INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 10

Decay-Heat Removal Simulator for Constraint Control of a Depleting Heat Sink

Pages 261–289

Edited by John D. Hedengren

How to cite this chapter. "Decay-Heat Removal Simulator for Constraint Control of a Depleting Heat Sink," 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. 10, pp. 261–289.

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

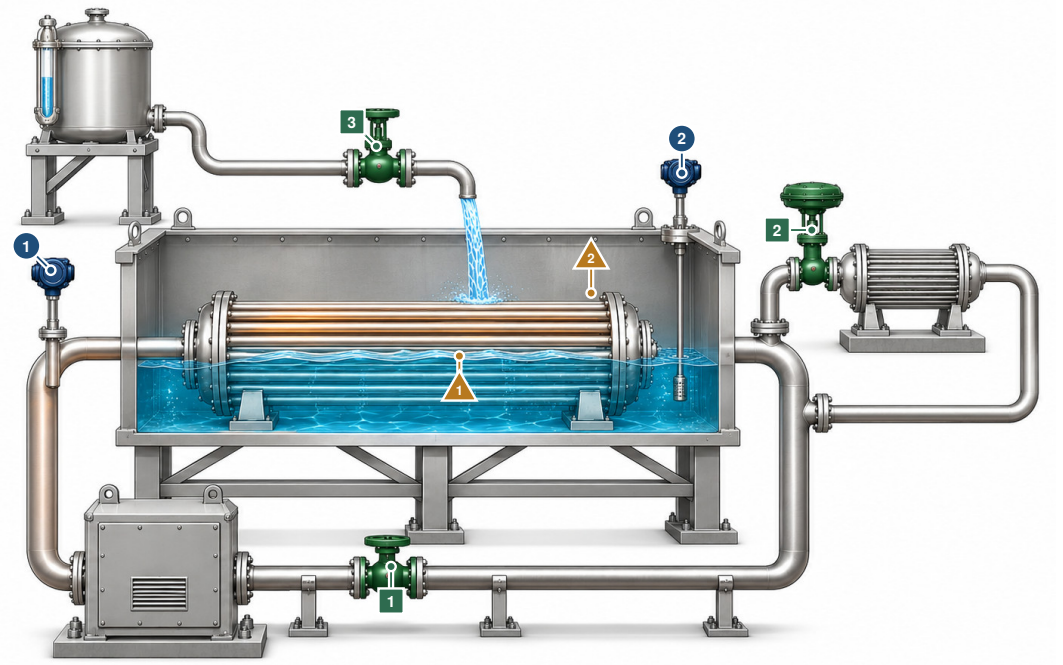
The app for this chapter
https://apmonitor.com/apps/decay-heat-removal/

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

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

CHAPTER 10

Decay-Heat Removal Simulator for Constraint Control of a Depleting Heat Sink



The controller moves 1 isolation valve on the cooled water returning to the reactor
2 auxiliary cooling train, which needs a diesel 3 makeup from a tank into the pool, opened by hand

The controller reads 1 core outlet temperature on the hot leg 2 pool level

What disturbs it 1 the pool boiling down until the tubes uncover 2 gas collecting in the condenser tubes

Decay heat boils the pool away, and once the level falls below the top of the tube bundle the condenser loses its grip on the reactor temperature. Illustration, not to scale.

IN PLAIN TERMS Rationing Water to Cool a Shut-Down Reactor

The main idea. A reactor that has been shut down keeps making heat for days, and nothing can switch that heat off. Here the heat goes into a pool of water that boils away, and the pool is the only place the heat can go.

The challenge. Cooling the reactor as hard as possible feels safe. But it boils off water the plant will need hours later. Keeping the reactor hotter saves water, yet it leaves less room below the temperature limit.

Why it matters. The same trade between a safety margin and a supply that runs out appears wherever a stored resource has to last for a set number of hours.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 10.2, the fact boxes of Sections 10.3 and 10.6, the try-it boxes of Section 10.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 10.6, Section 10.7, Section 10.8 and the open challenge

The control question A tripped reactor makes decay heat that no handle can reduce, and its only heat sink is a pool that boils away. With no setpoint to hold, how should an isolation valve, an auxiliary train and a finite makeup tank be used to keep the core outlet under its limit for a stated number of hours?

Learning objectives

- Compute the decay heat after a stated operating history and the pool water it boils.
- Derive the cube-root law of natural circulation, and explain why the valve barely moves the flow while its process gain varies by a factor of twenty-five.
- Predict from the stored-heat ledger what a cooldown costs in pool water before running a scenario.
- Decide from the two setpoints whether a margin override on the auxiliary train will run.
- Choose a structure, a setpoint pair and gains that meet a coping-time target, and judge them on the water left as well as on the peak.

Prerequisites Sections 1.3 (loop action), 1.4.5 (decay heat, coping time, the exchange rate of stored heat), 1.6.2 (velocity form, anti-windup) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/decay-heat-removal/

See also Volume 4, Chapter 2 for a protective override that rests on its lower clamp and is judged by its minimum margin. Volume 4, Chapter 14 for a reserve vault, Volume 4, Chapter 8 for a cash tank and Volume 4, Chapter 12 for a nitrogen tank, each spent against a hard limit.

10.1 Introduction

Almost every control problem taught in an introductory course has a setpoint. A furnace has an outlet temperature to hold, a column has a composition to hold, and a tank has a level to hold, and the task is to make the measurement equal the target with as little effort and as little overshoot as possible. The plant of this chapter has no target of any kind. A light-water reactor that has just tripped is making heat at a rate nobody can influence, its core outlet temperature will go wherever the heat balance puts it, and the operator has three handles and three things that must not happen. There is a temperature above which the vessel lifts its relief valves, a pool level below which the only heat sink stops working, and a number of hours the plant is required to survive without outside help. Nothing in that list is a setpoint, and the strategy that keeps the temperature most comfortable is measurably the one that runs out of water first.

The misconception the simulator corrects is the belief that the safe action is always to cool down as fast as possible. Running the isolation condenser wide open takes the core outlet from 300 °C to 212 °C in four hours and grows the temperature margin the whole way, which is exactly what a passive plant is designed to do. It also spends 14.6 t of a 55.0 t pool on a cooldown that no constraint required, because the 372.8 MJ/K of stored heat in the fuel, the water and the vessel steel has to be carried away by boiling water at 2257 kJ/kg. Holding the plant hot instead leaves 29.9 t in the pool at the same four hours and pushes the hour at which the tube bundle uncovers from 3.93 to 6.26. A block diagram cannot show that, because the quantity being traded does not appear in it. The error signal, the process gain and the actuator position are all identical in the two cases at the moment the decision is made, and what separates them is an integral of a flow over a horizon that has not happened yet.

The general class is control of a process with no setpoint, a hard constraint, and a consumable resource that has to be rationed over an unknown duration, where the aggressive answer and the conservative answer are both wrong. A battery on an islanded microgrid (Volume 17, Chapter 7) is the same problem, and the dispatch literature treats the state of charge as a resource to be scheduled against an uncertain restoration time rather than as a variable to be regulated [1]. So is a limited reagent or coolant inventory during a batch upset, a nitrogen supply on a purge system, and the oxygen in a closed environment. In each of them the constraint is one-sided, the resource integrates, and the correct trajectory depends on a horizon that is a judgment rather than a measurement. Control of a constraint rather than of a setpoint has a settled vocabulary in the process industries [2], and constraint

handling is the feature that carried model predictive control into practice in the first place [3]. This chapter works on a plant small enough to run in a browser, where the resource is visible, the constraint is hard, and the horizon can be changed by one slider.

🕮 **Pause.** Try Exercise 10.1 before reading on.

The physics the model needs is the physics an introductory course already supplies. Energy is conserved across five lumped capacities with a temperature difference driving every transfer term, mass is conserved in a pool that boils, and the circulation follows from a momentum balance between a buoyancy head and a quadratic friction loss. The consequence of that last balance is worth stating in advance, because it is the reason a passive plant is stable at all. Substituting the loop temperature rise into the momentum balance gives a flow proportional to the cube root of the power, so more heat raises the temperature difference, which raises the head, which raises the flow. It is a feedback loop nobody tuned and nobody can destabilize, and the cube root also makes it almost entirely without authority.

That this arrangement works is what makes the control question interesting. This chapter uses an interactive simulation of decay-heat removal by natural circulation into a boiling pool to separate what a passive system does on its own from the inventory decision it leaves to an operator. Two measured results carry the argument. The isolation valve has a factor of twenty-five in process gain across its travel and almost no authority over the flow it appears to control. The plant's stored heat trades against pool water at 165 kg per kelvin of avoided cooldown, and that exchange rate decides every architecture comparison in the chapter. The decay-heat correlation, the loop momentum balance and the controller arithmetic are all standard. Section 10.2 places decay heat and passive heat removal in their history. Section 10.3 sets out the fifteen states, the constants and the instrument chain in enough detail to rebuild the plant. Section 10.4 measures the loop and the valve at steady state and derives the stored-heat ledger. Section 10.5 states the four heat-removal architectures and the two proportional-plus-integral loops of opposite action. Section 10.6 runs the twelve scenarios, the last of them a graded challenge, and Section 10.7 sets the loop in its plant hardware. Sections 10.8, 10.9 and 10.10 give the limitations, a discussion and the conclusion, and the exercises follow in Section 10.11. The lesson throughout is that a limit and a horizon are not a setpoint.

10.2 Decay Heat and the Passive Answer to It

The residual hazard of a shut-down reactor was recognized while the first production reactors were being designed, and it was recognized as an arithmetic problem rather than a nuclear one. Way and Wigner reduced the aggregate beta and gamma emission of a mixed fission-product inventory to a two-term power law in the time since shutdown and the time at power, and that expression, evaluated with an assumed operating history, is still the shape every later correlation refines [4]. The exponent is the whole of the difficulty. A power law with an exponent near one fifth falls by a factor of twenty-five over six decades of time, so a plant that is making tens of megawatts one second after the trip is still making megawatts a day later, and no amount of waiting removes the requirement for a heat sink. Later measurement campaigns established the fission-product energy release directly rather than by summation over decay chains, and reviews of that work collected the systematic differences between fuels, enrichments and irradiation histories [5]. The result is now a consensus standard with tabulated uncertainties, and a licensing calculation is required to use it rather than a first-principles summation [6].

Removing that heat needs three things at once, which are something to move the coolant, something to take the heat out of the coolant, and somewhere to put it. The first generation of power plants answered all three actively, with residual heat removal pumps, a shutdown cooling heat exchanger and the ultimate heat sink, and the answer therefore depended on alternating current electrical power. Buoyancy-driven circulation was studied from the beginning as the backup for the first of the three. Reviews of natural circulation loops collected the steady-state scaling, the transition to the turbulent regime and the several instability mechanisms that a loop can exhibit [7], dimensionless groups were

established that collapse the steady flow of geometrically dissimilar loops onto one curve [8], and experiments on loops of widely different diameter and orientation confirmed the same scaling and mapped the stability boundaries [9]. The reactor-specific version of that body of work, including the phenomena a lumped model cannot reproduce, has been collected by the International Atomic Energy Agency [10], and the loop momentum balance and its cube-root consequence are standard textbook material [11].

The move from active to passive residual heat removal happened for the other two of the three requirements, and it happened in the design generation that followed Three Mile Island. The distinguishing property is not the absence of moving parts but the absence of a dependence on external power and on operator action within a stated period, and the IAEA classification of passive systems by degree of passivity was written to make that distinction precise [12]. Two families of hardware appeared. A submerged tube bundle carrying primary coolant and condensing into a pool open to atmosphere is the isolation condenser, present on the earliest direct-cycle boiling-water reactors, dropped from the intermediate designs, and returned in the current ones. A heat exchanger hung in a large in-containment water tank is the passive residual heat removal exchanger of the AP1000, which is credited for the same function on a pressurized plant [13]. Both were qualified by integral experiment rather than by analysis alone, and the experiments established what degrades them, which is principally the accumulation of non-condensable gas against the condensing surface [14].

> ⧖ HISTORY **A Condenser Tube That Cracked (1976)**
>
> In February 1976 a boiling water reactor in the United States scrammed from full power on a turbine generator trip, and the pressure spike as the turbine control valves closed ruptured a tube in its isolation condenser. The stainless steel tube showed chloride-assisted stress corrosion cracking that had started inside it, probably from seawater injected into the condenser years before, and the condenser was retubed with a nickel alloy. The regulator's bulletin that followed required plants to show that their condenser tubes stayed sound in service: leak detection that watches the shell-side water temperature, volume and isotopic content, and periodic nondestructive examination, or a hydrostatic test where that cannot be done [15].

What the passive answer changed was the shape of the safety argument. An active system is graded on availability, because a running pump removes heat indefinitely. A passive system is graded on a coping time, because a pool of water is finite and the argument is that the plant survives a stated number of hours on the inventory that is already inside the containment boundary [12]. That single change is what turns the subject of this chapter into a control problem. A coping time is computed against an assumed operating history, an assumed set of degradations and an assumed set of operator actions, and the operator actions are the part that a simulator can put in a student's hands.

The Fukushima Daiichi accident supplied the case study, and reconstructions treat the operation of the isolation condensers of its Unit 1 as one of the determining factors in how the transient developed [16, 17], as the history box below tells. The app of this chapter does not attempt to reproduce that event. It reproduces the decision, which is whether to throttle a passive heat sink that is working, and it prices the decision in kilograms.

HISTORY **An Isolation Condenser Stopped Before the Tsunami (2011)**

On 11 March 2011 Unit 1 at Fukushima Daiichi, the only unit of the site built with isolation condensers, was cut off from its turbine, and both condenser loops started on high reactor pressure. They cooled the reactor so fast that the operators stopped them, as the procedures required, to spare the vessel thermal stress. The operators then held the cooling rate by opening and closing the isolation valves on the condensate return lines. The valves had just been closed when the tsunami arrived and the power was lost. With no indication of valve position, the first attempt to restart the condenser came about three hours later, and the condenser did not work. Core cooling at Unit 1 had been lost from the moment it was stopped [17].

10.3 Plant Model

The plant is a 530 MW thermal light-water reactor at the instant the trip breakers open. Coolant at 270.0 °C enters the core, leaves it at 300.0 °C, rises through the riser and returns down the downcomer annulus, and 3330 kg/s of pumped flow carries exactly the 530 MW the core is making into a steam generator that is about to be isolated. The average fuel temperature is 685.0 °C. Above the vessel a tube bundle hangs in 55.0 t of water at 32.0 °C, open to atmosphere, with its condensate return throttled by an isolation valve. An auxiliary cooling train rated 2.0 MW can be lined up on the downcomer if its diesel starts, and a 12.0 t makeup tank can be spent into the pool once, by hand. Figure 10.1 draws the circuit with every symbol the equations below use, and marks the two thermal centers whose height difference Δz is the lever arm the buoyancy head works on. Working units are minutes for the clock, MW for power, MJ/K for heat capacity, °C, kg, kg/s, Pa and kW/K.

DID YOU KNOW? **Eight Hours of Pool Water**

The isolation condenser pools of the Fukushima Daiichi unit in the history box held enough water, both trains together, for eight hours of cooling before they needed refilling [17]. The coping time required of this chapter's plant is also eight hours (Table 10.2).

Table 10.1 classifies what the loops interact with. Fifteen states are integrated. Five annunciations bound the transient, which are the relief setpoint at 316 °C, the start of tube uncovery (a submerged fraction below 0.98), a pool below 5 % of its full inventory, a cooldown rate above the technical-specification limit of 55 K/h, and an average fuel temperature above 1200 °C. None of the five stops the run, because a transient is more legible past the point at which a real plant would have changed character than it is if the trace ends there.

10.3.1 Decay Heat and the Collapse of Fission

Decay power is the two-term power law of Equation (1.9) in the time since shutdown t_s and the time at power T_0, written here as Equation (10.1).

$$\frac{Q_d}{P_0} = A\left(t_s^{-n} - (t_s + T_0)^{-n}\right), \qquad A = 0.066,\ n = 0.2 \tag{10.1}$$

where t_s and T_0 are in seconds, P_0 is rated thermal power, and T_0 is floored at 1 h. Both t_s and T_0 are exposed, the first as an initial condition so that a scenario can join the transient already settled, and the second as a disturbance slider. Prompt and delayed fission are carried as two decaying states rather than as full point kinetics, as shown in Equation (10.2).

$$\begin{aligned} \tau_P \frac{\mathrm{d}n_P}{\mathrm{d}t} &= -n_P, \qquad \tau_D \frac{\mathrm{d}n_D}{\mathrm{d}t} = -n_D \\ Q_{fiss} &= P_0\,(n_P + n_D) \end{aligned} \tag{10.2}$$

where $\tau_P = 2.4$ s collapses the prompt flux and $\tau_D = 75$ s carries the longest delayed-neutron precursor. The two are initialized to $(1 - f_{d0})(1 - \delta)$ and $(1 - f_{d0})\delta$ of rated, where f_{d0} is the decay fraction

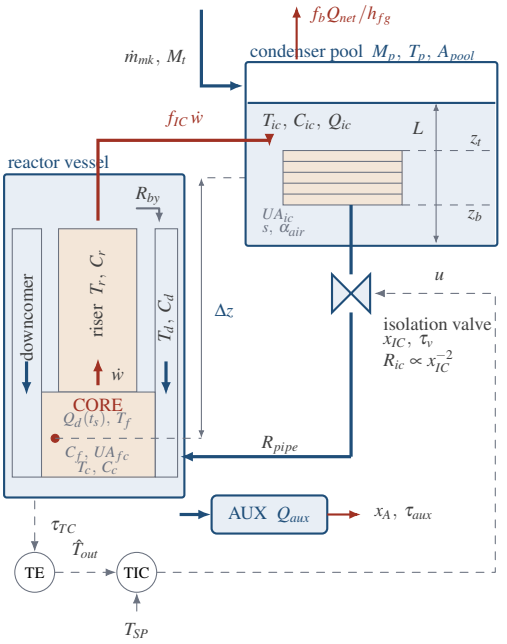


Figure 10.1 Natural-circulation loop, isolation condenser and condenser pool of the modeled plant. Decay heat Q_d raises the fuel node T_f, which passes heat to the core coolant node T_c through UA_{fc}, and the core outlet $T_{out} = 2T_c - T_d$ feeds the riser. A fraction f_{IC} of the circulation $\dot{w}$ passes through the condenser tube side T_{ic} and the rest returns by the internal bypass R_{by}. The condenser rejects Q_{ic} to the pool through conductance UA_{ic}, which the isolation valve sets by flooding the lower tubes and which falls with the submerged fraction s once the level L passes the top of the bundle z_t. The pool boils at $f_b Q_{net}/h_{fg}$, and $\dot{m}_{mk}$ can be spent into it from tank M_t. Marked dots are the two thermal centers whose height difference Δz drives the flow. TIC reads the core outlet through thermocouple lag τ_{TC} and drives the isolation valve.

of (10.1) at $t_s(0)$ and $\delta = 0.010$, so that the core is making exactly P_0 at $t = 0$. Total core power is $Q_{core} = Q_d + Q_{fiss}$, and the split matters only for the first minute. At 30 s the measured core power is 19.8 MW, of which 16.5 MW is decay heat and 3.3 MW is the delayed-neutron tail.

10.3.2 The Loop, and Why the Core Node Carries a Factor of Two

Five thermal capacities are integrated, which are the fuel, the core coolant, the riser, the condenser tube side and the downcomer, as shown in Equations (10.3) to (10.7). The factor 60 converts a rate per second into a rate per minute.

$$C_f \frac{dT_f}{dt} = 60\left[Q_{core} - UA_{fc}(T_f - T_c)\right] \tag{10.3}$$

$$C_c \frac{dT_c}{dt} = 60\left[UA_{fc}(T_f - T_c) + 2\dot{w}c_p(T_d - T_c)\right] \tag{10.4}$$

$$C_r \frac{dT_r}{dt} = 60\left[\dot{w}c_p(T_{out} - T_r) - Q_{main}\right] \tag{10.5}$$

$$C_{ic} \frac{dT_{ic}}{dt} = 60\left[f_{IC}\dot{w}c_p(T_r - T_{ic}) - Q_{ic}\right] \tag{10.6}$$

$$C_d \frac{dT_d}{dt} = 60\left[\dot{w}c_p(T_{mix} - T_d) - Q_{aux}\right] \tag{10.7}$$

Table 10.1 Classification of Variables. Nominal Values Are Those of the Design Point at the Instant of the Trip, Where Scenarios 1, 4, 5, 7, 9 and 12 Start.

Symbol	Role	Nominal	Range
T_{out} (TE)	CV, constrained	300.0 °C	limit 316 °C
$T_{lim} - T_{out}$	CV, override	16.0 K	4–60 K SP
x_{IC}	MV	2 %	0–100 %
x_A	MV, override	0 %	0–100 %
$\dot{m}_{mk}$	MV, manual only	0 kg/s	0–20 kg/s
$M_p(0)$	DV, measured	55.0 t	11.0–55.0 t
T_0 hours at power	DV, known	8760 h	2–17 520 h
gas fraction g	DV, unmeasured	0 %	0–80 %
aux availability	DV, discrete	1	0–1
valve seized	DV, discrete	0	0 or 1
TE noise	DV, unmeasured	0 °C	0–4 °C 1σ
n_P, n_D	states	0.927, 0.00936	≥ 0
t_s	state	1 s	—
T_f	state	685.0 °C	—
T_c	state	285.0 °C	—
T_r, T_d	states	270.0, 270.0 °C	—
T_{ic}	state	269.3 °C	—
T_p	state	32.0 °C	5–100 °C
M_p	state	55 000 kg	0–71 500 kg
M_t	state	12 000 kg	≥ 0
H_p	state	586.5 kPa	≥ 0
x_M	state	1	0–1
x_{IC}, x_A	states	0.02, 0	0–1

where $\dot{w}c_p$ is the loop thermal capacity rate in MW/K, obtained as $\dot{w}c_p/1000$ from $\dot{w}$ in kg/s and c_p in kJ/kg/K. The core node is a single mean temperature whose outlet is $T_{out} = 2T_c - T_d$, so its enthalpy balance is $\dot{w}c_p(T_d - T_{out}) = 2\dot{w}c_p(T_d - T_c)$, which is where the factor of two in (10.4) comes from. The mixed downcomer inlet is $T_{mix} = T_r - f_{IC}(T_r - T_{ic})$, because the bypass branch returns riser water unchanged and the two streams meet at the downcomer inlet. An isolated condenser therefore leaves a circulating loop with no heat sink on it rather than a stopped loop, which is the mechanism behind scenario 7. The steam generator duty $Q_{main} = x_M P_0$ decays with $\tau_{main} = 2.0\,\text{s}$ from $x_M(0) = 1$, and the auxiliary duty is $Q_{aux} = x_A Q_{aux,\max}$ with availability applied as a multiplier.

Evaluating the right-hand sides at the initial condition of Table 10.1 returns zero on four of the five nodes and −0.059 K/min on the downcomer, the residual being the two percent standby leakage of the isolation condenser. The trip transient therefore starts from a genuine full-power equilibrium rather than from a numerical bump, and the application verifies that on startup.

10.3.3 Circulation From a Momentum Balance

Flow is not a state. It is obtained algebraically each step by balancing the buoyancy head against a quadratic friction loss, as shown in Equation (10.8).

$$g\rho\beta\,\Delta z\,(T_r - T_d) + \Delta p_{pump} = R(x_{IC})\,\dot{w}^2 \tag{10.8}$$

where g is gravitational acceleration, ρ is coolant density, β is the thermal expansion coefficient, Δz is the height difference between the two thermal centers marked in Figure 10.1, and $\Delta p_{pump} = H_p$ is the coasting pump head. With the constants of Table 10.2 the buoyancy gain $g\rho\beta\Delta z$ is 174.2 Pa/K. Two branches carry the flow in parallel and both have quadratic resistance, so the equivalent resistance

and the split are closed form, as shown in Equations (10.9) to (10.13).

$$R_{ic} = \frac{R_{ic0}}{\max(x_{IC}, x_{leak})^2} \tag{10.9}$$

$$R_{by} = R_{by,r} + (R_{by,s} - R_{by,r})\,(1 - H_p/H_{p0}) \tag{10.10}$$

$$R = R_{pipe} + \left(R_{ic}^{-1/2} + R_{by}^{-1/2}\right)^{-2} \tag{10.11}$$

$$f_{IC} = \frac{1}{1 + \sqrt{R_{ic}/R_{by}}} \tag{10.12}$$

$$\dot{w} = \max\left(\dot{w}_{\min},\ \sqrt{\Delta p / R}\right) \tag{10.13}$$

where $\Delta p = \max(0,\ \Delta p_b + \max(0, H_p))$ and Δp_b is the buoyancy term of (10.8), and the condenser branch carries $f_{IC}\dot{w}$. Flooring the flow at $\dot{w}_{\min} = 4\,\mathrm{kg/s}$ rather than letting it reverse is the one place where the hydraulics is not solved but asserted, and Section 10.3.9 says where that fails. Substituting the loop temperature rise $Q/(\dot{w}c_p)$ into (10.8) with the pump stopped gives Equation (10.14).

$$\dot{w}^3 = \frac{g\rho\beta\,\Delta z}{c_p R}\,Q \tag{10.14}$$

Nothing in the code asserts (10.14). It is a consequence of solving the energy balance and the momentum balance separately, and Section 10.4 measures the exponent.

10.3.4 The Condenser and the Four Things That Spoil It

The isolation valve sits on the condensate return, which is how these units are modulated in practice, so throttling floods the lower tubes and takes heat-transfer surface out of service as well as adding resistance. The conductance is one nominal value with four multipliers, as shown in Equation (10.15).

$$UA_{ic} = UA_0\, x_{IC}\, f_{pool}(T_p)\, \frac{1}{1 + g/g_{50}}\left[s + (1-s)\alpha_{air}\right] \tag{10.15}$$

where g is the non-condensable gas fraction in percent, $g_{50} = 25\,\%$ halves the conductance, $\alpha_{air} = 0.05$ is the conductance of an uncovered tube relative to a wet one, and s is the submerged fraction of the bundle. The pool-side and coverage terms are Equation (10.16).

$$\begin{aligned} f_{pool} &= f_{conv} + (1 - f_{conv})\,\sigma\left(\frac{T_p - (T_{sat} - 10)}{12}\right) \\ L &= \frac{M_p}{\rho_p A_{pool}}, \qquad s = \mathrm{clamp}\left(\frac{L - z_b}{z_t - z_b}, 0, 1\right) \end{aligned} \tag{10.16}$$

with $\sigma(u) = \hat{u}^2(3 - 2\hat{u})$, $\hat{u} = \mathrm{clamp}(u, 0, 1)$, the smooth step used throughout. Because T_p is clamped at T_{sat} the pool-side multiplier reaches at most $\sigma(10/12) = 0.926$ and therefore $f_{pool} \le 0.972$, so the maximum conductance the model can deliver is 53.45 kW/K rather than the nominal 55 kW/K. Heat leaves the tube side as $Q_{ic} = UA_{ic}(T_{ic} - T_p)$.

10.3.5 The Pool, in Three Regimes

The pool is a single well-mixed node that heats, then boils, then falls, as shown in Equation (10.17).

$$\begin{aligned} Q_{net} &= Q_{ic} - \frac{\dot{m}_{mk} c_w (T_p - T_{mk})}{1000} - Q_{loss} \\ Q_{loss} &= h_{loss} \max(0, T_p - 25) \\ M_p c_w \frac{\mathrm{d}T_p}{\mathrm{d}t} &= 60\,000\,(1 - f_b)\,Q_{net} \\ \frac{\mathrm{d}M_p}{\mathrm{d}t} &= 60\,\dot{m}_{mk} - \frac{60\,000\,f_b\,Q_{net}^{+}}{h_{fg}} \end{aligned} \tag{10.17}$$

where $Q_{net}^{+} = \max(0, Q_{net})$ and $f_b = \sigma((T_p - (T_{sat} - 2))/2)$ when $Q_{net} > 0$ and the pool is not dry, and zero otherwise. The factor 60 000 carries both the conversion from MW to kW, which is what c_w in kJ/kg/K and h_{fg} in kJ/kg require, and the conversion from seconds to minutes. The switch reaches exactly unity at saturation, so a boiling pool parks at 100 °C with every watt going into vapor instead of fighting a clamp, and M_p is clipped to $[0, 1.3M_{p0}]$ and T_p to $[5, T_{sat}]$ after each step. Makeup drains the tank at $\frac{dM_t}{dt} = -60\dot{m}_{mk}$ and stops when the tank is empty.

Level is the quantity that matters before mass is. With $A_{pool} = 13.1\,\mathrm{m}^2$ a full pool stands 4.198 m deep, and the bundle spans 0.20 to 1.25 m, so 16.4 t of the inventory, 29.8 % of it, lies below the level at which the conductance starts to fall, and 2.62 t of it lies below the bundle entirely.

10.3.6 Actuators and the Instrument Chain

The isolation valve, the auxiliary train and the coasting pump head are integrated as first-order lags, as shown in Equation (10.18).

$$\begin{aligned} \tau_v \frac{dx_{IC}}{dt} &= 60\,(1 - \delta_{stuck}) \left(\frac{u_{IC}}{100} - x_{IC} \right) \\ \tau_{aux} \frac{dx_A}{dt} &= \frac{u_A}{100} - x_A, \qquad \tau_{pump} \frac{dH_p}{dt} = -60 H_p \end{aligned} \tag{10.18}$$

where τ_{aux} is already in minutes and therefore carries no conversion factor, while τ_v and τ_{pump} are in seconds and do. A seized valve sets $\delta_{stuck} = 1$, which stops the position where it is rather than driving it shut, and the demand keeps moving so that the gap between demand and position is visible. After each step the three positions x_M, x_{IC} and x_A are clipped to $[0, 1]$ and the pump head is floored at zero.

The controller sees one measurement. The core outlet is reconstructed as $T_{out} = 2T_c - T_d$, noise is added as a first-order-filtered Gaussian sequence of correlation time 0.3 min scaled so its stationary standard deviation matches the specified value, and the sum is passed through the thermocouple lag, as shown in Equation (10.19).

$$\tau_{TC} \frac{d\hat{T}_{out}}{dt} = T_{out} + \nu(t) - \hat{T}_{out}, \qquad \Delta\hat{T} = T_{lim} - \hat{T}_{out} \tag{10.19}$$

The thermocouple lag advances by the exact zero-order-hold update $x_k = \alpha x_{k-1} + (1 - \alpha) u_k$ with $\alpha = \exp(-\Delta t / \tau)$, and the noise by $\nu_k = \alpha \nu_{k-1} + \sqrt{1 - \alpha^2}\, \sigma_\nu \xi_k$, with ξ_k a unit Gaussian draw and σ_ν the specified standard deviation, so both are stable at any sample period. A second filter of time constant 2.0 min on $\hat{T}_{out}$ supplies the cooldown rate reported in K/h, because the technical-specification limit is a thermal-stress limit on the vessel wall and is not tripped on a thermocouple wiggle. The pool level, the pool mass and the coping-time estimates are displayed without noise or lag.

10.3.7 Coping Time in Closed Form

Because (10.1) integrates analytically, the hours of rejection left in the pool can be predicted rather than only extrapolated, as shown in Equation (10.20).

$$\begin{aligned} M_p \left[h_{fg} + c_w (T_{sat} - T_p) \right] &= \int_{t_s}^{t_s + \Delta t} Q_d \, dt - Q_{aux} \Delta t \\ \int_{t_1}^{t_2} Q_d \, dt &= \frac{A P_0}{1 - n} \left[t^{1-n} - (t + T_0)^{1-n} \right]_{t_1}^{t_2} \end{aligned} \tag{10.20}$$

The application solves (10.20) for Δt by bisection, 24 iterations on a monotone function capped at 72 h, and displays the answer beside a second figure obtained by dividing the remaining pool energy by the present rejection rate. The two disagree by design. At the instant of the trip the closed-form estimate reads 8.85 h and the present-rate estimate reads 1.15 h, because the present rate is 33.9 MW of decay heat and the closed form knows that the curve is falling. Watching the two converge over the first hour is a more instructive display than either alone.

10.3.8 Parameters, Initialization and Numerics

The constants are listed in Table 10.2. Six are hidden and are never displayed, and the rest are adjustable from the model panel.

Total sensible capacity is $C_{plant} = C_f + C_c + C_r + C_{ic} + C_d = 372.76\,\mathrm{MJ/K}$, and it is the quantity that decides every result in this chapter. Scenarios that start at the trip initialize the pumped loop from $T_d = 270\,^{\circ}\mathrm{C}$ and the steam generator duty, with the riser set equal to the downcomer because the steam generator has just removed the whole 530 MW from it. Scenarios that join the transient an hour in set $t_s = 3600\,\mathrm{s}$, $T_p = 100\,^{\circ}\mathrm{C}$, and solve the settled loop analytically from the decay power and the valve position, so they open at an equilibrium whose worst node derivative is $2.4 \times 10^{-13}\,\mathrm{K/min}$.

The app runs on the framework of Section 1.6.1. It integrates with fourth-order Runge-Kutta at a fixed 0.01 min (0.6 s) step, executes the controllers every 0.05 min (3 s) and logs four samples per simulated minute, one every 15 s. The seed is 20260906, and the speeds run from $1\times$ to $200\times$. The exported record holds the true core outlet beside the thermocouple reading, the commanded valve beside the actual position, and the pool inventory beside both coping-time estimates.

10.3.9 Assumptions

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

Five lumped nodes and one fuel node. A real vessel has axial and radial temperature distribution, and a real natural-circulation loop can develop flow oscillations and flashing instabilities that no lumped model produces [7, 9]. The assumption is sound for the hours-long inventory question and unsound for anything that happens in seconds.

Single-phase coolant with a linear buoyancy term. The driving head in (10.8) is proportional to the riser-to-downcomer temperature difference. Natural circulation in a boiling-water reactor is driven mainly by void fraction in the riser, which is a much stronger and much less linear effect [10], and Chapter 6 models such a void-driven core. The single-phase form is right for a pressurized loop and is a substantial simplification for a direct-cycle plant.

Flow cannot reverse. When the riser is colder than the downcomer, which is genuinely the case for the first minutes after the trip, $\dot{w}$ is floored at 4 kg/s rather than run backwards. Real loops do reverse, and reversal is a recognized phenomenon in the natural-circulation literature [7].

One conductance with four multipliers. Equation (10.15) is not a condensing tube bundle with a correlation. The subcooled-to-boiling step is a fitted smooth transition rather than a Rohsenow calculation, and the non-condensable term is a one-parameter degradation rather than a diffusion-layer model, so it reproduces the direction and the rough magnitude of the effect seen in integral experiments and not their detail [14].

The pool is one well-mixed node. Real condenser pools stratify strongly on the way to saturation, so the top reaches boiling well before a mixed model says it does and the conductance changes earlier than shown here.

The auxiliary train is a lag to a fixed rating. There is no pump curve, no exchanger and no service-water side, and its real availability depends on a diesel, a switchgear line-up and a cooling-water source, all of which appear here as one slider.

Stylized disturbances. A pool that loses a quarter of its inventory instantly is not something a real pool does, but it isolates a failure that no temperature reports and that only a level gauge and a coping-time calculation can see. Non-condensable gas that arrives in three discrete steps is a caricature of a gradual accumulation, although it separates the moment the valve leaves its operating range from the moment the temperature starts to move, which a gradual ramp blurs.

The model stops at the relief setpoint. Above T_{lim} the application annunciates and continues integrating, but it does not model the relief valves lifting or the inventory lost through them. That is exactly where a real transient changes character, and the application stops at the boundary rather than pretending to be a loss-of-coolant code.

Table 10.2 Model Parameters. Hidden Constants Are Never Displayed.

Symbol	Meaning	Value	Unit	Visibility
P_0	rated thermal power	530	MW	adjustable
A	Way–Wigner coefficient	0.066	—	adjustable
n	Way–Wigner exponent	0.2	—	adjustable
τ_P	prompt flux collapse	2.4	s	adjustable
τ_D	delayed-neutron tail	75	s	adjustable
C_f	fuel capacity	7.26	MJ/K	adjustable
UA_{fc}	fuel to coolant	1.325	MW/K	adjustable
C_c	core node capacity	82.7	MJ/K	adjustable
C_r	riser capacity	83.6	MJ/K	adjustable
C_{ic}	tube-side capacity	33.2	MJ/K	adjustable
C_d	downcomer capacity	166.0	MJ/K	adjustable
c_p	coolant heat capacity	5.3	kJ/kg/K	adjustable
Δz	thermal-center height	8.0	m	adjustable
ρ	coolant density	740	kg/m^3	adjustable
β	expansion coefficient	3.0×10^{-3}	kK	adjustable
R_{pipe}	loop resistance	0.014	Pa s^2/kg^2	adjustable
R_{ic0}	condenser branch	0.08	Pa s^2/kg^2	adjustable
$R_{by,r}$	bypass, pumps running	0.040	Pa s^2/kg^2	adjustable
$R_{by,s}$	bypass, pumps stopped	0.65	Pa s^2/kg^2	adjustable
H_{p0}	pump head at rated flow	586.5	kPa	adjustable
τ_{pump}	pump coastdown	15	s	adjustable
τ_{main}	loss of the main sink	2.0	s	adjustable
UA_0	condenser conductance	55	kW/K	adjustable
α_{air}	uncovered-tube factor	0.05	—	adjustable
f_{conv}	subcooled pool factor	0.62	—	adjustable
g_{50}	gas fraction halving UA	25	%	adjustable
τ_v	valve full stroke	20	s	adjustable
A_{pool}	pool plan area	13.1	m^2	adjustable
M_{p0}	pool when full	55000	kg	adjustable
z_b, z_t	bundle elevations	0.20, 1.25	m	adjustable
c_w	pool heat capacity	4.18	kJ/kg/K	adjustable
h_{fg}	latent heat	2257	kJ/kg	adjustable
T_{sat}	pool saturation	100	°C	adjustable
h_{loss}	pool surface loss	4.0×10^{-4}	MW/K	adjustable
M_{t0}	makeup tank	12000	kg	adjustable
T_{mk}	makeup temperature	25	°C	adjustable
$Q_{aux,\max}$	auxiliary train	2.0	MW	adjustable
τ_{aux}	auxiliary loading	1.5	min	adjustable
T_{lim}	core outlet limit	316	°C	adjustable
T_{warn}	warning margin	10	K	adjustable
$\dot{T}_{\max}$	tech-spec cooldown rate	55	K/h	adjustable
$T_{f,trip}$	fuel-damage alarm	1200	°C	adjustable
t_{cope}	coping time required	8	h	adjustable
τ_{TC}	thermocouple lag	12	s	adjustable
δ	delayed share of flux	0.010	—	hidden
$t_{\min}$	earliest power-law time	1.0	s	hidden
$\dot{w}_{\min}$	circulation floor	4	kg/s	hidden
x_{leak}	valve seat leakage	0.002	—	hidden
x_{sb}	valve standby position	0.02	—	hidden
ρ_p	pool density for level	1000	kg/m^3	hidden

Table 10.3 Settled Loop Against Isolation Valve Position, Measured at $t_s = 1\,\text{h}$ With $Q_d = 5.694\,\text{MW}$, a Boiling Pool and a Covered Bundle. The Gain Column Is a Backward Difference Quotient Over 0.1 % of Travel.

x_{IC} (%)	UA_{ic} (kW/K)	$\dot{w}$ (kg/s)	T_{out} (°C)	$\partial T_{out}/\partial x_{IC}$ (K/%)
20	10.69	87.7	666.4	−28.2
25	13.36	92.6	554.0	−18.1
40	21.38	106.3	385.3	−7.05
50	26.73	114.6	329.0	−4.51
60	32.07	122.2	291.5	−3.13
80	42.76	136.0	244.5	−1.76
100	53.45	147.9	216.3	−1.13

10.4 Static Behavior of the Loop and the Valve

10.4.1 The Cube Root, Measured

Equation (10.14) is a prediction about a model that never states it. The measurement feeds a power into the energy balance, solves the settled loop temperatures, and hands only those temperatures to the momentum balance (10.8), which never sees the power. Varying the power from 2 to 16 MW returns an exponent of 0.33333, and varying Δz from 4 to 16 m at fixed power returns 0.33333 as well. The consequence is that the loop is stiff. An hour after the trip with the valve wide open the circulation is 147.9 kg/s with 7.26 K around the loop, which is 4.44 % of the 3330 kg/s the pumps used to deliver, and opening the valve from 60 to 80 % moves the flow only from 122.2 to 136.0 kg/s, a rise of 11.3 %. Going the whole way from half open to wide open, a factor of two in the condenser branch resistance, is worth 29.1 %.

Two lessons follow, and they point in opposite directions. The actuator has almost no authority over the flow it appears to control, which is a useful warning about inferring a valve's function from a drawing. The same insensitivity makes the loop self-regulating without anyone having designed it that way, because a rise in power raises the temperature difference, the head and the flow together, and the exponent of one third guarantees that the correction is monotone and bounded.

10.4.2 A Manipulated Variable With a Factor of Twenty-Five in Gain

What the valve does move strongly is the heat-transfer surface. At a fixed decay power the condenser must reject $Q_d = UA_{ic}\,(T_{ic} - T_p)$, and because UA_{ic} is proportional to x_{IC} by (10.15) the settled temperature is close to a rectangular hyperbola in the valve position, which Table 10.3 measures an hour after the trip on a boiling pool with the bundle covered.

The gain runs from −1.13 K/% wide open to −28.2 K/% at a fifth open, a ratio of 25 across the actuator's own travel, and the position that holds 300 °C at these conditions is 57.4 %. That is the reason one set of proportional-plus-integral settings cannot be right everywhere on this valve, and the practical consequence, which the shipped tuning reflects, is to operate in the open half of the travel and treat the tight-throttling end as a place to visit rather than to live. It is also the reason the fully passive architecture is well behaved. A latched-open valve sits at the flattest point of Table 10.3 and cannot be moved off it.

10.4.3 The Arithmetic That Decides Everything

At steady state the pool receives exactly the decay heat whatever the valve is doing, because the valve is not in the energy balance. What the valve moves is the other term, as shown in Equation (10.21).

$$\int Q_{ic}\,\mathrm{d}t = \int Q_d\,\mathrm{d}t + C_{plant}\,\Delta T_{plant} - \int Q_{aux}\,\mathrm{d}t \tag{10.21}$$

Everything on the right except the middle term is fixed by physics or by the availability of a diesel. The middle term is the plant's own stored heat, and its exchange rate into pool water, introduced

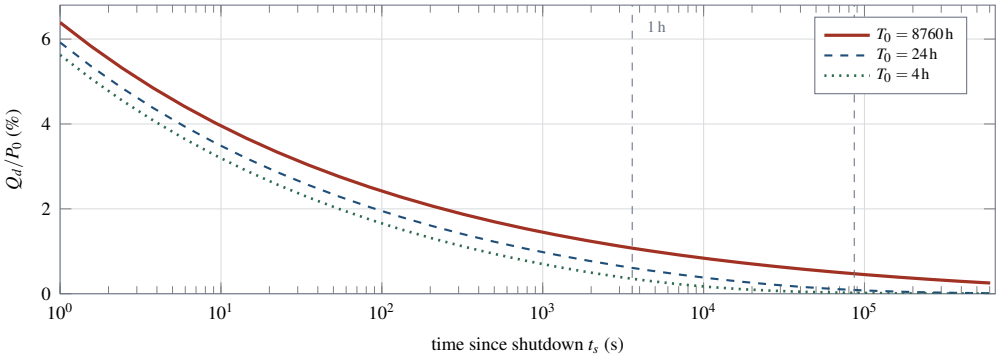


Figure 10.2 Decay heat fraction against time since shutdown for three operating histories: a year at power (solid), a day (dashed) and four hours (dotted). The year-long core falls only from 6.391 to 0.2525 % over six decades of time. Four hours at power supplies 32.9 % of the year-long fraction at one hour and a full day supplies 56.7 %, because a short run has not built the long-lived part of the fission-product inventory.

in Section 1.4.5, is $C_{plant}/h_{fg} = 165.2\,\mathrm{kg/K}$. Each megawatt rejected boils 0.4431 kg/s, which is 1595 kg per megawatt-hour, so the 2.0 MW auxiliary train is worth 3.19 t of pool water an hour whenever it runs. Those three numbers, together with the 16.4 t that sits below the top of the bundle, are enough to predict every architecture comparison in Section 10.6 before running one.

10.4.4 Operating History

The second term of (10.1) is the part a coping-time analysis has to assume, and it is not a small correction. Figure 10.2 plots the decay fraction on a logarithmic time axis for three operating histories.

The measured fractions for the year-long core are 6.391 % (33.87 MW) at 1 s, 2.701 % at 1 min, 1.074 % (5.694 MW) at 1 h, 0.6378 % (3.380 MW) at 8 h, 0.4709 % at 1 d and 0.2525 % at one week. At $t_s = 1\,\mathrm{h}$ a core that had been at power for four hours makes 32.9 % of that, one day gives 56.7 % and thirty days gives 87.4 %. The same hardware therefore has a completely different coping time depending on a number that no instrument on the plant reads, which is why every published coping time is quoted against a stated operating history.

10.5 Four Architectures and Two Controllers

The heat-removal architecture is a segmented control with four settings, and the interesting result is that the two obvious ones are both wrong. MANUAL gives the operator three stations, which are the isolation valve, the auxiliary train and the makeup pump, with nothing watching anything. FULLY PASSIVE latches the isolation valve wide open and does nothing else, which is the plant as designed with no electrical power and no operator. THROTTLED PASSIVE puts a proportional-plus-integral controller on the isolation valve holding a core outlet setpoint. PASSIVE PLUS AUXILIARY adds a second controller whose process variable is the temperature margin and whose output is the auxiliary train.

10.5.1 The Implemented Algorithm

Both loops use the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement, executed every $\Delta t_c = 0.05\,\text{min}$. In the notation of this chapter the update is Equation (10.22).

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D, \qquad c_k = s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k,\ u_{\min},\ u_{\max}) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k) \end{aligned} \tag{10.22}$$

where Δu_k^P and Δu_k^D are the changes in the proportional and derivative terms of Equation (1.13), $\Delta u_k^I = (K_c/\tau_I)\,e_k\,\Delta t_c$, s_k is the running sum S_k of Equation (1.14), c_k is the output before limiting v_k of Equation (1.15) with no bias, u_k is the output the loop emits, and T_t is the anti-windup tracking time. The integral is seeded with the actuator position at reset, so each loop starts from the position it finds. Anti-windup is back-calculation through the $(\Delta t_c/T_t)(u_k - c_k)$ term, which corrects the integral alone and never truncates s_k. Both loops ship with $\tau_D = 0$, so Δu_k^D is identically zero and the derivative filter is inactive, and neither loop carries a slew limit on its output because the valve and the train supply their own rate limits through (10.18).

The loop actions, in the convention of Section 1.3, are opposite. Opening the isolation valve makes the core colder, so the temperature loop has a negative process gain and is direct acting with $a = -1$. The override's process variable is a margin that grows when the train loads up, so that loop is reverse acting with $a = +1$. The tuned settings, which the automatic scenarios 4 to 6, 8, 10 and 11 load, are $K_c = 1.2\,\%/\text{K}$, $\tau_I = 6.0\,\text{min}$, $T_t = 1.5\,\text{min}$ and $u \in [0, 100]\,\%$ for the temperature loop, and $K_c = 4.0\,\%/\text{K}$, $\tau_I = 6.0\,\text{min}$, $T_t = 1.0\,\text{min}$ and the same output span for the override. The manual scenarios and the graded challenge load deliberately weak gains of 0.04 %/K with 30 min and 0.3 %/K with 45 min, with the same tracking times. The tracking constant on the temperature loop is deliberately four times faster than its integral time, because that valve spends stretches on a limit. In scenario 5 it is shut for a minute while the trip transient carries the core outlet below the setpoint and wide open from 37 to 59 min while the plant heats back up, and in scenario 10 it is wide open from 95 min to the end. A tracking constant as slow as τ_I leaves the loop late coming off the stop, by 4.7 min in a counterfactual run of scenario 5 with $T_t = \tau_I$.

In the fully passive architecture the controller is not removed. Its output is forced to 100 % and a tracking routine sets I_k each scan so that $s_k + I_k$ equals the applied position exactly, the bumpless transfer of Section 1.6.2 that MANUAL uses as well. A throttled architecture selected during a fully passive run therefore starts from a valve commanded wide open, and one selected before the run starts from the 2 % standby position, which matters in the graded challenge. Makeup has no automatic mode anywhere in the app, on purpose, because deciding when to spend twelve tonnes is the operator's decision.

10.5.2 The Override Is a Selector in Everything but Name

The auxiliary loop is a plain proportional-plus-integral controller whose setpoint is a margin rather than a temperature. While the valve is holding more margin than the override asks for, the error is large and negative, the output rests on its low limit and the back-calculation term in (10.22) holds the integral there, so the train stays off and comes on only when the margin is actually threatened. That behavior is a low selector implemented with an ordinary loop, and it is the standard way a protective constraint controller is built [2]. The recycle override of Volume 4, Chapter 2 rests on its lower clamp in the same way.

The consequence is a design rule that is invisible on the operator's display. The valve holds a margin of $T_{lim} - T_{SP}$. If the override setpoint ΔT_{SP} is below that value the error never goes positive in steady operation and the train never runs. If it is above, the override can never be satisfied and the train runs continuously. At the settings scenario 6 loads, the valve holds 16.0 K and the override asks for 10 K, so the train runs only when a disturbance drives the margin toward 10 K. In that scenario

Table 10.4 Scenario Sequence and the Mechanism Each Exercises

#	Scenario	What it exercises
1	The scram, by hand	decay-heat scaling, pumps to buoyancy
2	Driving the valve by hand at $t_s = 1\,\mathrm{h}$	hyperbolic characteristic, gain ratio of 25
3	Step test on the passive loop	cube-root insensitivity of the flow
4	Fully passive, valve latched open	it works, and the pool pays
5	Throttled passive at 300 °C	inventory bought with margin
6	Margin override on the auxiliary train	selector behavior, the setpoint pair
7	Isolation valve seizes at 2 %	no strong handle, stored heat as a debt
8	Auxiliary train unavailable	an override with nothing to override
9	Four hours at power instead of a year	operating history and coping time
10	Non-condensables to 50 %	disturbance hidden in the actuator
11	A quarter of the pool is lost	failure visible in no temperature
12	Graded challenge, 8 h on 46.8 t	structure, setpoint pair and adequate gains

it first leaves zero at 49.7 min, with the margin at 10.8 K and falling after the non-condensable gas has arrived, because the proportional term rises faster than the tracking integral follows. It takes 0.54 GJ in the first three hours with its demand never above 14.4 %. Moving the pair to 308 °C and 14 K inverts the relationship, and the measured effect at eight hours is 18.2 t of pool remaining and 31.4 GJ taken by the train against 6.39 t and 3.49 GJ with the original pair, with the bundle never uncovering instead of uncovering at 5.59 h. The process gain of the override is 0.70 K of margin per percent of train at the valve position that holds 300 °C an hour after the trip, and it falls as the valve opens, so this loop inherits a gain schedule from the other one.

🕮 **Pause.** Try Exercise 10.2 before reading on.

10.6 Measured Results

Twelve scenarios are arranged as a laboratory sequence, listed in Table 10.4. A scenario fixes the starting mode, the architecture, the two setpoints, the four tuning constants, the disturbance settings and any timed events. Every number in this section was obtained by driving the shipped application headlessly through the same stepping routine the browser uses, at the 0.01 min physics step, with each setting changed after the case is loaded and before the run starts. Peaks and crossing times are read at the physics step, tube uncovery is taken as a submerged fraction below 0.999, and every figure that is still moving is quoted with the time at which it was read. No scenario carries thermocouple noise, so every run is deterministic.

10.6.1 The First Two Minutes, and the Next Eight Hours

Scenario 1 starts at the trip and makes its point with arithmetic before any controller is involved. The prompt flux collapses with a 2.4 s time constant, the steam generator is isolated in 2.0 s and the pumps coast down with a 15 s time constant, so the core outlet first rises to a peak of 309.9 °C at 6.6 s while the main heat sink disappears faster than the fission power does, and then falls to 293.2 °C by 30 s as the decay heat drops to 16.5 MW. Circulation falls from 3330 kg/s to 395 kg/s at 30 s and 147.5 kg/s at 60 s, and from that point onward the water moves because the riser is hotter than the downcomer and for no other reason.

Then nothing dramatic happens for the rest of the shift, and that is what makes the scenario worth eight hours of attention. With the condenser placed in service at one minute the core outlet holds between 294.7 and 295.3 °C for half an hour, the pool reaches saturation at 46.1 min, and the level reaches the top of the tube bundle at 3.93 h. After that the conductance falls with the level, the plant heats up again, and the core outlet crosses the relief setpoint at 9.62 h. The pool's 139.8 GJ against the integrated decay curve predicts 8.85 h to the last kilogram by (10.20), and the last kilogram is not the deadline.

Table 10.5 Three Architectures on the Identical Trip Transient, Each Line Read at the Time It Names Because None of the Three Settles. The Third Column Holds 300 °C With the Override Setpoint Above the 16 K the Valve Delivers, Which Latches the Auxiliary Train In.

	Fully passive	Throttled 300 °C	Override 300 / 24 K
T_{out} at 2 h (°C)	250.7	298.8	299.5
T_{out} at 4 h (°C)	211.9	299.8	299.8
M_p at 2 h (t)	35.90	43.44	49.35
M_p at 4 h (t)	15.72	29.92	42.13
M_p at 8 h (t)	2.67	7.61	31.76
Stored-heat deficit at 4 h (GJ)	32.86	0.07	0.08
Auxiliary energy to 8 h (GJ)	0.00	0.00	56.71
Valve position at 4 h (%)	100.0	37.4	18.0
$\dot{w}$ at 4 h (kg/s)	138.2	91.6	75.7
Peak T_{out} to 8 h (°C)	305.7	314.2	309.8
Minimum margin to 8 h (K)	10.3	1.8	6.2
Minutes over limit to 8 h	0.0	0.0	0.0
Uncovery (h)	3.93	6.26	never
Limit first crossed (h)	9.62	9.24	never
Minutes over limit to 12 h	142.9	165.9	0.0

TRY IT IN THE APP 1 · The Scram, and the Pool That Has Hours

Load scenario 1 and press Run at 20 or 60 times speed. Watch the core outlet temperature first: it rises to 309.9 °C as the main heat sink disappears, then falls as the fission power dies away. Then watch the pool strip for hours: the level reaches the top of the tube bundle at 3.93 h.

10.6.2 Fully Passive, and the Bill It Runs Up

Scenario 4 latches the valve open and adds nothing. The result is that the plant is comfortable and the pool is not, as Table 10.5 reports against two throttled architectures. All three runs share the same initial state, which is the full-power equilibrium of Table 10.1 with 55.000 t of pool water at 32.0 °C and the valve at its 2 % standby position. The first two are scenarios 4 and 5 as loaded, and the third is scenario 5 with the architecture set to passive plus auxiliary and the override setpoint raised to 24 K before the run.

The passive column reads as a success on every temperature line. The core outlet goes the wrong way for a temperature limit, reaching 250.7 °C at two hours and 211.9 °C at four, and the minimum margin over the whole eight hours is 10.3 K, which occurs in the first minutes when the plant is still hot. It also does not violate the cooldown rate. Excluding the first ten minutes, during which the collapse of fission power itself produces a filtered rate of up to 365.3 K/h and the rate stays above the limit for 1.85 min, the peak cooldown rate is 32.2 K/h, comfortably inside the 55 K/h technical-specification limit. On this plant the honest argument for throttling is not vessel thermal stress. It is water. Figure 10.3 draws four lines of Table 10.5 as bars.

The cost appears in the inventory lines. A wide-open condenser does not stop when it has taken the decay heat, and the 88.1 K of unasked-for cooldown between 300.0 °C and 211.9 °C is 32.86 GJ out of the plant by (10.21) and 14.56 t out of the pool at 165.2 kg/K. That is a quarter of the original inventory, spent on a cooldown no constraint asked for, and it is why the pool reaches the top of the bundle at 3.93 h rather than 6.26 h. Beyond that point the run is one-way. The pool is dry at 11.09 h and the core outlet is at 372.8 °C by twelve hours after 142.9 min above the limit. Figure 10.5 shows the same run at eight hours, after it has turned around and before it fails.

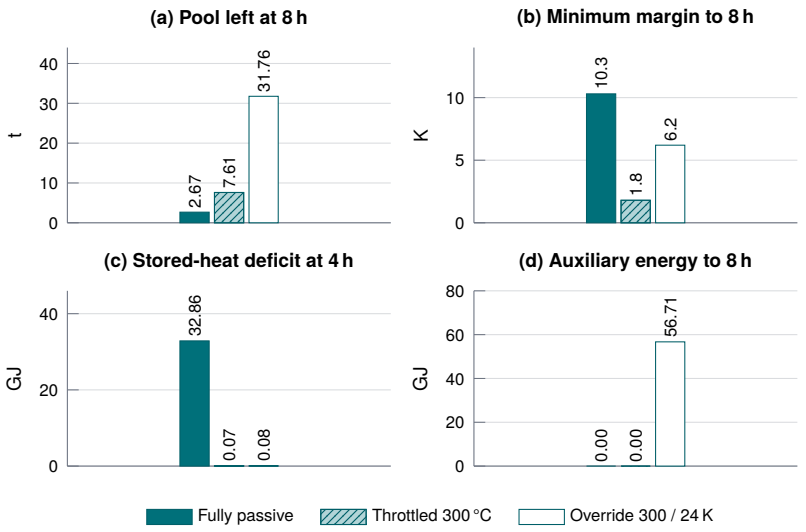


Figure 10.3 The three architectures of Table 10.5 on the identical trip transient: (a) pool water left at eight hours, (b) the minimum temperature margin over eight hours, (c) the stored-heat deficit at four hours and (d) the auxiliary energy taken over eight hours. Solid: fully passive. Hatched: throttled at 300 °C. White: the override with the train latched in. Notice that the architecture with the most margin keeps the least water, and that its deficit in stored heat is the water it has spent.

10.6.3 Throttling a Passive System on Purpose

Scenario 5 changes one control. Holding the core outlet at 300 °C leaves 29.92 t in the pool at four hours against 15.72 t passive, a saving of 14.20 t. The stored-heat ledger of (10.21) predicts it to within 2.3 %, because the difference in stored energy between the two runs at four hours, divided by h_{fg}, is 14.53 t. Throttling also pushes uncovery from 3.93 h to 6.26 h. The mechanism is one line of arithmetic. The pool receives the same decay heat either way, so the only term the valve can move in (10.21) is the stored heat, and it moves it by refusing to cool the plant down. Read at four hours the stored-heat deficit is 0.07 GJ against 32.86 GJ, and its peak over the whole run is 3.36 GJ against 33.84 GJ.

What it costs is margin, held continuously. The minimum margin over eight hours is 1.8 K against 10.3 K passive, the peak core outlet is 314.2 °C against the 316 °C limit, and both of those occur in the first hour while the loop is still recovering from the trip. Neither architecture spends a single minute over the limit within eight hours, and the observation window matters here, because continuing the throttled run to ten hours gives a peak of 333.9 °C and 45.9 min over the limit once the bundle has uncovered. The throttled answer is not a stable end state. It is a way of moving the deadline.

The throttled loop also shows the gain problem of Table 10.3 in operation. Holding 300 °C puts the valve near 37 % at four hours, where the settled process gain is about −5.8 K/%, and the circulation falls to 91.6 kg/s against 147.9 kg/s wide open. Between two and six hours the valve position swings from 29.4 to 52.4 % and back to 37.4, then settles near 37 % in a damped oscillation with a period near 115 min, while the core outlet stays inside 298.2 to 303.1 °C. The commanded position accumulates 232 % of travel over the first 400 min at the controller scan, which the scorecard, sampling every 15 s, reports as 226 %. The cycle is a consequence of a loop tuned for one end of a hyperbola operating at the other while its load falls, and it costs nothing here only because the process it is regulating is slower still.

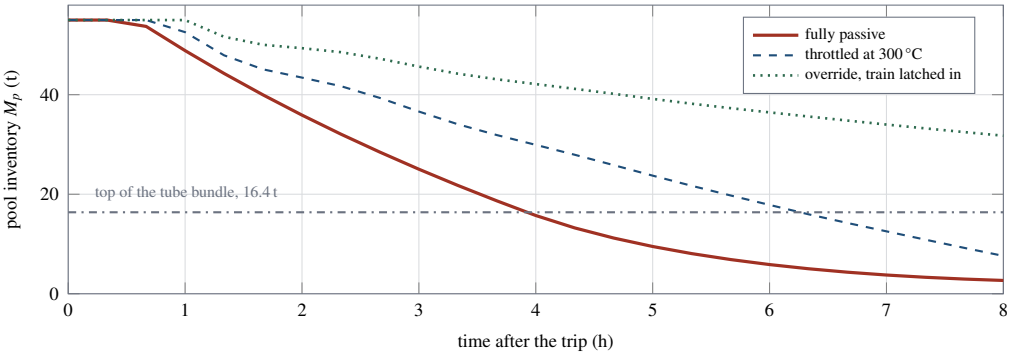


Figure 10.4 Condenser pool inventory for the three architectures of Table 10.5 on the identical trip transient: fully passive (solid), throttled (dashed) and the override with the train latched in (dotted). The dash-dotted line marks the top of the tube bundle. The passive trace crosses the top of the tube bundle at 3.93 h, the throttled trace at 6.26 h, and the override with the auxiliary train latched in never reaches it. All three are flat for at least the first half hour because a subcooled pool absorbs 15.63 GJ of sensible heat before it boils. Boiling starts at 98 °C, at 34, 50 and 60 min on the three paths, and the pool reaches saturation at 45.3, 60.6 and 74.6 min.

10.6.4 Inventory Trajectories

Figure 10.4 plots the pool inventory for the three architectures of Table 10.5, with the level at which the bundle begins to uncover marked.

The three traces separate for one reason each. The passive trace falls fastest because it is boiling the plant's stored heat as well as the decay heat, the throttled trace falls at the decay heat alone, and the override trace falls at the decay heat minus the 2.0 MW the train removes. The last of the three closes (10.21) to within a few percent without any further assumption. The train takes 56.71 GJ over eight hours, which at 2257 kJ/kg is 25.1 t of water not boiled, and the measured difference in remaining inventory is 31.76 t against 7.61 t, or 24.2 t.

> **✦ DID YOU KNOW? The Water Saver Circulates the Least**
>
> From the simulator: at four hours the override with the train latched in has its valve 18.0 % open and circulates 75.7 kg/s, against 100.0 % and 138.2 kg/s fully passive, yet it keeps 42.13 t in the pool against 15.72 t (Table 10.5).

10.6.5 Three Ways the Plant Takes the Handle Away

Scenario 7 seizes the isolation valve at its two percent standby position. With 372.8 MJ/K of plant and 9.3 MW of average core power from 5 to 10 min, the capacity-weighted mean temperature rises at 1.46 K/min against the 1.49 K/min that Q_{core}/C_{plant} predicts, the difference being the 0.17 MW that still leaks through the seated valve, and the core outlet reaches the relief setpoint at 10.7 min. The auxiliary train would not have saved it, because 2.0 MW against 9.3 MW is not a heat sink. The second lesson arrives when the valve is freed at 25 min. The peak of 332.5 °C is thirty-odd kelvin of stored energy that now comes out of the pool on top of the decay heat, and by five hours the run has spent 31.99 GJ on cooldown. Time spent hot is water.

Scenarios 6 and 8 exercise the override on a base that joins the transient an hour after the trip with 53.350 t in the pool already boiling, the valve holding 300 °C at 57.4 %, and non-condensable gas rising to 30 % at 30 min. Table 10.6 separates the setpoint pair from the availability of the diesel by changing one thing at a time from that base.

Table 10.6 The Override, With One Change at a Time From a Common Base at $t_s = 1$ h. Columns One and Two Differ Only in the Setpoint Pair Applied at 3 h, and Columns Two and Three Differ Only in the Availability of the Auxiliary Train. All Lines Are Read at $t = 8$ h Unless Stated.

	308 / 14 K aux available	300 / 10 K aux available	300 / 10 K no aux
M_p at 4 h (t)	27.53	25.43	25.10
M_p at 8 h (t)	18.19	6.39	6.25
Auxiliary energy (GJ)	31.42	3.49	0.00
Train position at 4 h (%)	99	0	0
Peak T_{out} (°C)	308.3	309.8	318.0
Minimum margin (K)	7.7	6.2	−2.0
Minutes over limit	0.0	0.0	7.5
Uncovery (h)	never	5.59	5.55
Limit first crossed (h)	never	never	7.88

The first two columns are identical trajectories until 3 h, when the setpoint pair is applied. Giving away eight more kelvin of margin and raising the override setpoint above what the valve then delivers latches the train in, and the effect is 18.19 t of pool remaining at eight hours against 6.39 t, with the bundle never uncovering instead of uncovering at 5.59 h. Nothing about the block diagram changed.

The third column is the one that corrects an easy inference. With the loose pair the train barely runs, taking only 3.49 GJ in eight hours, so removing it entirely costs almost no inventory at all, 6.25 t against 6.39 t. What it costs is temperature. The peak core outlet rises from 309.8 °C to 318.0 °C, the minimum margin goes from 6.2 K to −2.0 K, and the plant spends 7.5 min above the limit from 7.88 h. The value of an auxiliary train is therefore whichever resource its setpoint pair makes it spend, which is inventory when it is latched in and margin when it is a last-resort protective action, and the same hardware delivers one or the other depending on two numbers.

Scenario 8 also shows what a control system that cannot act is still good for. The isolation valve is first commanded wide open at 70.5 min. The override demand reaches 26 % near 95 min, is zero for most of the rest of the run, and climbs from 7.05 h to reach 100 % at 7.66 h, while the train delivers nothing throughout. The train's position follows the demand, because availability scales only the duty, so the diagnostic gap is between demand and duty rather than between demand and position. It says the passive path is carrying the whole load and the plant has no reserve behind it.

Scenario 10 lets non-condensable gas collect in the bundle in three steps to 50 %, which by (10.15) leaves a third of the conductance. For the first hour and a quarter the core outlet stays within 5.0 K of the setpoint, with a peak of 305.0 °C at 53 min, because the controller holds it, and nothing on the temperature strip suggests the heat sink is failing. What moves is the valve, from 55.7 % before the gas arrives to 88.3 % at one hour and full open at 78.5 min. Then it stops, and from there the temperature drifts upward, reaching 304.3 °C at four hours, crossing the limit at 7.22 h and reaching 327.0 °C by eight. On a controlled loop the disturbance appears in the manipulated variable, so the actuator position is the diagnostic and the process variable is not, and a valve running out of travel is the moment the plant stops having a control system.

🕮 **Pause.** Try Exercise 10.3 before reading on.

10.6.6 The Failure That Shows Up in No Temperature

Scenario 11 removes a quarter of the pool at 25 min on the throttled architecture. For three hours nothing happens. The heat balance is unchanged, the bundle is still covered, the core outlet stays within 298.9 to 300.1 °C, the valve only follows the falling decay heat from 54.8 to 40.6 %, and no measurement on the temperature side of the plant reports the event at all. Then the level reaches the top of the bundle at 3.18 h instead of the 5.50 h the same case reaches without the leak, the conductance falls with it, the valve runs to 100 % and stops, and the run is one-way from there.

Table 10.7 Graded Challenge, Read at $t = 8$ h. All Four Runs Use the Same Disturbances, and the Last Three Use the Suggested Gains.

	Passive	Throttled 306	Override 306 / 14 K	Override 306 / 10 K
Peak T_{out} (°C)	341.8	346.5	313.2	313.5
Minutes over limit	69.9	120.8	0.0	0.0
Minimum margin (K)	−25.8	−30.5	2.8	2.5
M_p at 8 h (t)	4.81	5.56	11.56	8.59
Auxiliary energy (GJ)	0.00	0.00	28.28	20.70
Uncovery (h)	4.30	4.90	6.665	5.965

The limit is crossed at 5.87 h and the core outlet reaches 370.3 °C by eight hours after 128.0 min above it. Nothing, nothing, nothing, and then everything at once is the characteristic shape of an inventory-limited failure, and it is the argument for treating the coping-time calculation of (10.20) as a first-class instrument rather than as paperwork.

> **TRY IT IN THE APP** **11 · A Quarter of the Pool Disappears**
>
> Load scenario 11 and press Run. A line fails at 25 min and a quarter of the pool is lost. Watch the core outlet: it stays between 298.9 and 300.1 °C for three hours. Then watch the pool strip and the valve: the level reaches the top of the bundle at 3.18 h, the valve runs to 100 %, and the limit is crossed at 5.87 h.

The makeup decision inherits the same structure. Twelve tonnes divided by the boil-off rate buys 1.32 h if it is spent at one hour, when the decay heat is 5.694 MW, and 2.06 h if it is spent at six hours, when the decay heat is 3.646 MW. Later is worth more, and too late is worth nothing, because once the conductance has fallen the temperature is already climbing and the option of choosing the moment has been lost.

For contrast, scenario 9 changes only the operating history. Four hours at power instead of a year gives 0.3532 % at one hour against 1.074 %, and the bundle that uncovers at 3.93 h in scenario 4 never uncovers at all. At eight hours the pool still holds 20.94 t and the core outlet is 115.9 °C, having spent 68.64 GJ on cooldown, which is more than twice what the year-long core spent and does not matter because there was enough water.

10.6.7 The Challenge, and Why Both Intuitions Fail

The graded challenge starts the same trip with 46.8 t in the pool instead of 55.0 t, no makeup available, an auxiliary train limited to 1.0 MW, and non-condensable gas at 10 % from the start, rising to 30 % at two hours and 45 % at five. It loads the fully passive architecture with the weak gains of the manual scenarios. The target is eight hours with the core outlet below 316 °C throughout, no makeup used and water still in the pool, judged on the case as given. Table 10.7 reports four configurations, the last three with the suggested gains, which are those of scenarios 4 to 6 except for a temperature-loop gain of 3.0 %/K in place of 1.2.

Both intuitive answers fail, and they fail differently. Rejecting as hard as possible uncovers the bundle at 4.30 h and spends 69.9 min over the limit. Hoarding by throttling to 306 °C and leaving the train alone is worse, at 120.8 min over, 15.0 min of it in the first hour and the rest from 6.24 h, because the margin was given away in advance and there is nothing left when the gas arrives. The answer is both at once, which is a deliberate permanent sacrifice of most of the temperature margin in exchange for inventory, plus an override configured so that it spends the one resource that is not water. Holding 306 °C with the override asking for 14 K peaks at 313.2 °C at 26.6 min with nothing over the limit, keeps the train at full output from 14.8 min for 451 of the 480 min, takes 28.28 GJ on it, and finishes with 11.56 t in the pool and uncovery at 6.665 h.

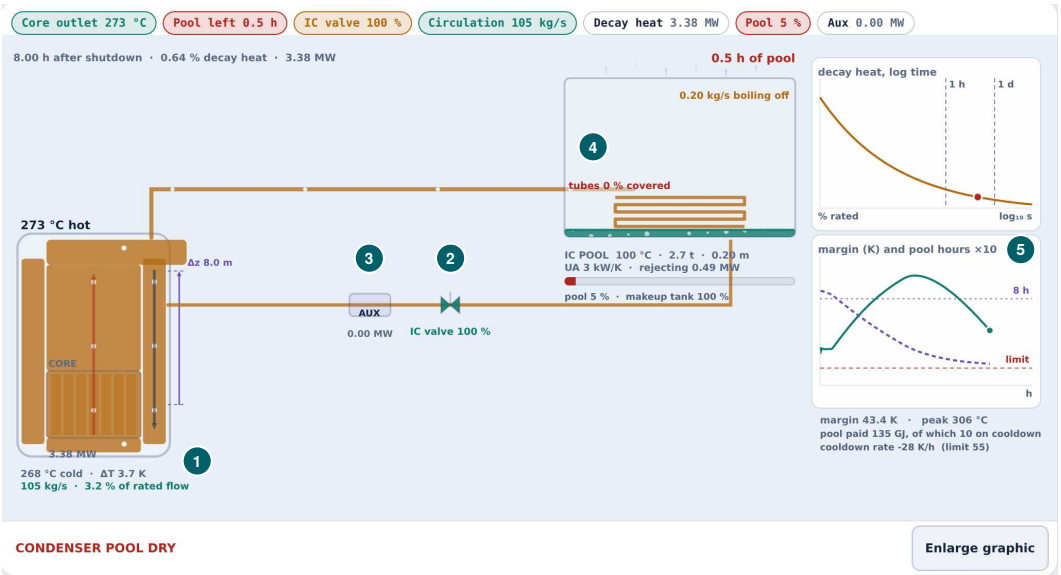


Figure 10.5 The fully passive architecture of scenario 4 at $t = 8.00$ h, with the trends spanning the whole run. Nothing has gone wrong with the temperature. Decay heat is down to 3.38 MW, which is 0.64 % of rated, the core outlet sits at 273 °C on 3.7 K around the loop, and the margin to the limit is still 43.4 K with a peak of 306 °C behind it. The isolation valve has been latched wide open for eight hours, and buoyancy alone is moving 105 kg/s, or 3.2 % of rated flow, on the 8.0 m between the thermal centers. What has been spent is the sink. The pool is at 5 %, boiling off at 0.20 kg/s, the tube bundle is 0 % covered, the pool annunciation has fired below 5 % of the full inventory, and half an hour of usable inventory remains after the 135 GJ the pool has absorbed. The core outlet is climbing at 28 K/h from its 209 °C minimum at 4.5 h, so the plant was cooled faster than any requirement asked for and the sink has paid for it. The view is cut to the plant panel of the application. Callouts: (1) the reactor vessel with the core and the natural-circulation loop, (2) the isolation valve latched open, (3) the idle auxiliary train, (4) the pool with its tubes uncovered, (5) the margin and pool-hours inset.

The gains have to be strong enough to open the valve in the first hour. Every configuration starts with the valve at its 2 % standby position, and the throttling loop keeps it shut for two minutes while the core outlet climbs back toward the setpoint, with the pool still subcooled. At the 1.2 %/K of scenarios 4 to 6 the valve opens too slowly, and the core outlet peaks at 319.2 °C at 36.5 min with 41.1 min over the limit, which fails. On the weak gains the case loads, the valve command never exceeds 71.2 % and the core outlet is over the limit from 10.8 min to the end. The setpoint pair the case loads, 300 °C and 20 K, already latches the train because 20 K exceeds the 16 K the valve holds, and on the suggested gains it also passes, with 10.69 t left.

The fourth column is the instructive one. Moving the override setpoint from 14 to 10 K, which leaves the block diagram, the gains and the valve setpoint untouched, still meets the temperature target, because the trip transient and later the gas drive the margin below 10 K often enough to load the train. It is at full output from 21.9 to 60.2 min and again from 433.8 min, about 85 min in all. It takes only 20.70 GJ on the train instead of 28.28 GJ and finishes with 8.59 t instead of 11.56 t, with uncovery 0.70 h earlier. The temperature scorecard cannot distinguish the two configurations and the inventory scorecard separates them by 2.97 t, which is the whole argument for reporting a resource beside an error.

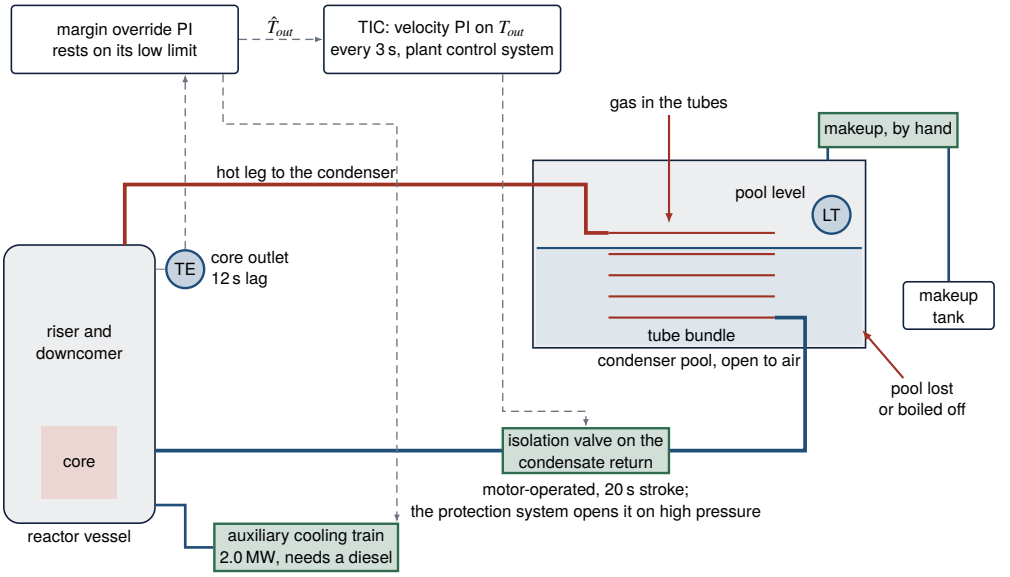


Figure 10.6 The loop of Figure 10.1 with the hardware of a plant. A core outlet thermocouple (TE) feeds the temperature controller TIC, which every 3 s moves the motor-operated isolation valve on the condensate return, and the margin override acts on the auxiliary cooling train from the same reading. In a plant the protection system, not the control loop, starts the condenser. The pool level (LT) is displayed only, and makeup is spent by hand. Thick solid lines carry the hot leg (upper) and the cold condensate return (lower), thin dashed arrows signals.

10.7 Technology in Practice

The symbols of Figure 10.1 stand, in a plant, for a temperature element at the core outlet, a valve on the condensate return of the condenser, an auxiliary cooling train and a makeup supply into the pool. Figure 10.6 draws the loop with that hardware in place of the schematic's symbols.

What Measures It

The chapter's controller reads one temperature. Plants measure temperature mostly with thermocouples and resistance temperature detectors, field devices placed in contact with the item measured or in thermowells that protrude into the fluid [18]. The condenser itself is watched by other instruments. Shell-side level must be kept for the system to operate within its technical specifications, a rising shell temperature points to a leaking tube or condensate return valve, and a radiation monitor sits on the shell vent [15]. The chapter carries the outlet reading through the thermocouple lag τ_{TC} of Equation (10.19) and shows the pool level without noise or lag.

> **TECHNOLOGY** **Sensor: The Core Outlet Thermocouple**
>
> A thermocouple is a junction of two dissimilar metal wires that produces a voltage roughly proportional to temperature. It is more rugged than a resistance thermometer and spans a wider range, though usually less accurately [19]. A thermowell adds a lag, and an air gap at the bottom of the well adds a second one that can destabilize a loop. The common failure is an open circuit. Without burnout detection the instrument then reads electrical noise as a wildly varying temperature, and with it the reading goes to the end of scale chosen as the safer one [19]. The chapter's TE is a 12 s lag with optional noise.

What Moves It

Three final elements act on the heat balance. The isolation valve sits on the condensate return, which is where operating plants start and stop the condenser [20]. Valves in this service are motor-operated gate valves [15], and at Fukushima Daiichi the cooling rate was held by opening and closing them [17], where the chapter modulates its valve continuously. Makeup to the pool comes in existing plants from the condensate, demineralized water or fire water systems, and one plant has gravity-fed tanks that carry the condenser for eight hours [20]. The auxiliary train stands for an active cooling system that needs a diesel, a switchgear line-up and a cooling-water source, as Section 10.3.9 says.

> **TECHNOLOGY** **Final Element: The Condensate Return Valve**
>
> In operating plants the steam supply and condensate return valves of an isolation condenser are motor-operated gate valves of 25.4 to 35.6 cm, the inboard ones on alternating current and the outboard ones on direct current, which can also be worked by hand. On a loss of power they stay where they are. Packing leaks are the most frequent valve failure in service, and valves that bind thermally are cycled during cooldown [15]. Opening the condensate return valve is what starts the condenser [20]. The chapter's isolation valve modulates with a 20 s full stroke, and scenario 7 seizes it where it stands.

Where the Controller Runs

Starting the condenser is a protection function. The reactor protection system supplies the pressure and level signals to the initiation logic, and if automatic initiation fails the operators open the outboard condensate return valve themselves [15]. Safety equipment is kept independent of non-safety equipment, with priority logic so that safety actuations always override non-safety demands [18]. A temperature loop like the chapter's is therefore a non-safety function of a plant control system, whose loops run with cycle times from milliseconds to minutes [21]. The chapter's 3 s control period and 0.6 s physics step sit inside that range. Each of its two loops is one PID block with output limits and back-calculation.

> **IN PRACTICE** **Keep the Standby Condenser Ready**
>
> An isolation condenser waits with its steam valves open and one condensate return valve shut, so its tubes sit at reactor pressure. Vent valves bleed air and other gases to the main steam lines, because gas collected in the loop could stop natural circulation. The system starts when high reactor pressure has lasted 15 seconds, a delay that rides out turbine trips. The shell water then lasts 20 to 90 minutes, depending on the plant, before makeup must keep the tubes covered. The largest contributor to the system's unreliability is the condenser itself, followed by its valves and the instrument air [20].

What Is Modernizing It

Two developments touch this loop, and a third is proposed. Fiber optics are being applied to temperature measurement in nuclear plants, and on-line monitoring can show that a sensor has not drifted since its last calibration, so the calibration interval can be extended [18]. Digital electric actuators with a servo motor and a roller screw calibrate themselves over their whole stroke and take a position demand in digital form, free of the drift of an analog current loop, on modulating and open-close

valves alike [22]. A predictive controller with the pool mass as a state, which Section 10.9 proposes, would compute the inventory trade that the chapter leaves to judgment.

10.8 Limitations

The model is not a safety analysis and cannot be made into one. With five lumped nodes, single-phase buoyancy, a fitted condenser conductance, a well-mixed pool and no relief-valve, containment or hydrogen modeling, it cannot support any claim about a real plant's coping time, peak cladding temperature or accident progression. The natural-circulation phenomena that a system code exists to capture, which are two-phase driving head, flow oscillation, flashing instability and flow reversal, are all outside it [7, 10]. The decay-heat expression is the Way-Wigner form rather than the tabulated standard with its uncertainties [4, 6], so it is correct in shape and approximate in magnitude, and it carries no actinide or structural-activation contribution. The non-condensable term reproduces the direction of an effect that integral experiments resolve in far more detail [14]. Disturbances arrive as commanded steps and timed events rather than as the drifting behavior of a plant, and the repeatability is itself unrealistic, because freezing the integration step and the controller period reproduces one trajectory on every machine while a real plant computer does no such thing.

The four features that survived the simplification all survived because a control decision has to be made around each of them, and they are a source term nobody can move, a manipulated variable whose gain varies by a factor of twenty-five across its travel, an inventory that integrates and never comes back, and a constraint that is one-sided. Section 10.3.9 states where each individual assumption breaks.

> ⚠ **OPEN CHALLENGE** **Proving That a Passive System Will Work**
>
> Passivity is not the same as reliability, and an active design with variable controls can meet a safety function more precisely [12]. The reliability of a natural-circulation system rests on phenomena rather than components: non-condensable gas, thermal and mass stratification, pool heat transfer and others. No commonly accepted practice yet exists for estimating it or for bringing it into a probabilistic safety assessment [10]. The gas fraction and the level-dependent conductance of this chapter are two such phenomena reduced to multipliers. Progress would be an auditable, traceable method that turns the physics into a failure probability [10].

10.9 Discussion

The result that generalizes beyond reactors is the shape of the scorecard. Every architecture in Table 10.5 that looks acceptable on the temperature lines within eight hours is distinguished only by an inventory line, and the two challenge configurations of Table 10.7 that both score zero minutes over the limit differ by 2.97 t of water and 0.70 h of coping time. A controller for this class of plant therefore cannot be graded on an error measure at all. Integrated absolute error keeps accumulating for as long as the run continues and says nothing about whether the resource will last, and the useful measures are the peak approach to the constraint, the resource remaining at a stated horizon, and the time at which the constraint is first violated. The application shows those beside each other rather than combining them, because combining them requires the horizon, and the horizon is the input the operator does not have.

The same reasoning explains why, past a threshold of gain, the correct answer here is a pair of setpoints rather than a set of gains. Both controllers in Section 10.6 are ordinary proportional-plus-integral loops with unremarkable settings. Once the gains are strong enough to open the valve from standby during the first hour, every remaining difference in Table 10.7 comes from where the two setpoints sit relative to T_{lim} and to each other. That is the general character of override and selector structures, whose behavior is decided by the relationship between the limits they are given rather than by their tuning [2]. It also suggests where a formal method would help, because a model predictive formulation with the pool mass as a state, the temperature limit as an inequality and the coping-time requirement as a terminal constraint would compute the trade rather than requiring the operator to

guess it [3], and the closed-form decay integral of (10.20) makes the horizon cheap to evaluate. The corresponding formulation on an islanded microgrid, with the state of charge in place of the pool and a restoration time in place of the coping time, is already standard [1].

Two features of the plant are worth carrying away independently of the control question. The first is that a passive system’s driving force is a feedback loop nobody tuned, and (10.14) shows why it is unconditionally stable and almost entirely without authority at the same time. The second is that the deadline for an inventory-limited system is set by a geometry rather than by a mass, because 16.4 t of the 55.0 t pool lies below the level at which the conductance starts to fall and the last kilogram is never the number that matters. Both properties recur wherever a stored quantity is consumed against a hard limit, as in the reserve vault of Volume 4, Chapter 14, and both are easier to see on a plant whose whole transient fits in eight hours of simulated time than on the systems where they were first learned.

10.10 Conclusion

This chapter treated decay-heat removal by natural circulation into a boiling pool as an interactive laboratory for constraint control against a depleting resource. The simulator integrates fifteen states, which are two flux states, the decay-heat clock, five thermal capacities totaling 372.8 MJ/K, the pool temperature and mass, the makeup tank, the coasting pump head and three actuator positions, with the circulation obtained each step by inverting a loop momentum balance so that the cube-root scaling is a measured consequence at an exponent of 0.33333 rather than a correlation. An hour after the trip the loop carries 147.9 kg/s on 7.26 K of driving temperature difference, which is 4.44 % of the pumped flow, and the isolation valve faces a process gain running from −1.13 K/% wide open to −28.2 K/% at a fifth open. The plant has no setpoint, and the exchange rate that decides every architecture is 165.2 kg of pool water per kelvin of avoided cooldown. A latched-open condenser holds the core outlet at 211.9 °C at four hours and leaves 15.72 t in the pool, throttling to 300 °C leaves 29.92 t at the same hour and moves tube uncovery from 3.93 to 6.26 h, and a margin override on a 2.0 MW auxiliary train leaves 18.19 t at eight hours with the bundle never uncovering. Neither of the intuitive strategies meets the graded challenge, and two passing configurations one setpoint apart finish 2.97 t of pool water and 0.70 h of coping time apart. Further work would add a two-phase driving head, a stratified pool, relief-valve inventory loss above the limit, and a predictive controller that computes the inventory trade against a stated horizon instead of leaving it to judgment. The app lets a reader learn to control a hard limit with a resource that does not come back.

KEY POINTS

- A tripped reactor makes heat that no handle can reduce, and its only heat sink is a pool that boils away.
- The isolation valve has a factor of twenty-five in process gain across its travel and almost no authority over the flow it appears to control.
- Stored heat trades against pool water at 165.2 kg/K.
- Throttling to 300 °C leaves 29.92 t at four hours against 15.72 t passive and moves tube uncovery from 3.93 to 6.26 h.
- The useful measures are the peak approach to the constraint, the resource left at a horizon and the first violation.

Where else this applies. A finite resource spent against a hard limit also sets the reserve vault of Volume 4, Chapter 14 and the battery of Volume 17, Chapter 7, and an override resting on its clamp guards the compressor of Volume 4, Chapter 2.

References

[1] A. Parisio, E. Rikos, and L. Glielmo, “A model predictive control approach to microgrid operation optimization,” *IEEE Transactions on Control Systems Technology*, vol. 22, no. 5, pp. 1813–1827, 2014.

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

[3] 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.

[4] K. Way and E. P. Wigner, "The rate of decay of fission products," *Physical Review*, vol. 73, no. 11, pp. 1318–1330, 1948.

[5] A. Tobias, "Decay heat," *Progress in Nuclear Energy*, vol. 5, no. 1, pp. 1–93, 1980.

[6] American Nuclear Society, "Decay heat power in light water reactors," ANSI/ANS-5.1-2014 (R2023), La Grange Park, IL, USA, 2014, reaffirmed 2023.

[7] Y. Zvirin, "A review of natural circulation loops in pressurized water reactors and other systems," *Nuclear Engineering and Design*, vol. 67, no. 2, pp. 203–225, 1982.

[8] P. K. Vijayan and H. Austregesilo, "Scaling laws for single-phase natural circulation loops," *Nuclear Engineering and Design*, vol. 152, no. 1–3, pp. 331–347, 1994.

[9] P. K. Vijayan, "Experimental observations on the general trends of the steady state and stability behaviour of single-phase natural circulation loops," *Nuclear Engineering and Design*, vol. 215, no. 1–2, pp. 139–152, 2002.

[10] International Atomic Energy Agency, "Natural circulation in water cooled nuclear power plants: Phenomena, models, and methodology for system reliability assessments," International Atomic Energy Agency, Vienna, Austria, Tech. Rep. IAEA-TECDOC-1474, 2005.

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

[12] International Atomic Energy Agency, "Safety related terms for advanced nuclear plants," International Atomic Energy Agency, Vienna, Austria, Tech. Rep. IAEA-TECDOC-626, 1991.

[13] T. L. Schulz, "Westinghouse AP1000 advanced passive plant," *Nuclear Engineering and Design*, vol. 236, no. 14–16, pp. 1547–1557, 2006.

[14] D. Paladino, O. Auban, M. Huggenberger, and J. Dreier, "A PANDA integral test on the effect of light gas on a passive containment cooling system (PCCS)," *Nuclear Engineering and Design*, vol. 241, no. 11, pp. 4551–4561, 2011.

[15] R. D. Orton, "Pre-phase I aging assessment of the BWR isolation condenser system," Pacific Northwest Laboratory, for the U.S. Nuclear Regulatory Commission, Richland, WA, Tech. Rep., Aug. 1995.

[16] F. Tanabe, "Analysis of core melt accident in Fukushima Daiichi-Unit 1 nuclear reactor," *Journal of Nuclear Science and Technology*, vol. 48, no. 8, pp. 1135–1139, 2011.

[17] International Atomic Energy Agency, *The Fukushima Daiichi Accident*. Vienna, Austria: International Atomic Energy Agency, 2015, sTI/PUB/1710, report by the Director General with technical volumes 1–5.

[18] ——, *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.

[19] 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.

[20] Z. Ma, K. Kvarfordt, and T. Wierman, "System study: Isolation condenser 1998–2022," Idaho National Laboratory, Idaho Falls, ID, Tech. Rep. INL/RPT-24-77016, Dec. 2023.

[21] 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.

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

10.11 Exercises

10.1 ○ Concept check · 2 min

The reactor of this chapter has tripped, the pumps have coasted down and the isolation condenser is the only heat sink. What is the control task for the next eight hours?

(a) Cool the plant as fast as the cooldown-rate limit allows, so that the temperature margin grows
(b) Keep the core outlet below its limit and make the pool water last the required hours
(c) Hold the decay heat at a setpoint by moving the isolation condenser valve
(d) Keep the condenser valve shut so that no pool water is boiled before it is needed

10.2 Concept check · 2 min

In the Passive + aux architecture the isolation valve holds a core outlet setpoint of 300 °C, 16 K below the 316 °C limit. The override on the auxiliary train is set to a margin of 10 K. What does the train do?

(a) It runs continuously, because its setpoint is below the margin the valve holds
(b) It runs at a fixed share of its rating, set by the ratio of the two setpoints
(c) It stays off, and loads only when a disturbance drives the margin toward 10 K
(d) It stays off whatever happens, because back-calculation holds its integral at zero

10.3 Concept check · 1 min

In scenario 10 non-condensable gas collects in the condenser bundle while the throttled loop holds the core outlet at 300 °C. Where does an operator first see that the heat sink is failing?

(a) In the core outlet temperature, which rises in step with the gas fraction
(b) In the pool level, which falls faster as the conductance is lost
(c) In the loop flow, which stops once the gas has blanketed the bundle
(d) In the valve position, which opens toward its stop to hold the setpoint

10.4 Theory and calculation · 6 min

The reactor of this chapter is rated $P_0 = 530\,\mathrm{MW}$. Use Equation (10.1) with $A = 0.066$ and $n = 0.2$, and $h_{fg} = 2257\,\mathrm{kJ/kg}$.

(a) Find the decay-heat fraction and the decay power one hour after the trip for a core that ran for a year ($T_0 = 8760\,\mathrm{h}$) before it. *Checkpoint:* 1.074 % and 5.694 MW.

(b) Repeat for a core that ran for four hours, and give the result as a share of part (a). *Checkpoint:* 0.3532 %, which is 32.9 % of the year-long value.

(c) A saturated pool takes the decay heat of part (a). Find its boil-off rate in kilograms per second and in tonnes per hour. *Checkpoint:* 2.52 kg/s and 9.08 t/h.

10.5 Theory and calculation · 6 min

The five thermal capacities of the model are $C_f = 7.26$, $C_c = 82.7$, $C_r = 83.6$, $C_{ic} = 33.2$ and $C_d = 166.0\,\mathrm{MJ/K}$, and $h_{fg} = 2257\,\mathrm{kJ/kg}$.

(a) Find the plant's total sensible capacity and its exchange rate into pool water. *Checkpoint:* 372.76 MJ/K and 165.2 kg/K.

(b) The fully passive run cools the core outlet from 300.0 to 211.9 °C in four hours. Treat that as a bulk cooldown of the whole plant and estimate the pool water it spends. *Checkpoint:* 32.84 GJ, or 14.55 t.

(c) The 2.0 MW auxiliary train runs at full output. Find the pool water it saves per hour, and the saving for the 56.71 GJ it takes over eight hours in the third column of Table 10.5. *Checkpoint:* 3.19 t per hour and 25.1 t.

10.6 Theory and calculation · 8 min

An hour after the trip the pumps have stopped and the bypass resistance is $R_{by,s} = 0.65\,\mathrm{Pa\,s^2/kg^2}$. Take $g = 9.81\,\mathrm{m/s^2}$, $\rho = 740\,\mathrm{kg/m^3}$, $\beta = 3.0 \times 10^{-3}\,\mathrm{K^{-1}}$, $\Delta z = 8.0\,\mathrm{m}$, $c_p = 5.3\,\mathrm{kJ/(kg\,K)}$, $R_{pipe} = 0.014$ and $R_{ic0} = 0.08\,\mathrm{Pa\,s^2/kg^2}$, and a decay power of 5.694 MW.

(a) Find the buoyancy gain $g\rho\beta\Delta z$ and the loop resistance R of Equations (10.9) to (10.11) with the valve wide open. *Checkpoint:* 174.2 Pa/K and 0.0578 Pa s^2/kg^2.

(b) Find the circulation from Equation (10.14) and the temperature difference around the loop. *Checkpoint:* 147.9 kg/s and 7.26 K.

(c) Repeat part (b) with the valve half open, and compare the change in flow with the change that doubling the decay power would give. *Checkpoint:* 114.6 kg/s. Wide open carries 29.1 % more, and doubling the power gives 26.0 %.

10.7 App challenge · 10 min

Load scenario 4, *Fully passive: it works, and look what it costs*, change no setting, and run it in Auto for 480 min at 200x with the default seed. Export the CSV file, whose time column is Time (h). Then load scenario 5, *Throttled passive: refusing to cool down*, and do the same.

(a) In each file read the pool inventory, column `Mp_t`, and the core outlet, column `Tout_C`, in the row at 4.0 h. *Checkpoint:* 15.72 t and 211.9 °C passive, 29.92 t and 299.8 °C throttled, within 0.01 t and 0.1 °C.

(b) In each file find the first row in which the submerged fraction of the bundle, column `submergence`, is below 0.999. *Checkpoint:* 3.93 h passive and 6.27 h throttled, within 0.01 h.

(c) Read the stored-heat deficit, column `metric_Ecool_MJ`, in the 4.0 h row of each file. Divide the difference by $h_{fg} = 2257\,\mathrm{kJ/kg}$ and compare the result with the difference in `Mp_t` from part (a). *Checkpoint:* 32 857 MJ and 73 MJ, so 14.53 t predicted against 14.20 t measured.

10.8 App challenge · 8 min

Load scenario 8, *The diesel does not start*, change no setting, and run it in Auto for 480 min at 200x with the default seed. The auxiliary train is unavailable, and the setpoint pair is 300 °C and 10 K.

(a) Read Peak core outlet, Minimum margin and Over limit on the scorecard at 480 min. *Checkpoint:* 318 °C, −2.0 K and 7.5 min.

(b) Export the CSV file and read `Mp_t` in the last row, at 8.0 h. *Checkpoint:* 6.25 t, within 0.01 t.

Compare the four values with the column of Table 10.6 for the same pair with the train available, and state what the train was worth at this setpoint pair.

10.9 App challenge · 10 min

Load the challenge, *Challenge: Eight hours on the water that is in the tank*. Run each case below for 480 min at 200x with the default seed. For each run read Peak core outlet and Over limit on the scorecard, and the pool inventory `Mp_t` in the last row of the exported file. Accept an Over limit within 0.1 min and an inventory within 0.01 t.

(a) The case as loaded is the Passive column of Table 10.7. Read the same three values there instead of running it. *Checkpoint:* 342 °C, 69.9 min and 4.81 t.

(b) Set Heat-removal architecture to Passive + aux, Core outlet temperature setpoint to 306 °C and Temperature loop gain to 3.00 %/K with an integral time of 6.0 min. In the Advanced view set Margin at which auxiliary cooling comes in to 14 K and Auxiliary override gain to 4.0 %/K with an integral time of 6.0 min. Reveal suggested tuning applies the same settings. Press Reset and run. Also read Auxiliary energy. *Checkpoint:* 313 °C, 0.0 min, 11.56 t and 28.3 GJ.

(c) Set Temperature loop gain to 1.20 %/K, press Reset and run again. *Checkpoint:* 319 °C, 41.1 min and 11.14 t.

10.10 Application area · 4 min

The makeup tank holds 12 t and can be spent into the pool once. The decay heat of the year-long core is 5.694 MW one hour after the trip and 3.646 MW six hours after it, and $h_{fg} = 2257\,\mathrm{kJ/kg}$.

(a) Counting latent heat only, find how many hours of boil-off the tank replaces if it is spent at each of the two times. *Checkpoint:* 1.32 h and 2.06 h.

(b) Say why spending it later is worth more, and what sets the latest moment at which it is still useful.

10.11 Application area · 3 min

In scenario 4 the core outlet falls from 300.0 to 211.9 °C in four hours. A shift supervisor argues that the isolation valve must be throttled to protect the vessel from the 55 K/h technical-specification cooldown rate.

(a) Find the average cooldown rate over the four hours. *Checkpoint:* 22.0 K/h.

(b) Using the results of the chapter, say whether the supervisor's argument holds on this plant, and give the argument that does.

10.12 Review the agent · 5 min

An AI assistant reviews a passing answer to the challenge (Passive + aux, 306 °C and 14 K, temperature loop gain 3.00 %/K) and replies as follows. "With 14 K the train runs at full output for 451 of the 480 min. Lower the override setpoint to 10 K. Both settings meet the target with no minutes over the limit and peaks of 313.2 and 313.5 °C, so nothing is lost, and the train is held in reserve for the end of the shift, when the non-condensable gas is worst." Find the flaw in the proposal and say how it could be checked.