step follows the order of Section 1.6.1: due scenario events, the controller when a full 0.2 s

period has elapsed (at $t = 0$ and every period after that), the integration, the post-step corrections and the log. The post-step corrections clip the pump and the rod against their stops and push the seven delayed quantities into the history buffer. The history is pushed from that post-step hook rather than from inside the derivative function, because a fourth-order integrator evaluates the derivative four times per step at intermediate states.

7.3.6 Assumptions

Each simplification below is stated with the place it fails.

Point kinetics with the prompt jump. Equation (7.1) removes a millisecond mode so that a 277 s graphite transient can be integrated in the same run, and it is invalid at and above prompt critical. The application annunciates $\rho \geq \beta$ as a modeling failure rather than integrating through it.

The precursors are uniformly mixed in the core. A point model carries no flux shape, so a precursor about to leave through the outlet nozzle and one just born at the inlet are treated as equally valuable. Resolving the distribution against a real velocity profile changes the computed effective fraction [18], and nothing here is a design value.

Transit times change instantly with pump speed. Salt already in the loop is treated in (7.3) as having always traveled at the new speed, so the returning term in (7.2) is approximate during the first loop transit after a large pump change. That is the situation during a trip coastdown, which is why the trip scenario is a demonstration of a mechanism rather than a licensing calculation.

Reactivity coefficients are constants and there is no burnup. The salt coefficient moves with composition as fission products accumulate and the graphite coefficient moves with fast neutron damage over a few years, and the application exposes α_g as an adjustable parameter instead of modeling that drift.

Xenon-135 is not modeled at all. A molten salt reactor sparges the noble gases out of the fuel salt continuously in the pump bowl, so the equilibrium xenon worth is a small fraction of the 2600 pcm a solid-fuel core carries and it does not set an operating restriction [15]. That is an advantage of liquid fuel and also a claim about a stripping efficiency a real plant would have to measure. Nothing here should be used to reason about a xenon transient.

The whole external transport sits on the return leg. A real loop divides the transit into a hot leg, an exchanger and a cold leg, which would place part of the delay ahead of the exchanger rather than all of it behind. The assumption makes the dead time measured in Section 7.4 a clean 16 s where a distributed model would smear it.

The secondary loop is quasi-static. Equation (7.12) represents the entire steam plant, so there is no steam generator, no drum level, no feedwater loop and no tertiary salt loop.

Natural circulation is one number. The 6 % floor in (7.3) is not a buoyancy calculation and does not depend on loop height, temperature difference or power. It keeps the transit times finite and it is not a natural circulation prediction.

The disturbances are idealized. A rod that drops 100 pcm in a single step at an instant of the operator's choosing is not a fault a plant delivers on request, but it isolates one measurement, the same reactivity valued against two different delayed fractions. The fuel pump trip is run at 9 % power because Section 7.6 shows the same mechanism is not observable at full power behind a faster one.

The freeze plug is a single threshold. Above 750 °C at the core outlet the fission term is set to zero and the model declares itself invalid, whereas a real drain takes minutes and carries decay heat into the drain tank.

7.4 Selection of the Manipulated Variable

A reactor has a control rod and a heat exchanger has a secondary flow, and neither is used here to hold the core outlet temperature. That choice follows from steady-state and step measurements on the model, and the measurements are reported because two of the three candidates fail for reasons that generalize to other processes.

Table 7.4 The Three Candidate Manipulated Variables Measured at the Design Point, 100 % Load and a 665 °C Core Outlet.

	Fuel pump	Secondary flow	Rod trim
Step applied	−10 %	+25 %	+1 %
ΔT_{fo} steady (°C)	+7.57	+4.65	+3.44
Gain (°C per %)	−0.757	+0.186	+3.44
Apparent dead time (s)	0	16	0
ΔT_{fo} at 1 s (°C)	0.231	0.000	—
ΔT_{fo} at 8 s (°C)	4.24	0.000	—
Power at steady state (%)	101.1	110.2	101.9
Actuator rate limit	none	none	0.5 %/s

HISTORY The Experiment That Held Its Outlet With a Rod (1965)

The design report of the Molten Salt Reactor Experiment set out how its power would be run. Below one megawatt, in the starting mode, the nuclear power was held constant. Above it, in the run mode, the temperature of the fuel salt leaving the reactor was held constant. One of the three control rods was on a servo and the other two on manual switches. Load was set at the radiator that passed the heat of the coolant salt to air: a single spring-return switch opened the radiator doors, then closed a bypass damper, then started a second blower. The fuel pump was specified at about 1,160 revolutions per minute to deliver 1,200 gallons per minute [9]. The plant of this chapter holds the same outlet temperature with the pump instead of the rod.

7.4.1 Steady-State and Step Response of Three Candidates

Table 7.4 reports each candidate at the design condition. The steady-state gains come from the closed-form steady state of the model and the dynamic figures come from open-loop steps run on the engine itself, each timed from the controller scan at which the step reaches the plant.

The **control rod** has the largest gain per unit of travel and the slowest response. Criticality pins the weighted core temperature, so the rod is the only actuator that can shift the mean at all. One percent of rod travel at the design position is worth 15.5 pcm of integral worth, which the two temperature coefficients absorb as 2.51 °C on the mean fuel salt and 2.98 °C on the graphite, and the outlet moves 3.44 °C because the power also rises to 101.9 % and widens the core rise. The rod travels at 0.5 %/s, so the full 3.44 °C takes two seconds of continuous motion to command, and the rod is a coarse actuator for setting an operating point rather than for rejecting a disturbance. It is provided as an operator trim.

The **secondary salt flow** is the apparent candidate and it fails on both the sign and the dynamics. Opening it takes more heat out of the primary, so the reactor makes more power, so the core temperature rise widens, so the outlet gets *hotter* rather than colder. Taking the secondary flow from 100 to 125 % moves the steady outlet from 665.00 to 669.65 °C while the power moves from 100.0 to 110.2 %, a gain of +0.186 °C per percent against a controller that would need it negative. The dynamic behavior is worse than the sign. Stepping the secondary flow leaves the core outlet at exactly its starting value for 16 s, because the only path from the heat exchanger to the core outlet runs through the cold leg and the cold leg is a full loop transit away. There is no inverse response anywhere in that record. The outlet holds at 0.000 °C of deviation through 16 s, then moves 0.94 °C by 20 s, 1.45 °C by 25 s and 2.17 °C by 30 s, rising monotonically to its steady value. Pure dead time followed by a monotone rise is a different defect from inverse response and it calls for a different remedy, so the two are distinguished here rather than grouped together. A controller on this pairing would also hold full reactor power against every load change, which defeats the purpose of load following.

Pause. Try Exercise 7.3 before reading on.

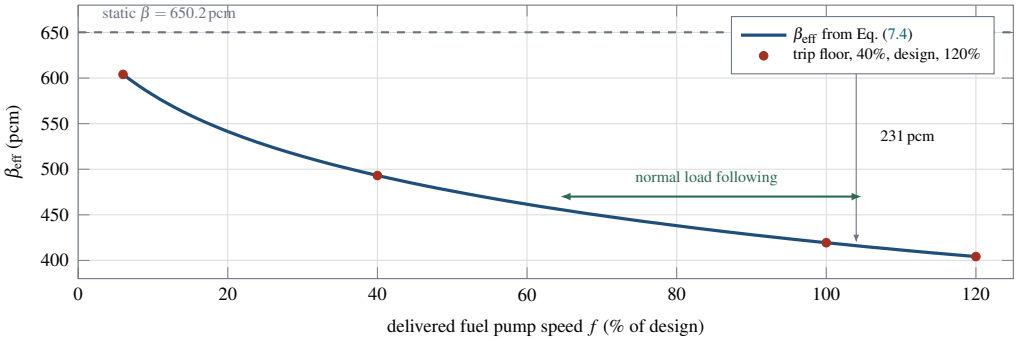


Figure 7.2 The effective delayed neutron fraction as a function of delivered fuel pump speed. The curve is Equation (7.4) evaluated with $\tau_c = 9/f$ and $\tau_L = 16/f$ in seconds. Holding the core outlet at 665 °C while the turbine moves from 110 down to 30 % load runs the pump between 104.8 and 64.4 %, so ordinary load following moves the stability margin between 416 and 456 pcm with no operator action.

The **fuel pump** carries neither defect. The core temperature rise is $P_0 n/W_f$ and $W_f = C_f/\tau_c$ falls with the pump, so slowing the pump widens the rise and raises the outlet directly. Because $T_{fo} = 2T_f - T_{fi}$ and the flow capacity rate appears in (7.7) at once, the response begins immediately. A step from 100 to 90 % moves the outlet 0.231 °C in the first second, 4.24 °C by 8 s and 7.57 °C at steady state, which is a gain of −0.757 °C per percent of pump measured across that step. The relation is monotone across the whole actuator range. It is not, however, linear. The local slope of the steady-state relation is −0.686 °C per percent at 100 % pump speed and −1.180 °C per percent at 75 %, so the process gain is 72 % larger at the pump speed a partially unloaded plant runs at. Any single tuning is therefore a statement about a pump speed, which is the standard condition under which gain scheduling is applied [1, 2].

7.4.2 Effect of Pump Speed on the Stability Margin

The cost of the pairing is that the manipulated variable is also the argument of Equation (7.4). Figure 7.2 plots the effective delayed fraction across the whole actuator range.

Two features of Figure 7.2 matter for control. The first is the vertical gap of 231 pcm at the design point, which is reactivity the rod holds permanently and which the plant hands back the moment circulation stops. The second is the marked span, which is the part of the curve an ordinary load schedule visits. Holding the outlet on setpoint from 110 down to 30 % turbine load takes the pump from 104.8 to 64.4 % and moves β_{eff} from 416 to 456 pcm, about 10 % of the margin the plant has against every reactivity disturbance. The three classical structures do not constrain that number, and the tuning that spends the least integrated error will spend the most margin, because the fastest way to hold a temperature here is to run the pump hard. The effective delayed fraction and the percentage of precursors lost are therefore reported alongside the outlet error, because a trade that is not measured is still made.

🕮 **Pause.** Try Exercise 7.1 before reading on.

7.5 Controller Design

7.5.1 Velocity Form PID Algorithm

The temperature controller TIC is the velocity-form PID of Section 1.6.2, Equations (1.13) to (1.16) [19], with the measurement $\hat{T}_{fo}$ as PV, the pump command u as output and the control period $\Delta t_c = 0.2$ s, independent of the 0.02 s physics step. The loop is direct acting, $a = -1$ in the convention of Section 1.3, because a hot outlet calls for more flow. The app's Show Model text names

it the same way. Derivative acts on the measurement, after a first-order filter with τ_f defaulting to 1.0 s, so a setpoint move gives the pump no derivative kick. The filter advances in MANUAL as well as in AUTO so that the estimate is current the instant AUTO is selected.

The output is clamped to 20–120 % of design pump speed and carries no rate limit, because the pump lag τ_p of 4 s in (7.13) already sets how fast flow can change. Anti-windup is the back-calculation of Equation (1.16), with a tracking time T_t that the application sets equal to τ_I, so at $\tau_I = 20$ s each scan removes 1 % of the clipped excess. No classical run in the original comparisons of Section 7.6 reaches either limit inside its observation window (the first contact, at $K_c = 2.2$, comes at $t = 615$ s), so the tracking time does not affect a reported result. An integral time of 300 s is read as $\tau_I \to \infty$, so proportional only is a selectable setting rather than a limit that is approached, and the integral update is then skipped. MANUAL tracking and the copy of each AUTO output to the operator's pump station make transfer bumpless in both directions, as Section 1.6.2 describes.

Two tunings recur through the scenarios. The cautious pair is $K_c = 1.0\,\%/^\circ$C with no reset action, and the tuned pair is $K_c = 1.2\,\%/^\circ$C with $\tau_I = 20$ s and $\tau_D = 0$.

7.5.2 Integrator Count and Steady-State Offset

Whether this loop requires reset follows from the structure of the plant, and the answer is the opposite of the one a rod-driven reactor gives [1]. The rods of Chapter 2 are driven in speed, so the drive supplies the integrator there. The manipulated variable here is a pump *position*, the pump follows it through the first-order lag of (7.13), and no element downstream of the controller integrates. A proportional controller can hold its output away from its bias only by holding an error, so the loop settles at an offset. Counting the integrators around the loop settles the question before any test is run, and Section 7.6 measures the cost of getting it wrong.

7.5.3 Load Feedforward and Dynamic Compensation

The turbine load is known the instant an operator asks for it, and the outlet thermocouple is the last measurement in the plant to register it, because the disturbance propagates through the cold leg, the heat exchanger, a 16 s loop transit, the core and a thermowell. Load feedforward therefore enters as the controller bias u_b, written $u_{\mathrm{bias},k}$ here, as shown in Equation (7.16), so that the clamp of Equation (1.15) and the bumpless transfer both account for it.

$$\begin{aligned} u_{\mathrm{bias},k} &= 100 + K_{ff}\,\tilde{L}_k \\ \tilde{L}_k &= \alpha_{ff}\,\tilde{L}_{k-1} + (1-\alpha_{ff})(L_k - 100) \end{aligned} \tag{7.16}$$

where L is the ramp-limited turbine load in percent, $\tilde{L}$ is that load deviation after the feedforward lag, and $\alpha_{ff} = \exp(-\Delta t_c/\tau_{ff})$ is the same exact zero-order-hold coefficient used on the instruments. The feedforward gain is derived from the steady state rather than tuned. Holding the outlet means holding the core rise, and the rise follows the power, but the power does not follow the load one for one, because a 30 % load reduction moves the cold leg 30 °C and takes the power only to 80.0 %. Solving the steady state for the pump speed that puts the outlet on setpoint at each load gives a local requirement of 0.48 % of pump speed per percent of turbine load at full load, rising to 0.51 % across a reduction to 30 %, and the default gain is 0.50. The lag τ_{ff} defaults to 20 s and it carries most of the benefit, for reasons Section 7.6 quantifies.

7.5.4 Inner Flow Cascade

The third architecture closes a flow loop around the pump. The temperature controller output becomes a flow setpoint, and an inner controller with $K_{c,F} = 2.5\,\%/\%$ and $\tau_{I,F} = 6$ s acts directly on the flow transmitter. It is reverse acting, $a = +1$, because a higher command gives a higher flow, and it has the same 20 to 120 % limits and the same algorithm with $T_t = \tau_{I,F}$. The inner flow loop of the level-to-flow cascade of Volume 3, Chapter 2 meets a supply pressure swing that changes what its valve delivers, as the delivery fault of scenario 8 changes what this pump delivers. Switching

architecture in mid-run changes what the outer output *means*, from a pump command in two of the three forms to a flow setpoint in the third, and moves the bias by the whole feedforward term. The switch is therefore handled as a transfer. For one execution every controller tracks the pump command the plant already has, which is the same mechanism that makes MANUAL to AUTO bumpless.

7.5.5 Model Predictive Control

The fourth structure is the predictive option of Section 1.6.3, written for this plant. It plans fuel pump commands from filtered outlet temperature $\hat{T}_{fo}$, filtered flow $\hat{f}$, current load, received load order and known pump, secondary-flow and rod-trim commands. It has no reactor-power measurement. Its constants are frozen at the defaults in Table 7.3, so plant changes remain model mismatch. The nominal equilibrium outlet $F(f,\ell)$ follows the criticality and heat balances of Equations (7.5)–(7.11), with load $\ell = 100L$ in percent. The initial rod position is inferred from the initial measured temperature, flow and load through those balances. It is never read from the hidden rod state. Known rod trim shifts that inferred position.

The static map is separated into pump and load contributions, $g(f,\ell) = F(f,\ell) - F(f_0,\ell)$ and $h(\ell) = F(f_0,\ell) - F_0$, where F_0 is the map value at initial flow, load and commands. Subscript 0 denotes initialization. At each decision, $K_f = \partial g/\partial f$ is approximated by symmetric 0.1 % differences about measured flow, clipped to 21–119 %, and $c_f = g(\hat{f},\ell) - K_f\hat{f}$. The design pump/load tangents are $-0.685620/+0.329341$ °C per flow/load percentage point. Independent symmetric ± 5 point steps give pump gain -0.701421, lag 11.864 s, delay 1.739 s and RMS residual 0.072 813 °C. The load fit gives $+0.342366$, 10.874 s, 19.927 s and 0.062 597 °C, including load slew. Rounded dynamics produce the five-state predictor in Equation (7.17).

$$\begin{aligned} d_{1,j+1} &= u_j, \qquad d_{2,j+1} = d_{1,j}, \\ p_{j+1} &= a_p p_j + (1-a_p)(K_f d_{2,j} + c_f), \\ q_{j+1} &= a_L q_j + (1-a_L) h(\ell_{j+1-20}), \\ b_{j+1} &= b_j, \qquad y_j = T_0 + p_j + q_j + b_j. \end{aligned} \tag{7.17}$$

Here j counts 1 s prediction steps, $a_p = \exp(-1/12)$, $a_L = \exp(-1/11)$, d_1, d_2 are pump command slots in percent, and p, q, b, y, T_0 have units of °C. The initial state is $[u_0, u_0, 0, 0, 0]$ and T_0 is initial measured outlet. Past delayed loads come from observed history. Future loads hold the current value with preview Off. Preview On projects only the remaining slew toward an order already received, at the nominal one percentage point per second. A declared error adds 20 load points to that projected order, clipped to 0–120 %. Neither setting reads a later scenario event. The fast-ramp plant's four-point-per-second slew is explicit model mismatch.

The saved model propagates the applied command between decisions. Its one-step residual updates the bias as $b^+ = b^- + [1 - \exp(-\Delta t/10)](\hat{T}_{fo} - T_0 - p^- - q^- - b^-)$, where Δt is elapsed seconds. This output-bias correction supports measured offset removal [20]. It is not a Kalman filter. Four held-out flow/load responses have raw RMS residuals 0.0967–0.4091 °C. Full precursor and graphite transients remain outside this predictor.

A receding-horizon quadratic program [21, 22] is solved at initialization and every 1 s, with the first command held over five ordinary scans. Its 120 s horizon uses ten blocks with lengths $(1,1,2,3,5,8,12,20,28,40)$ seconds, as shown in Equation (7.18).

$$\begin{aligned} \min J = &\sum_{j=1}^{120} (y_j - T_{\mathrm{SP}})^2 + 10(y_{120} - T_{\mathrm{SP}})^2 \\ &+ 10 \sum_{m=1}^{10} (v_m - v_{m-1})^2 + 20 s^2, \\ 20 \le v_m \le \bar{u}_\beta, &\quad |v_m - v_{m-1}| \le 2, \\ T_{\mathrm{SP}} - 3 - s \le y_j \le T_{\mathrm{SP}} + 3 + s, &\quad s \ge 0. \end{aligned} \tag{7.18}$$

Table 7.5 Scenario Sequence and the Mechanism Each Exercises

#	Scenario	What it exercises
1	Pump parked, load to 70 %	self-regulation, the pinned mean
2	Drive the pump by hand	the flow the steady state asks for
3	Proportional only, $K_c = 1.0$	offset on a position actuator
4	Reset added, $\tau_I = 20\,\mathrm{s}$	the gain limit set by transport
5	Rod drop at 120 and 40 % flow	reactivity worth against a moving β_{eff}
6	Fuel pump trip at 9 % power	circulation as a reactivity device
7	45 % load ramp at 4 %/s	feedforward timing
8	Pump delivery falls to 85 %	a disturbance inside the inner loop
9	Heat exchanger fouled by a quarter	a capacity limit behind a well-behaved loop
10	Graphite coefficient made positive	a requirement on the coefficient sum
11	Thermocouple noise and a 3 °C bias	actuator travel and a bias
12	Graded challenge, three events	tuning for the final operating point

Here v_m is a block command in percent and s is one shared temperature slack in °C across all 120 samples. Numerical weights apply to Celsius and percentage-point values. Commands and moves are hard constraints. The $\pm 3\,^{\circ}\mathrm{C}$ band is soft. Equation (7.4) decreases with flow, so a minimum β_{eff} imposes a maximum command, not a pump floor. Bisection of the nominal relation gives $\bar{u}_\beta = 105.412504\%$ at 415 pcm, intersected with the physical 20–120 % span. An infeasible previous command is clipped into that span before defining v_0 and its first move limit.

The shared solver's complete plan, transformed temperatures, objective and slack must be finite, correctly shaped, optimal and consistent with every box and move limit. Rejection invokes the tracked outlet PI, with the nominal ceiling retained even during sustained failure. Both PIs track the applied command under MPC. A forced 180 s failure test reaches 105.412504% with zero first-transfer and subsequent Manual-transfer bump. Manual follows the requested pump station through the physical clamp. Existing pump trip, coastdown, drain and annunciations still act.

7.6 Results

Table 7.5 lists the twelve original scenarios in the order the mechanisms arise. A thirteenth case, described in Section 7.6.9, precedes the graded challenge in the extended sequence. The app's selector numbers that case 12 and lists the challenge after it without a number. Each one presets the mode, the architecture and the tuning and schedules its own events. Integrated absolute error is quoted in °C s and pump travel is the total distance the pump command moved, in percent of design flow. Both measures accrue for as long as a run continues, so the standard observation window is a 600 s run from the settled design point and every integrated error and pump travel below is read at that instant unless another window is named. Three places name another one, the noise comparison at 400 s, the coefficient sweep of Table 7.9 at 1800 s because the slowest case has not settled sooner, and the graded challenge over its own 900 s.

7.6.1 Self-Regulation Without Control

Scenario 1 holds the pump at 100 % and takes the turbine to 70 % with no operator action. Reactor power falls to 78.2 % in ten minutes with nothing commanded, because the cold leg rose from 480 to 510 °C by (7.12), the duty fell with it, the salt heated, and the fuel temperature coefficient of −5 pcm/°C removed the reactivity that fission needed. The control problem appears in the temperatures rather than in the power. The mean fuel salt ends at 615.8 °C, which is 0.8 °C from where it started, while the outlet ends at 655.0 °C and the inlet at 576.6 °C. The core temperature rise closed from 100 to 78 °C. Load following occurs without any control action and holding the outlet temperature does not.

Table 7.6 Gain Sweep on a 30 % Load Reduction at $t = 80$ s, Every Column Read at $t = 600$ s. The First Column Has No Reset Action. Pump Travel Is the Total Distance the Command Moved.

K_c (%/°C)	1.0	1.2	1.8	2.0	2.2
τ_I (s)	—	20	20	20	20
Peak deviation (°C)	6.89	4.18	3.14	7.77	24.5
Error at 600 s (°C)	6.20	0.00	0.68	4.19	10.3
Integrated error (°C s)	3043	262	386	748	1563
Pump travel (%)	10.3	22.3	150	407	1002

Scenario 2 gives the same disturbance to an operator with the pump as the manipulated variable. The steady state requires 85.2 % pump speed, and moving the pump there raises β_{eff} from 419 to 433 pcm, so the corrected plant carries 14 pcm more margin against a reactivity disturbance than the uncorrected one.

7.6.2 Proportional Offset and the Transport-Limited Gain

Table 7.6 sweeps the proportional gain against the same 30 % load reduction, on the schedule of scenario 4 with the load step at $t = 80$ s and every gain set before the run starts.

The proportional-only column settles 6.20 °C below setpoint and remains there. The pump ends at 93.8 %, which is the 100 % bias plus a gain of 1.0 acting on the 6.20 °C error, against the 85.2 % the plant requires. The controller is 8.6 points short, for the structural reason given in Section 7.5. Scenario 3, which runs the same controller with the load step at $t = 60$ s, reads 6.19 °C at the same instant. Adding reset removes the offset completely and improves the peak from 6.89 to 4.18 °C, with the integrated error falling to 262 °C s against 3043 °C s for the proportional-only run over the same window. Pump travel rises from 10.3 to 22.3 %, which is the cost of reset.

Reading Table 7.6 across the remaining columns gives the characteristic signature of a loop limited by dead time rather than by process gain [1]. The pellet grill of Volume 5, Chapter 2 also commands its controlled variable through a transport delay. The peak deviation is still improving at $K_c = 1.8$, at 3.14 °C, but the pump travel has risen from 22.3 to 150 % to buy 1.04 °C, and the oscillation is already growing slowly, so the error at 600 s is 0.68 °C rather than zero. At 2.0 the travel is 407 % and the swing is still growing at the end of the run, and at 2.2 the loop is unstable, reaching a 24.5 °C deviation and 1002 % of travel by 600 s and driving the pump onto a stop at 615 s. The limit is transport. The heat exchanger cools the salt and that salt does not reach the core inlet for a full 16 s loop transit by (7.10), so roughly half of what the controller does returns to it a quarter of a minute late. Halving $\tau_{L,0}$ to 8 s, set before a Reset so that the plant starts lined out on the shorter loop, moves the gain limit up and moves β_{eff} up with it. At $K_c = 2.0$ the pump then travels 45 % instead of 407 %, it takes $K_c = 2.8$ to spend a comparable 448 %, and β_{eff} at design flow rises from 419 to 458 pcm, because on this plant the loop dead time and the stability margin are two readings of one pipe length.

> ✦ DID YOU KNOW? **More Gain, a Larger Swing**
>
> From the simulator: on the same 30 % load cut, raising the proportional gain from 1.2 to 2.2 with $\tau_I = 20$ s takes the peak deviation from 4.18 to 24.5 °C and the pump travel in 600 s from 22.3 to 1002 % (Table 7.6). The 16 s loop transit, not the pump, sets the limit.

7.6.3 Feedforward Timing on a Fast Load Ramp

Scenario 7 applies a 45 % load reduction at 4 %/s, four times the normal ramp rate. Table 7.7 compares the three architectures against it.

The middle column carries the result. A feedforward of the correct steady-state size, applied without dynamic compensation, is *worse* on the peak than no feedforward at all, 7.42 against 7.05 °C, because it moves the pump the moment the turbine moves while the disturbance has not yet traveled

Table 7.7 A 45 % Load Reduction at 4 %/s, $K_c = 1.2$ and $\tau_I = 20\,\text{s}$ Throughout, Every Column Read at $t = 600\,\text{s}$.

	Feedback alone	Static feedforward	Lagged feedforward
K_{ff} (%/%)	0	0.50	0.50
τ_{ff} (s)	—	0	20
Peak deviation (°C)	7.05	7.42	2.64
Integrated error (°C s)	416	329	188
Pump travel (%)	44.6	73.3	59.9

Table 7.8 Pump Delivery Cut to 85 % of Command, Read at $t = 600\,\text{s}$. Settling Is to Within 0.5 °C of Setpoint and Is Measured From the Step.

	Cascade $K_c = 0.6$	Single $K_c = 1.2$	Single $K_c = 1.8$
Peak deviation (°C)	1.84	5.79	5.20
Integrated error (°C s)	77	314	299
Settling time (s)	89	151	272
Pump travel (%)	55	45	116

through the cold leg and round the loop. A correction that arrives before its disturbance has to be withdrawn, and the withdrawal is itself a disturbance [1]. Lagging the same feedforward by 20 s takes the peak to 2.64 °C and the integrated error from 416 to 188 °C s, which is 63 % off the peak and 55 % off the integrated error. On the normal 1 %/s ramp the same arrangement takes the peak from 4.18 to 1.26 °C and the integrated error from 262 to 82 °C s.

> **TRY IT IN THE APP** **7 · Load Feedforward, and When It Should Arrive**
>
> Load scenario 7 and run it to 600 s with the feedforward lag at 20 s: the outlet peaks 2.64 °C from setpoint. In Advanced, press Reset, set the feedforward lag to 0 before the ramp at 60 s and run again. The peak grows to 7.42 °C, worse than the 7.05 °C of feedback alone.

7.6.4 Cascade Performance and Noise Response

Scenario 8 drops the pump delivery to 85 % of its command, which is a disturbance inside the inner loop, because it changes the flow without changing any variable the temperature loop measures. The inner loop measures the flow directly and returns it to 99.1 % of design within 10 s and 100.2 % at 20 s. The single loop is still at 90.2 % at that same 10 s mark, 94.6 % at 30 s and 99.0 % only after 90 s. Table 7.8 carries the effect on the controlled variable.

Raising the single-loop gain does not close the gap, because the disturbance enters where the outer measurement cannot register it for a heat exchanger and a loop transit. The outer loop also has to be retuned, and in the opposite direction to the one a faster actuator suggests. Closing the inner loop makes the actuator faster, so an outer gain that was correct on the single loop now oscillates. At an outer gain of 1.2 the cascade spends 257 % of pump travel and never settles inside 0.5 °C, because a swing of about 1 °C persists through the whole window and is still growing, to 2.2 °C by $t = 1500\,\text{s}$, against 55 % and 89 s at 0.6. Figure 7.3 draws the feedforward and cascade comparisons side by side.

The cost appears in scenario 11, which puts 0.6 °C of noise on the thermocouple and 1.0 % on the flow transmitter with the plant at steady state. At the default single-loop tuning the pump travels 96.7 % in 400 seconds, and the cascade at $K_c = 0.6$ and $\tau_I = 40\,\text{s}$ travels 1256 % over the same 400 seconds, because the inner loop is fast and follows the transmitter noise at full speed. Derivative action is worse still. Adding 5 s of derivative with the filter removed takes the travel to 1606 %, a

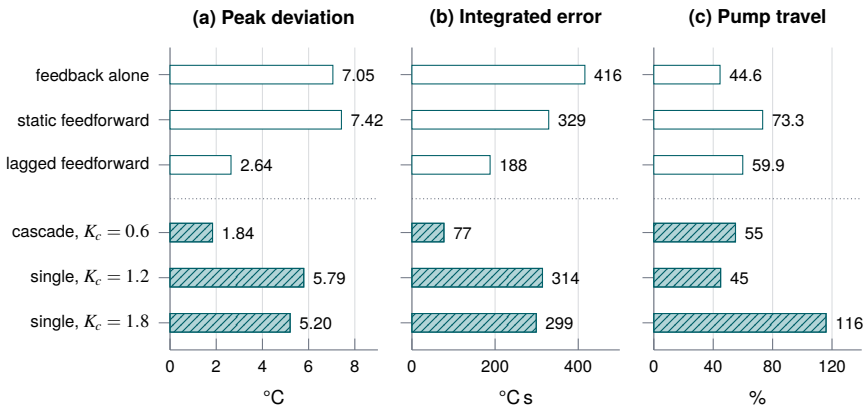


Figure 7.3 Tables 7.7 and 7.8 drawn as bars, every value read at $t = 600$ s. White bars, above the dotted line: the 45 % load reduction at 4 %/s of scenario 7 under three architectures. Hatched bars, below it: the pump delivery cut to 85 % of command of scenario 8, on the cascade and on the single loop at two gains. Notice that the static feedforward has the worst peak of the upper three, and that raising the single-loop gain from 1.2 to 1.8 does not close the gap to the cascade against the delivery cut, while it more than doubles the pump travel.

seventeenfold increase, and moves the true outlet as far as 1.32 °C from setpoint in the first 200 s, against 0.49 °C without derivative, while filtering the same derivative at 4 s recovers the travel only to 282 %. Five further noise seeds give 90 to 103 %, 1215 to 1248 %, 1515 to 1749 % and 232 to 355 % for the same four cases, so every comparison holds well outside the seed spread. The thermocouple bias that the same scenario applies at $t = 200$ s is a failure no tuning corrects. An integrating controller drives the measurement to setpoint, so with the instrument reading 3 °C high the true outlet settles at 662.0 °C, 3.0 °C low, and the pump averages 104.3 % over the last 100 s instead of 100. Every controlled variable reads correctly throughout.

7.6.5 Rod Insertion at Four Pump Speeds

Scenario 5 drops a rod worth −100 pcm into the core in MANUAL at four pump speeds. Each slower speed is set before a Reset, so the run starts lined out at that flow, because a full-power core stepped straight to 40 % flow reaches the freeze plug. The insertion is 0.247 dollars at 120 % flow, where β_{eff} is 404 pcm, and 0.203 dollars at 40 %, where it is 493 pcm. The transients scale with the insertion in dollars. Power falls to 58.6 % of where it started at 120 % flow, 59.6 % at 100, 62.4 % at 60 and 63.9 % at 40, monotonically across the whole pump range. A liquid-fueled reactor is therefore analyzed at its fastest pump speed, because that is the condition with the least margin. Figure 7.4 shows the run at 120 % flow as the plant returns to criticality.

7.6.6 Fuel Pump Trip

Scenario 6 trips the fuel pump. The precursors stop leaving the core, so every one of them decays where its neutrons are multiplied, β_{eff} climbs from 419 toward 604 pcm as the pump coasts down on its 10 s flywheel time constant, and the reactivity the core must hold falls from 231 to 46 pcm. The difference appears as a net insertion peaking at +100 pcm, or 0.172 dollars. Power rises from 9.0 to 16.8 % thirty-six seconds after the trip, an increase of 86 %, with no rod motion and no valve motion anywhere. The fuel salt then heats, the fuel temperature coefficient takes the reactivity back, and the power turns over.

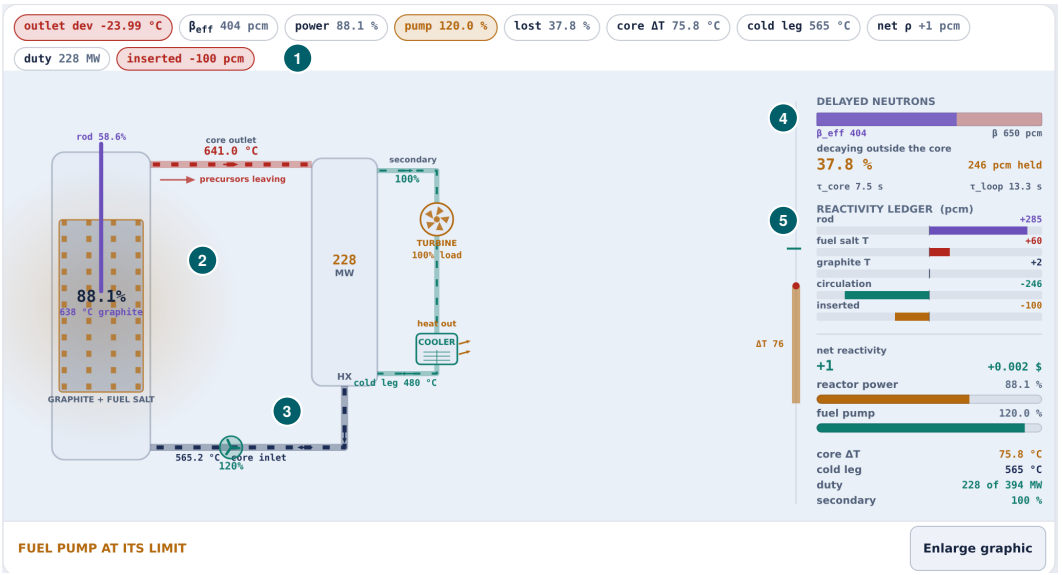


Figure 7.4 Recovery from a −100 pcm rod drop at the top of the pump range, in the plant view of the app. Callouts: (1) the status badges, (2) the core with its power and graphite temperature, (3) the fuel pump on the cold leg, (4) the delayed neutron panel and (5) the reactivity ledger. Scenario 5 runs in MANUAL with the fuel pump held at 120 % of design flow, the rod is dropped at $t = 60$ s, and the state shown is $t = 200$ s. At this pump speed β_{eff} is 404 pcm against the static 650 pcm, with 37.8 % of the precursors decaying outside the core and transit times of 7.5 s inside and 13.3 s outside. The reactivity terms are +285 pcm of rod, −246 pcm of circulation, −100 pcm of dropped rod, +60 pcm from the cooler fuel salt and +2 pcm from the graphite, summing to a net +1 pcm as the plant returns to criticality at a lower power of 88.1 %. The core outlet is 641.0 °C against a 665 °C setpoint, because no loop is holding the temperature.

> **TRY IT IN THE APP** **6 · Fuel Pump Trip at Low Power**
>
> Load scenario 6 and press Run. The reactor idles at 9 % power, and at 30 s the fuel pump trips and coasts down. Watch the reactor power: with no rod and no valve moving, it climbs to 16.8 % thirty-six seconds after the trip, then turns over as the hotter salt takes the reactivity back.

The scenario is set at 9 % power deliberately, and the reason is a result in itself. Returning the load, the secondary flow and the setpoint to their design values before a Reset starts the same case at full power, with the same trip at 30 s. At 250 MW with the flow gone the core cannot carry the heat away and it heats at $(1 - f_G)P_0/C_f = 10.4\,^{\circ}\text{C/s}$. A coefficient of −5 pcm/°C is then removing 52 pcm every second, which overwhelms the 231 pcm long before the coastdown has delivered all of it, so the power falls instead of rising and the core outlet reaches the 750 °C freeze plug 19.8 s after the trip. The mechanism is present at full power and is not observable there, because a faster effect masks it. Testing a mechanism at the operating point where it is observable and then reasoning about the operating point where it is not is what the pump experiments at the Molten Salt Reactor Experiment did, and they were run at zero power for the same reason [11, 12].

7.6.7 Heat Exchanger Fouling and a Positive Graphite Coefficient

Scenario 9 fouls the heat exchanger, taking UA from 3.20 to 2.40 MW/°C. The outlet dips 4.42 °C while the loop works, returns to within 0.12 °C by the end of the 600 s window, and the plant loses 13.7 % of its power, which is not a failure of the loop but the loop doing what it was asked. It holds

Table 7.9 Graphite Coefficient Sweep on a 30 % Load Reduction at $t = 80$ s, With α_f Held at −5 pcm/°C. Open-Loop Rows Hold the Pump Fixed, Closed-Loop Rows Use $K_c = 1.2$ and $\tau_I = 20$ s and Are Read at $t = 1800$ s.

α_g (pcm/°C)	−1	+2	+4	+6
Outlet deviation at 1 h (°C)	9.9	16.9	69.5	155
Steady deviation (°C)	9.9	16.9	128.6	unstable
Power at 1 h (%)	78.4	74.7	45.2	0
Integrated error (°C s)	269	299	378	914
Pump travel (%)	25	28	36	170

the core rise by matching the flow to whatever power the heat sink allows, and the heat sink now allows less. Power settles at 86.3 % of rated and the pump reads 89.8 % at 600 s on its way to 90.0 %. The secondary pump cannot buy all of the power back, because K in (7.11) saturates at UA. At 100 % secondary flow the fouled exchanger gives 1.575 MW/°C and 86.3 % power, at 125 % it gives 1.704 and 92.7 %, and at 150 % it gives 1.798 and 97.4 % once settled (the 200 s holds of the scenario's completion steps read 92.8 and 97.5 %), so half again as much pumping recovers 11 of the 14 points that were lost. The controlled variable reads correctly through the whole event, and the evidence of the fault is in the power, the duty and the pump position, none of which carries a controller.

Scenario 10 makes the graphite temperature coefficient positive while leaving the fuel salt coefficient at −5 pcm/°C, and Table 7.9 reports the sweep. The plant still follows the turbine, because the salt coefficient alone is sufficient, and what changes is the cost. The open-loop rows hold the pump fixed through the same 30 % load reduction. At the design −1 pcm/°C and at +2 pcm/°C the plant is settled within an hour, at 9.9 and 16.9 °C below setpoint. At +4 pcm/°C it is not settled after an hour, reading 69.5 °C low with the power at 45.2 % and still falling toward an eventual 128.6 °C at 14.3 % power, because the sum of the coefficients is only −1 pcm/°C and the restoring stiffness has almost gone. At +5 pcm/°C the sum is exactly zero, the temperature has no stiffness at all, and the plant walks down to 510 °C until fission stops. The requirement is on the sum. In AUTO at the tuned settings the loop absorbs the first two cases at unchanged cost and fails on the last, with the integrated error going from 269 °C s at −1 pcm/°C to 914 °C s at +6 and the pump travel from 25 to 170 % over the same 1800 s. The graphite coefficient of a graphite-moderated reactor moves with fast neutron damage over a few years, so a control system designed against a coefficient measured at startup carries an undeclared dependency on the sum.

7.6.8 Graded Challenge

The challenge applies a load reduction to 75 % at 60 s, an unannounced −100 pcm rod drop at 260 s and a return to full load at 460 s, and asks for the outlet held within ±2 °C through the last third of the run, the cold leg never below 500 °C, the pump never pinned at a limit and total travel under 150 %. The rod drop is 100 pcm against a combined coefficient of −6 pcm/°C, which is 16.7 °C of mean fuel temperature that has to go away and that no pump setting can prevent, so every tuning takes a large excursion. Across 1320 tunings swept over both architectures (K_c from 0.4 to 2.8, τ_I from 10 to 300 s, and with feedforward K_{ff} of 0.5 or 1.0 and τ_{ff} from 0 to 60 s) the smallest peak deviation is 16.2 °C, and the suggested tuning peaks at 22.8 °C. What is graded is the recovery. Feedforward at $K_c = 0.6$, $\tau_I = 30$ s, $K_{ff} = 0.5$ and $\tau_{ff} = 20$ s settles to within 0.77 °C and holds, for 102 % of pump travel, with the cold leg never below 530.4 °C, while the proportional-only tuning is up to 14.05 °C off over the last third, seven times the band, and still 13.2 °C off at 900 s, with 158 % of travel, although its inlet stays above 540.4 °C and its pump is never pinned. The suggested tuning passes with 1.23 °C to spare. The intermediate tuning is the worst of the three, because the settings that were correct in scenario 4 are unstable once the rod drop has moved the plant to about 73 % pump speed, and they finish swinging 19.8 °C and spending 1718 % of pump travel. The process gain and the loop dead

Table 7.10 Fixed-Setting Comparisons Over 900 s. Peak, IAE, Command Travel and Saturation Use the Shared 5 Hz Metrics. t_β and D_β Integrate Actual Margin Deficit at the 0.02 s Physics Step.

Case	Structure	Peak °C	IAE °C s	Travel %	Sat. %	Min. β pcm	t_β s	D_β pcm s	Solves / rejects	Mean µs
Normal load cut	Cascade	4.089	389.52	123.79	0.00	419.42	0.00	0.00	0/0	0.0
	FF	1.257	87.24	27.64	0.00	419.42	0.00	0.00	0/0	0.0
	Fixed FF	1.471	86.75	19.75	0.00	419.42	0.00	0.00	0/0	0.0
	MPC On	2.494	129.13	16.84	0.00	419.42	0.00	0.00	900/0	131.2
	MPC Off	2.150	98.90	16.82	0.00	419.42	0.00	0.00	900/0	93.9
	MPC +20	1.886	712.75	19.05	0.00	418.09	0.00	0.00	900/0	107.8
Fast load ramp	Cascade	96.847	4545.18	3737.18	1.58	419.42	0.00	0.00	0/0	0.0
	FF	2.640	234.97	77.64	0.00	419.42	0.00	0.00	0/0	0.0
	Fixed FF	3.061	169.45	37.19	0.00	419.42	0.00	0.00	0/0	0.0
	MPC On	5.457	253.37	27.99	0.00	419.42	0.00	0.00	900/0	95.7
	MPC Off	5.376	235.98	27.92	0.00	419.42	0.00	0.00	900/0	90.0
	MPC +20	4.061	695.21	30.64	0.00	418.09	0.00	0.00	900/0	104.9
Rod, 120% start	Cascade	191.807	28588.28	3342.76	32.98	404.17	66.60	694.74	0/0	0.0
	FF	18.771	1140.93	281.07	6.69	404.17	68.52	710.12	0/0	0.0
	Fixed FF	20.460	867.14	101.55	6.69	404.17	70.58	723.27	0/0	0.0
	MPC On	21.838	1302.20	44.16	0.00	404.17	65.06	44.68	900/0	131.7
	MPC Off	21.838	1302.20	44.16	0.00	404.17	65.06	44.68	900/0	116.9
	MPC +20	22.124	1404.62	44.38	0.00	404.17	65.06	44.68	900/0	112.5
Challenge	Cascade	23.497	3008.17	484.98	0.00	419.42	0.00	0.00	0/0	0.0
	FF	22.821	1597.81	101.85	0.00	419.42	0.00	0.00	0/0	0.0
	Fixed FF	18.783	1747.45	284.80	0.00	419.42	0.00	0.00	0/0	0.0
	MPC On	23.502	1672.03	86.74	0.00	419.42	0.00	0.00	900/0	132.6
	MPC Off	23.502	1651.28	86.75	0.00	419.42	0.00	0.00	900/0	109.6
	MPC +20	24.187	1833.73	91.25	0.00	418.09	0.00	0.00	900/0	117.3
Margin case	Cascade	6.659	786.74	234.68	0.04	409.12	439.92	2385.47	0/0	0.0
	FF	4.953	299.24	62.59	0.00	409.29	476.76	2423.64	0/0	0.0
	Fixed FF	4.971	320.75	53.24	0.00	409.29	477.90	2441.66	0/0	0.0
	MPC On	4.978	2324.30	39.35	0.00	415.00	0.00	0.00	900/0	106.3
	MPC Off	4.986	2262.85	39.37	0.00	415.00	0.00	0.00	900/0	107.3
	MPC +20	4.809	2459.75	38.57	0.00	415.00	0.00	0.00	900/0	122.3
Challenge	Shipped P	27.370	11054.41	158.31	0.00	419.42	0.00	0.00	0/0	0.0

FF is shipped feedforward, or suggested feedforward on the challenge. Fixed FF uses one predetermined model rule. MPC +20 biases the received order by 20 load percentage points. The final row preserves the shipped proportional-only challenge. $t_\beta = \int \mathbf{1}_{\beta_{\text{eff}}<415}\, dt$ and $D_\beta = \int \max(0, 415 - \beta_{\text{eff}})\, dt$. The mean is an observed postpriming solver wall time, not full control cost.

time both move with the actuator, so the gain belongs to the operating point the plant reaches after the disturbance rather than to the one it starts from.

7.6.9 Predictive Control and the Margin Constraint

Table 7.10 extends the observation window to 900 s for all five cases. Runs use seed 20260829, identical scenario events, instruments and noise, and unchanged MPC settings. Challenge comparisons first apply the suggested settings, except the explicit shipped P-only row. Classical feedforward receives current ramped load through its original lag. MPC Off therefore supplies the adjacent comparison with matching load information. Preview On has additional received-order projection, without access to future orders or rod insertions.

The fixed classical row uses a rule declared before its results: response time 20 s from the load delay, rounded pump lag/delay 12/2 s, $K_c = 12/[0.701421(20+2)] = 0.777642$ %/°C, $\tau_I = 12$ s, $\tau_D = 0$, $K_{ff} = 0.342366/0.701421 = 0.488103$ and $\tau_{ff} = 20$ s. This one setting is reused across cases. No tuning search or global optimality claim is made.

For normal load cuts, fixed FF IAE 86.75 °C s beats MPC On's 129.13 °C s and shipped FF's 87.24 °C s, with travel 19.75/16.84/27.64%. On the fast ramp, fixed FF also beats MPC On in IAE

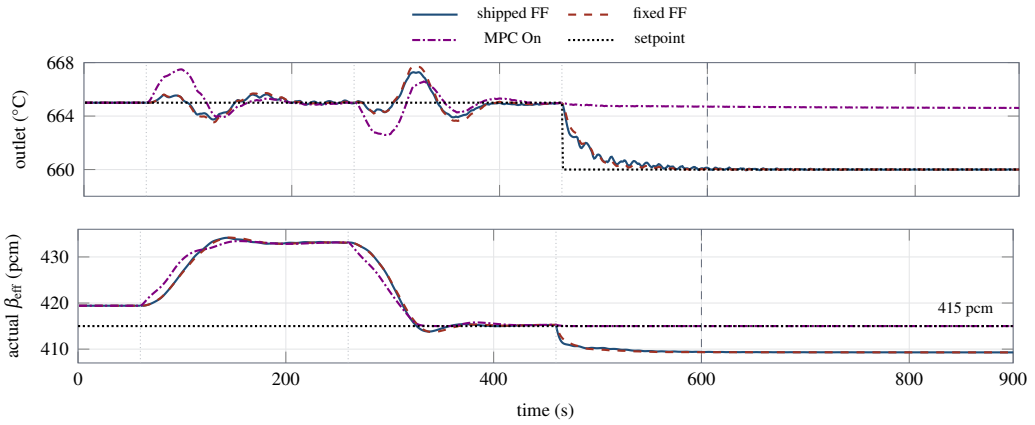


Figure 7.5 Outlet tracking (upper panel) and actual delayed-neutron margin (lower panel) over 900 s. Shipped FF is the solid line, fixed FF the dashed line and MPC On the dash-dotted line. The thick dotted line is the setpoint above and the 415 pcm margin below. Thin dotted event markers locate the received orders and setpoint change, and the thin dashed marker is 600 s. FF follows the lower setpoint by using more pump speed, while MPC remains at its nominal ceiling. MPC Off is retained in Table 7.10 to keep the three plotted trajectories legible.

(169.45 against 253.37 °C s). Preview Off lowers MPC IAE to 98.90 and 235.98 °C s on the load cases, so preview supplies no measured benefit. The +20 error reduces normal-load peak to 1.886 °C while raising IAE to 712.75 °C s.

Fixed FF improves challenge peak from MPC's 23.502 to 18.783 °C, but fails the unchanged target through last-third error 3.66 °C and travel 284.80%. Suggested FF and MPC On/Off/+20 pass, with last-third errors 0.77/0.95/0.95/0.94 °C. Cascade and shipped P fail at 4.86/14.05 °C. Coldest inlets are 521.5/524.7/540.4 °C for fixed FF, Cascade and shipped P, and 527.4–530.4 °C for passing rows. No challenge row is physically pinned. Whole-run peak and IAE do not replace this recovery verdict.

The margin case orders 70% load at 60 s, 110% at 260 s and 660 °C setpoint at 460 s. MPC On gives IAE 2324.30 °C s and zero seconds below 415 pcm, against FF's 299.24 °C s and 476.76 s. At 600 s, outlet/margin are 664.707 °C/415.000 pcm for MPC, 660.129 °C/409.397 pcm for FF and 660.036 °C/409.342 pcm for fixed FF, as shown in Figure 7.5. The ceiling leaves persistent temperature error.

At delivery gain 1.15, feasible MPC commands still give 492.54 s below the actual margin, deficit 2560.77 pcm s and minimum 409.008 pcm. The rod case starts at 120% flow and 404.174 pcm. Command clipping requires 65.06 s for actual margin recovery, with deficit 44.68 pcm s. Scorecard travel 44.16% excludes the initial clip, whereas cumulative pump travel includes it and reads 58.75%. Nominal command feasibility does not remove initial actuator lag.

Ten full 900 s refinements at half the physics step preserve tested comparison orderings and challenge verdicts, with maximum IAE change 0.000 146 °C s. All 13 MPC scenarios retain finite outputs and 20–120% commands through 900 s, with graphite run for 3600 s. The pump-trip run still reaches 750.030 °C and spends 572.40 s after drain.

Warm Chrome observations at 1×/20× give solver means 280/282 µs, full decision-scan means 590/535 µs and all-control load 0.081/1.195% of a core near 60 frames/s. Each uses ten wall seconds after 200 s simulated warm-up. Guard stops are zero, and guard-off playback also approaches 60 frames/s, establishing no guard pacing benefit. Playback through the 60 s order reproduces 326 rows and 65 solves at both speeds, excluding wall-time columns.

7.7 Discussion

The protective function in this plant is not a control system. A length of the drain line is kept solid by a cooling fan, and when the salt gets hot enough the plug melts and the entire fuel inventory runs into a tank shaped so that it cannot go critical. The application fires that event at a core outlet of 750 °C, which a full-power pump trip reaches 19.8 s after the trip with no signal, no power supply and no operator action. As an engineering argument it is stronger than anything a control system can offer, because there is no logic to fail and no actuator to stick, and it is one of the reasons liquid fuel is being revisited [5, 6]. As a licensing argument it is awkward, because the regulatory apparatus for reactors is built around actuated protection systems with test intervals, redundancy, single-failure criteria and demonstrated failure rates, and none of those concepts has an obvious meaning for a plug of salt. A safety case that moves work out of a system and into a material moves it into materials qualification rather than removing it.

The second implication concerns what a control system reports. Ordinary load following moves 10 % of the plant's margin against every reactivity disturbance while the temperature loop is minimizing an unrelated objective, so the stability margin is reported beside the control error rather than left implicit. The effective delayed fraction, the fraction of precursors lost, the cold leg margin and the pump travel are therefore presented together with the outlet deviation.

The third is about evidence. One reactor of this kind has ever operated, for four years, ending in 1969 [10], and every number in this application came out of a model. The circulation loss depends on a loop transit time a real plant would have to measure, the graphite coefficient on a material property that changes in service, and the omission of xenon on a pump bowl stripping efficiency measured on one machine half a century ago [15]. Control work on such a project therefore has to record, for every number the design rests on, whether it was calculated, measured on a rig, or measured on an operating plant.

7.8 Technology in Practice

The tags of Figure 7.1 are, in a plant, a thermocouple on the hot leg, a flow measurement on the cold leg, a pump motor on a speed drive and the controllers of a plant control system. Figure 7.6 draws the loop with that hardware in place of the schematic's symbols.

What Measures It

The most common temperature sensors in nuclear plants are thermocouples and resistance temperature detectors, field devices that sit in contact with what they measure or in thermowells that protrude into the fluid [23]. In a salt loop four things decide how long such a sensor lasts: the long exposure to high temperature, the corrosive salt, the transmutation of the sensor's own materials and the vibration of a thermowell in fast flow, while gamma radiation matters little [24]. Every thermowell adds to the first-order lag of the sensor through its mass and specific heat, and a lag that is large enough can compromise the stability of feedback control [25]. The chapter carries the outlet thermocouple and its well as one 2 s filter with a bias and noise, and the flow transmitter as a 0.5 s filter.

> **TECHNOLOGY Sensor: A Sheathed Thermocouple in a Thermowell**
>
> In the Molten Salt Reactor Experiment, thermocouples near the entrance of the reactor outlet line sensed the outlet temperature for the alarm and control circuits [9]. Above 600 °C most resistance detectors are out of their practical range, and the thermocouple is the best overall choice. Type N thermocouples drift little up to 1300 °C. The bead cannot touch the salt, which conducts electricity and would wick up the wires and short the measurement, so it sits in a sheath inside a thermowell [24]. The most common failure is an open circuit, which an instrument without a burnout mode reads as a wildly varying temperature [25]. The chapter's TT is a 2 s filter with noise and a bias.

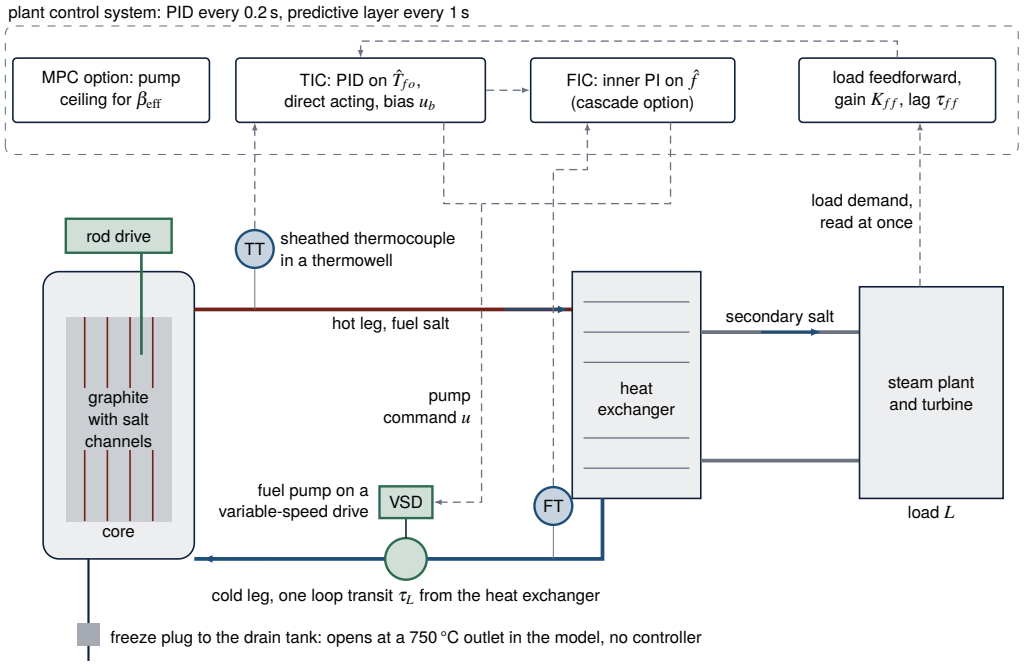


Figure 7.6 The temperature loop of the chapter with plant hardware in place of the symbols of Figure 7.1. Sensors are blue circles and final elements green: the core outlet thermocouple TT in its thermowell on the hot leg, the flow transmitter FT on the cold leg, the fuel pump on its variable-speed drive (VSD) and the rod drive, which is an operator trim. TIC acts on the pump directly, or sets the flow setpoint of the inner FIC in the cascade, with the load demand fed forward through a lag. The MPC option replaces both and caps the pump command to keep β_{eff} above its floor. Notice that the same cold leg carries the cooled salt and the returning precursors back to the core one loop transit later.

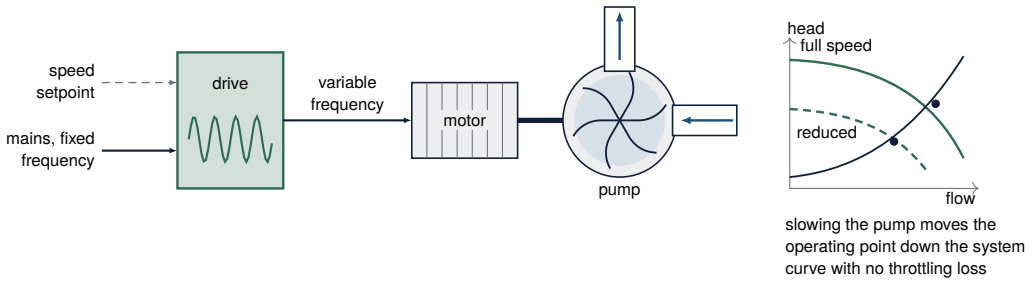


Figure 7.7 A variable-speed drive on a centrifugal pump, from the book's instrument library. Notice that the drive changes the pump curve, not the system curve, so the flow moves along the system curve as the speed changes. In this chapter the same speed also sets the time the fuel salt spends inside and outside the core.

What Moves It

The final element is the fuel pump, and in the chapter its command is the only one a controller writes. The rod is an operator trim rate limited at 0.5 %/s, and the secondary pump follows an operator command. The pump model is an idealization, with no dead band and no rate limit beyond its 4 s lag. Figure 7.7 shows how a speed drive moves a centrifugal pump's flow.

TECHNOLOGY Final Element: The Fuel Salt Pump on a Variable-Speed Drive

The fuel pump of the Molten Salt Reactor Experiment was a vertical centrifugal pump of the sump type, with an overhung impeller and a 75-horsepower motor. About 65 gallons per minute of its output was returned inside the pump bowl, part of it as a spray, to release xenon and krypton from the salt. Coastdown tests found that the motor took about 10 seconds to stop once its supply was lost [9]. To vary such a pump's flow, a variable-frequency drive sets the supply frequency of the motor and keeps the voltage in proportion, giving a continuous range of speeds rather than steps [26]. The chapter's pump follows its command through a 4 s lag and coasts down with a 10 s time constant on a trip.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, and nested loops in which the setpoint of one loop comes from the measurement of another [27]. The chapter's PID executes every 0.2 s over a 0.02 s integration step, and the predictive option plans every second. The structure is a set of function blocks: a direct-acting PID on the filtered outlet temperature, a first-order lag on the load demand added to its bias, an optional inner PI on the flow transmitter, and output clamps at 20 and 120 % of design pump speed.

IN PRACTICE Load Changes Planned Around Thermal Fatigue

A molten salt reactor follows load inherently, and its instruments and controls serve to smooth each transition and reduce the stress on the components. Pump speed, absorbers or reflectors are expected to act in an anticipatory scheme, and a key goal is to keep the temperature swings on the hot structural metal small, because fatigue from those swings is one of the main limits on its life. Measuring the fuel salt flow itself is hard and costly, so the plant is expected to watch the fuel pump speed as an approximate flow and to make its heat balance on the coolant side of the primary heat exchanger [24]. The lagged load feedforward of this chapter is such an anticipatory term.

What Is Modernizing It

Three developments reach this loop. Johnson noise thermometry reads the same thermocouple signal in parallel with the ordinary electronics and gives a slow but low-uncertainty temperature that recalibrates it, which virtually eliminates drift. It is not yet a commercial product, it is very sensitive to electromagnetic interference, and regulators have not yet accepted its self-calibration. Distributed fiber-optic Bragg thermometry places many temperature points along a single fiber, although gamma radiation can darken the fiber [24]. Operating plants are moving from analog to digital control through individual system upgrades [23], and control systems that adopt office computers, operating systems and network protocols have far less isolation from the outside world and a greater need to be secured [27].

7.9 Limitations

Nothing in this chapter is a reactor analysis tool. With no spatial flux, no burnup, constant reactivity coefficients, no protection system, no secondary plant and a natural circulation floor that is one fixed number, the application cannot support any claim about a real core's behavior, its safety margins, or what a particular design would do on a pump trip. Section 7.3.6 names the place each individual simplification fails.

The control scope is deliberately narrow. There is one temperature loop with one feedforward, one cascade and a separate nominal MPC, no reactivity control loop on the rod, and no classical gain scheduling on pump speed even though Section 7.4 measures the process gain as a function of it. A real plant would coordinate the rod, the fuel pump and the secondary flow together, and the rod here is a manual trim. Nothing trips. The high flux limit and the deviation alarm annunciate and the simulation continues, so a transient stays readable past the point where a plant would have stopped. The instruments are a lag plus noise plus a bias, whereas a real thermocouple in a salt thermowell has

a nonlinear calibration, drifts, and fails in ways that do not look like a constant offset. The perfect repeatability of a fixed step and a fixed control period is a teaching convenience rather than a property of a control system.

The reduced MPC omits dynamic precursor inventories, graphite storage and explicit delivery-gain estimation. Its output bias cannot distinguish sensor bias from plant mismatch: a +3 °C thermocouple error gives measured error 0.000 441 °C but true error −2.999 545 °C after 900 s. Its margin is nominal, with no actual-margin guarantee during pump lag or a delivery fault. The ceiling is a control policy and does not establish reactor protection.

What was kept, and kept because a controller has to be designed around it, is the effective delayed fraction that moves with the manipulated variable, the transport delay shared between the neutronics and the thermal loop, the criticality condition that pins the mean, and the saturating heat sink.

OPEN CHALLENGE **Measuring Flow in a Molten Salt**

A 2018 survey of instruments for molten salt reactors found no gap that would prevent operating one, but named flow measurement as the area with the largest hurdles left. Penetrations and internal parts are unwelcome in salt piping, because a part that breaks loose is carried downstream, and joined parts must survive a range of about 700 °C. A clamp-on ultrasonic meter needs waveguides, because its transducers stop at 200 °C on piping at 700 °C, and the waveguides add up to several percent of error and can cool the salt toward freezing. Progress would be transducers that work at the pipe temperature without a waveguide, tested for their response to radiation [24].

7.10 Conclusion

This chapter studied a circulating-fuel reactor as an interactive laboratory for the class of plants whose manipulated variable changes the plant. The simulator integrates twelve states, six delayed neutron precursor groups carrying a delayed return term, the fuel salt and graphite temperatures, the heat exchanger primary temperature, the rod position and two pump speeds, with the external loop represented as a transport delay shared by the precursors and the heat. At the design pump speed the effective delayed neutron fraction is 419 pcm against the 650 pcm the fuel produces, so 35.5 % of the delayed neutrons decay outside the core and the rod holds 231 pcm of permanent excess reactivity to pay for pumping. The manipulated variable is the fuel pump rather than the heat sink, chosen from measurements that give the secondary flow pairing a gain of +0.186 °C per percent and 16 s of pure dead time against the fuel pump's −0.757 °C per percent and no dead time at all. A proportional controller settles 6.2 °C low because no element downstream of it integrates, reset removes that offset for 22.3 % of pump travel, and the gain limit is set by transport rather than by plant gain. A fuel pump trip at low power returns 231 pcm of uncommanded reactivity and takes the power from 9.0 to 16.8 %. A fuller model would add a spatial precursor distribution, a coordinated rod and pump strategy, a scheduled gain, and a secondary plant with its own dynamics. The plant remains a compact case of holding a temperature with an actuator that is also a safety margin.

The MPC comparison makes the margin policy explicit, with zero rather than 476.76 s below 415 pcm on the nominal margin case at the cost of 2324.30 °C s rather than 299.24 °C s of temperature error.

KEY POINTS

- Pumping the fuel salt carries some late neutrons out of the core, so the pump that holds the outlet temperature also sets the stability margin.
- The secondary flow pairing has the wrong sign and 16 s of dead time, and the fuel pump has neither.
- A proportional controller settles 6.2 °C low, and the gain limit is set by transport.
- A fuel pump trip at low power takes the power from 9.0 to 16.8 % with no rod motion.

Where else this applies. Volume 5, Chapter 2 also controls through a transport delay.

References

[1] 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.

[2] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[3] J. Lewins, "Neutron and precursor equations," in *Nuclear Reactor Kinetics and Control*. Elsevier, 1978, pp. 30–55.

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

[5] J. Serp, M. Allibert, O. Beneš, S. Delpech, O. Feynberg, V. Ghetta, D. Heuer, D. Holcomb, V. Ignatiev, J. L. Kloosterman, L. Luzzi, E. Merle-Lucotte, J. Uhlíř, R. Yoshioka, and D. Zhimin, "The molten salt reactor (MSR) in generation IV: Overview and perspectives," *Progress in Nuclear Energy*, vol. 77, pp. 308–319, 2014.

[6] D. LeBlanc, "Molten salt reactors: A new beginning for an old idea," *Nuclear Engineering and Design*, vol. 240, no. 6, pp. 1644–1656, 2010.

[7] E. S. Bettis, W. B. Cottrell, E. R. Mann, J. L. Meem, and G. D. Whitman, "The aircraft reactor experiment—operation," *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 841–853, 1957.

[8] W. B. Cottrell, H. E. Hungerford, J. K. Leslie, and J. L. Meem, "Operation of the aircraft reactor experiment," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL-1845, 1955.

[9] R. Robertson, "MSRE design and operations report part 1: Description of reactor design," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep., 1965.

[10] P. N. Haubenreich and J. R. Engel, "Experience with the Molten-Salt Reactor Experiment," *Nuclear Applications and Technology*, vol. 8, no. 2, pp. 118–136, 1970.

[11] B. Prince, "Period measurements on the Molten Salt Reactor Experiment during fuel circulation: Theory and experiment," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep., 1966.

[12] B. Prince, S. Ball, J. Engel, P. Haubenreich, and T. Kerlin, "Zero-power physics experiments on the Molten-Salt Reactor Experiment," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep., 1968.

[13] T. W. Kerlin, S. J. Ball, and R. C. Steffy, "Theoretical dynamics analysis of the Molten-Salt Reactor Experiment," *Nuclear Technology*, vol. 10, no. 2, pp. 118–132, 1971.

[14] T. W. Kerlin, S. J. Ball, R. C. Steffy, and M. R. Buckner, "Experiences with dynamic testing methods at the Molten-Salt Reactor Experiment," *Nuclear Technology*, vol. 10, no. 2, pp. 103–117, 1971.

[15] J. Engel and R. Steffy, "Xenon behavior in the Molten Salt Reactor Experiment," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep., 1971.

[16] S. J. Ball, "Simulators for training Molten-Salt Reactor Experiment operators," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep., 1966.

[17] W. Zeng, Q. Jiang, J. Xie, and T. Yu, "Core outlet fuel temperature control of liquid molten salt reactor during load following operation," *Progress in Nuclear Energy*, vol. 121, p. 103214, 2020.

[18] K. S. Kupriyanov, O. S. Feinberg, and V. V. Ignatiev, "Effective delayed neutron fraction in a molten salt reactor with circulating fuel," *Physics of Atomic Nuclei*, vol. 85, no. 8, pp. 1391–1399, 2022.

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

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

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

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

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

[24] D. E. Holcomb, R. A. Kisner, and S. M. Cetiner, "Instrumentation framework for molten salt reactors," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2018/868, 2018.

[25] 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.

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

[27] 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.

7.11 Exercises

7.1 Concept check · 2 min

Which statement describes the loop that holds the temperature of this plant?

(a) A control rod loop holds the mean fuel salt temperature, and the fuel pump stays at design speed

(b) The secondary salt flow holds the core outlet, because it sets how much heat leaves the primary

(c) The fuel pump holds reactor power at rated against every change in turbine load

(d) The fuel pump holds the core outlet temperature, and its speed also moves β_{eff}

7.2 Concept check · 2 min

In Equation (7.2), what does the fourth term, $W_i(t-\tau_L)\,\mathrm{e}^{-\lambda_i\tau_L}/\tau_c$, represent?

(a) Precursors of group i that left the core one loop transit ago and survived the trip back

(b) Precursors of group i that decay inside the core at the present instant

(c) Precursors of group i that the flow sweeps out of the core at the rate $1/\tau_c$

(d) Precursors of group i born by fission in the heat exchanger and the pipes

7.3 Concept check · 2 min

At the design point, with the fuel pump held at 100 %, the secondary salt flow is stepped from 100 to 125 %. What does the core outlet temperature do?

(a) It falls at once, because more heat is taken out of the primary salt

(b) It holds for 16 s, then rises monotonically to a hotter steady value

(c) It dips first and then rises, the inverse response of a heat exchanger

(d) It falls after 16 s of dead time, when the cooler salt reaches the core

7.4 Theory and calculation · 5 min

Use Table 7.3 at the design point, with $n = 1$ and the fuel pump at 100 %.

(a) Compute the flow capacity rate W_f and the core temperature rise $P_0 n/W_f$. *Checkpoint:* 2.5 MW/°C and 100 °C.

(b) Compute the time constants $C_f/(2W_f)$ of the fuel salt node and C_g/h_{fg} of the graphite. *Checkpoint:* 4.5 s and 277 s.

(c) After a pump trip at full power with the flow gone, the salt heats at $(1-f_G)P_0/C_f$. Compute that rate, the reactivity it removes each second through α_f, and the time this takes to remove the 230.8 pcm that the rod holds. Compare it with the 10 s coastdown. *Checkpoint:* 10.4 °C/s, 52 pcm per second and 4.4 s.

7.5 Theory and calculation · 6 min

Use the delayed neutron data of Table 7.2 and Equations (7.3) and (7.4).

(a) At the design flow ($\tau_c = 9\,\text{s}$, $\tau_L = 16\,\text{s}$), compute the share of its worth that group 1 ($\lambda_1 = 0.0124/\text{s}$) keeps and the share that group 6 ($\lambda_6 = 3.01/\text{s}$) keeps. *Checkpoint:* 38.3 % and 96.4 %.

(b) At 120 % pump speed, give τ_c and τ_L. With $\beta_{\text{eff}} = 404.2$ pcm there, compute the reactivity the rod must hold, the share of the delayed neutrons lost and the worth in dollars of a −100 pcm rod drop. *Checkpoint:* 7.5 s and 13.3 s, then 246 pcm, 37.8 %, 0.247 dollars.

(c) At 40 % the same drop is 0.203 dollars. By what factor is it worth more at 120 %? *Checkpoint:* 1.22.

7.6 Theory and calculation · 6 min

Scenario 1 shows that with the pump parked at 100 % the outlet reaches 655.0 °C after the cut to 70 % load, and scenario 2 shows that 85.2 % puts it back on 665 °C. Treat the outlet as linear in pump speed between those two points at this load. The proportional-only controller of scenario 3 has $K_c = 1.0\,\%/°\text{C}$, a bias of 100 % and the direct action of Section 7.5.1.

(a) Compute the secant process gain in °C per percent of pump. *Checkpoint:* −0.676 °C per percent.

(b) Predict the steady outlet offset and the pump speed, and compare them with the 6.20 °C and 93.8 % that the run of Table 7.6 settles at. *Checkpoint:* 5.97 °C low, pump at 94.0 %.

(c) Find the gain that would bring the predicted offset down to 1 °C, and compare it with the gain at which the loop becomes unstable. *Checkpoint:* $K_c = 13.3$, about six times the 2.2 %/°C at which the loop diverges.

7.7 App challenge · 8 min

Load scenario 1, The plant with nobody driving (MANUAL), and change no setting. Run to 600 s at the default seed at 20× playback, then export the CSV file (its time column is in minutes, so 600 s is the row at 10 min).

(a) Read the reactor power and the fuel mean, core out and core in temperatures at 600 s (columns `power_pct`, `tfuel_C`, `outlet_C` and `inlet_C`), to 0.1 °C and 0.1 %. *Checkpoint:* 78.2 %, then 615.8 °C, 655.0 °C, 576.6 °C.

(b) Compute how far each temperature moved from its design value (615, 665 and 565 °C) and the core rise. *Checkpoint:* +0.8, −10.0 and +11.6 °C, with a rise of 78.4 °C.

(c) Explain with Equation (7.5) why the mean barely moved, and compare the rise with $P_0 n / W_f$ at this power. *Checkpoint:* 78.2 °C from the power, 0.2 °C below the rise read.

7.8 App challenge · 8 min

Load scenario 4, Reset added, on a load change, at the default seed and 20× playback. Every run is 600 s and every setting is made before Run. Read the IAE and MV travel on the scorecard, the Pump travel and β_{eff} tiles, and the peak deviation as the last value of the CSV column `outlet_max_error_C`.

(a) Run with the shipped tuning, $K_c = 1.2$ and $\tau_I = 20$ s. *Checkpoint:* IAE 261.9 °C s (within 0.5), travel 22.3 %, peak 4.18 °C.

(b) Reset, set Proportional gain to 2.0 and run again. *Checkpoint:* IAE 747.5 °C s (within 1), travel 407 %, peak 7.77 °C.

(c) In Show Model, Parameters, set External loop transit time at design flow to 8 s, press Reset, keep $K_c = 2.0$, read the β_{eff} tile before Run, and run again. Explain both changes. *Checkpoint:* β_{eff} 458 pcm and travel 45 %.

7.9 App challenge · 8 min

Load scenario 7, Load feedforward, and when it should arrive, at the default seed and 20× playback. Make three 600 s runs, each set before Run: as shipped ($K_{ff} = 0.5$, $\tau_{ff} = 20$ s), then Reset with Feedforward lag 0, then Reset with Control architecture Single. Read the scorecard IAE and the last value of the CSV column `outlet_max_error_C` after each run.

(a) Report the three peaks and the three IAE values, to 0.05 °C and 1 °C s. *Checkpoint:* 2.64, 7.42 and 7.04 °C, then 188, 329 and 416 °C s.

(b) Compute how much the lagged feedforward takes off the peak and off the IAE of feedback alone. *Checkpoint:* 63 % off the peak and 55 % off the IAE.

(c) Explain why the static feedforward has the highest peak of the three. *Checkpoint:* It moves the pump before the disturbance has reached the core.

7.10 Application area · 4 min

The freeze plug drains the fuel into a tank that cannot go critical when the core outlet passes 750 °C. In the full-power pump trip of Section 7.6 it does so 19.8 s after the trip, with no signal, power supply or operator.

(a) Give one reason this is a strong engineering argument and one reason it is an awkward licensing argument. *Checkpoint:* No logic or actuator to fail, and no test interval, redundancy or failure rate to credit.

(b) Name two things the model's single threshold leaves out that a real safety case would have to demonstrate. *Checkpoint:* Melt timing of the plug and a drain that takes minutes with decay heat in the tank.

7.11 Application area · 4 min

The graphite temperature coefficient of a graphite-moderated reactor drifts with fast neutron damage over a few years. Use Table 7.9, with α_f fixed at −5 pcm/°C.

(a) At what value of α_g does the plant lose its temperature stiffness, and what does a parked plant do at +4 pcm/°C? *Checkpoint:* +5 pcm/°C, and at +4 it settles 128.6 °C low at 14.3 % power.

(b) What should an operating plant track, and what does the AUTO loop do as the coefficient rises? *Checkpoint:* The coefficient sum. The loop absorbs +2 at nearly the same cost and fails at +6.

7.12 Review the agent · 5 min

An AI assistant is asked which fuel pump speed bounds the rod-drop analysis and replies as follows. "Run the −100 pcm rod drop at 40 % pump speed. A slower pump sweeps fewer precursors back into the core, so β_{eff} is lowest at low flow, and the same rod is then worth the most dollars. The fast pump case is milder and need not be analyzed." Find the flaw and say how you would check it.