INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 1 Process Dynamics, Modeling, and Identification

CHAPTER 9

Spent-Fuel Pool Simulator for Integrating Inventories and Time-to-Limit Prediction

Pages 204–233

Edited by John D. Hedengren

How to cite this chapter. "Spent-Fuel Pool Simulator for Integrating Inventories and Time-to-Limit Prediction," in *Process Dynamics, Modeling, and Identification*, John D. Hedengren, Ed., vol. 1 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 204–233.

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

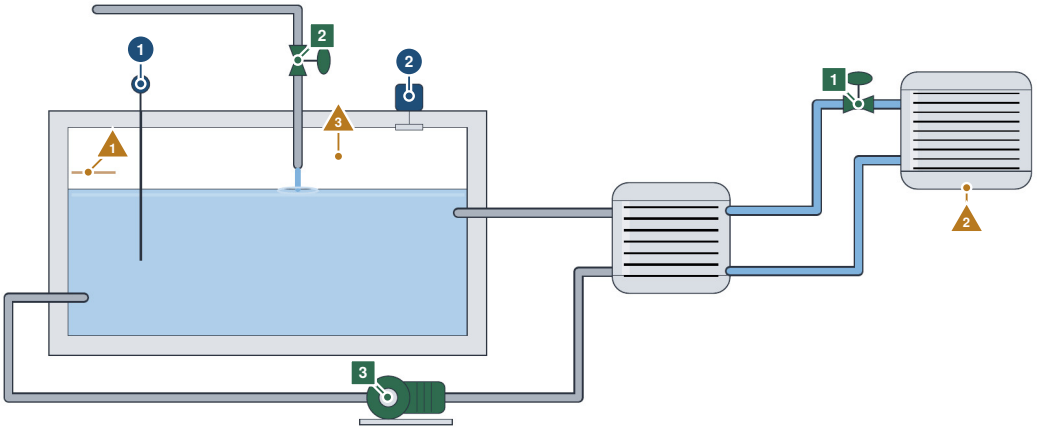
The app for this chapter
https://apmonitor.com/apps/spent-fuel-pool/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v01-dynamics-modeling/#spent-fuel-pool

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

CHAPTER 9

Spent-Fuel Pool Simulator for Integrating Inventories and Time-to-Limit Prediction



The controller moves 1 cooling water valve 2 makeup valve 3 cooling train pump

The controller reads 1 pool temperature 2 pool level

What disturbs it 1 water lost to evaporation or a leak 2 cooling water temperature 3 air over the pool

Notice the two separate paths: the cooling train carries heat out to the cooling water, while only the makeup line puts back the water the pool loses. Illustration, not to scale.

IN PLAIN TERMS **A Warm Pool That Can Run Out of Time**

The main idea. Used nuclear fuel keeps making heat for years, so it is stored under deep water in a pool. Pumps send the water through a cooler, and a valve adds water to replace what slowly evaporates.

The challenge. If the pumps stop, the pool warms so slowly that nothing seems wrong for hours. Warm water also swells, so the level can rise even while water is being lost.

Why it matters. The useful number is how many hours are left before the pool boils, and working it out step by step beats stretching a straight line.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of the dense racks and the fact on real pools (Section 9.2), the try-it of scenario 3 (Section 9.4), the fact on feedforward and the try-it of scenario 10 (Section 9.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question The level of a spent-fuel pool is a pure integrator, and its temperature is a self-regulating first-order plant only while a cooling train runs. How does one step test tell which plant the pool is, and what should an operator compute when cooling is lost and the useful number is the time left to a limit?

Learning objectives

- Explain from the two balances why the level integrates and why the temperature turns into an integrator when both trains stop.
- Compute the temperature time constant and the heat-up rate from the thermal capacity and the feedback of the exchanger and the surface.
- Identify a first-order and an integrating model from one step test, and state the tuning each one commits the loop to.
- Design an averaging level loop, and predict its offset when the disturbance is a leak rather than evaporation.
- Compare a marched time to a limit with a present-rate extrapolation, and feedforward from computed evaporation with level feedback.

Prerequisites Sections 1.4.2 and 1.4.3 (balances and linearization), 1.4.4 (an integrator as a long lag), 1.4.6 (effective dead time) and 1.6.2 (velocity form). Chapter 2 shows a tank that regulates itself.

The app https://apmonitor.com/apps/spent-fuel-pool/

See also Volume 6, Chapter 4 for averaging level control and fouling that walks toward a limit, Volume 6, Chapter 10 for leak detection by mass balance. Volume 18, Chapter 10 for decay heat on a pool that boils, Volume 18, Chapter 2 for regulation against self-regulation.

9.1 Introduction

A slow process is usually assumed to be a forgiving one. A pool of 1231 m^3 of borated water over 448 discharged fuel assemblies rejects 2148 kW through two cooling trains while losing 39.5 kg/h of water to evaporation, which is 0.37 mm/h of level. If both trains stop the pool warms at 1.49 °C/h. Nothing on the trend display moves during the first shift. The assumption that follows from that appearance is wrong twice over. It is wrong because an integrating plant does not forgive an error, it accumulates one, so the deviation at any moment is the imbalance multiplied by however long nobody acted. It is wrong again because the two rates just quoted are the rates now, and a straight line drawn through the last hour of either trend predicts the next ten badly. The pool boils 1.1 h later than the temperature ramp says it will, and reaches its shielding limit 21.9 h into a loss of cooling whose first 8.8 h carry a level trend that is rising.

A simulator supplies the part a transfer function cannot. The plant in question is not one plant. With a train running the exchanger duty rises with pool temperature at 423 kW/K, which is real negative feedback, and the pool is first order with a time constant of 3.33 h. With both trains stopped only 4.20 kW/K of surface loss remains and the time constant is 339 h, which no controller and no operator can use as a time constant. A student handed either model separately learns a correct thing about a plant that does not exist by itself, and a student handed both learns the thing that matters, which is that the identification depends on which equipment happens to be in service and that operating experience built during normal operation is the wrong preparation for the day the pumps stop. Surveys of what control graduates are short of keep returning to the same gap, which is practice at getting a model out of a plant instead of being handed one in a problem statement [1, 2].

The general class is broader than a nuclear plant. It is an integrating plant whose apparent dynamics depend entirely on whether its dominant negative feedback path is in service, carrying two coupled inventories, and for which the operative quantity during an upset is a time to a limit rather than an error. Large tank farms and surge inventories behave this way whenever the discharge is set by a pump rather than by hydrostatic head, and level loops on them are detuned on purpose so that the inventory absorbs upsets instead of passing them downstream [3]. Building thermal mass behaves

this way when the plant is off, which is why predictive control of a building is written against a stored-energy state and a forecast rather than against a temperature error [4]. The state of charge of a battery is the same integrator with a different unit. Tuning rules for integrating plants with dead time make the dead time the whole of the denominator of the gain [5], and the identification that supplies that dead time is a slope and a delay rather than a gain and a settling point.

The second difficulty is that the two inventories are coupled by a mechanism that appears on neither loop's block diagram. Evaporation from the pool surface takes heat and water away together, it is driven by the vapor-pressure difference to air on the refueling floor that nobody controls, and it grows steeply with pool temperature, from 39.5 kg/h at 35 °C to 217 kg/h at 55 °C and 1864 kg/h as the surface approaches saturation. The same disturbance is negligible in the energy balance and dominant in the mass balance, so the loop that could reject it does not need to and the loop that needs to cannot. However, a level measurement on a pool of 108 m^2 of surface is a very poor flowmeter, because 39.5 kg/h of evaporation is 8.8 mm of level a day against 3 mm of transmitter noise.

This chapter works with an interactive spent-fuel pool in which the same temperature measurement identifies two different plants depending on what is running, and in which the inventory loop is answered by computing a disturbance rather than by tuning a controller. It measures how far a marched time-to-limit calculation beats a present-rate extrapolation on this plant, and why the first-day level deviation under feedback contains no controller gain at all. The pool model and the tuning arithmetic are standard. Section 9.2 traces how at-reactor pools came to hold what they hold and how their instruments were judged after 2011. Section 9.3 sets out the seven states, the 31 constants and the two instrument chains closely enough that the pool can be rebuilt from the page, with decay heat summed over a mixed-age inventory and evaporation from a published mass-transfer correlation. Section 9.4 linearizes both balances and identifies the pool twice from one step test. Section 9.5 states the four control architectures, the tuning rules the identification commits to, the evaporation feedforward and the two time-to-limit calculations. Section 9.6 measures the twelve scenarios, from a pool in which nothing appears to happen to a graded challenge. Section 9.7 sets the model beside the instruments and valves of real pools, Section 9.8 states the limitations, Section 9.9 weighs makeup water against shielding and Section 9.10 draws the results together.

9.2 Background

The pool was never intended to hold what it now holds: water storage at the reactor was designed as a short buffer and became a long-term store, as the box below traces, and original open racks were replaced by high-density racks analyzed as a criticality problem in their own right [6], and the hazard of the dense inventory was reviewed by the National Academies [7].

> ⧖ HISTORY **Sized for Reprocessing, Packed Instead (1974 to 1982)**
>
> Pools in the United States, like those in Europe and Japan, were first sized to hold spent fuel for only a few years, until it was cool enough to travel to a reprocessing plant that would recover its plutonium and uranium. After India tested a nuclear explosive made with plutonium recovered for peaceful purposes in 1974, the United States halted in 1977 the licensing of an almost completed reprocessing plant. A 1982 act promised a deep underground repository that was to begin accepting fuel by 1998. Operators packed their pools instead. Racks once spaced in a loose square array gave way to dense-pack racks at 23 cm, barely above the 21.4 cm of a reactor core, with each assembly in a box whose walls hold neutron-absorbing boron, and virtually every pool in the country was re-racked [8].

Thermal analysis of the pool followed the same path from short buffer to long-term store. The early treatment is a lumped energy balance on a well-mixed pool against decay heat, with evaporation and surface loss as the only passive sinks, extended to the local temperature distribution around the racks and to the natural circulation that carries heat out of them [9]. Decay heat itself is the forcing function and it is the one quantity the analysis cannot choose, being fixed by burnup and cooling time,

and modern practice validates decay heat power for light water reactor fuel against depletion codes rather than against a single correlation [10]. Evaporation received much less attention than decay heat until it mattered, and the correlations available for undisturbed water surfaces were assembled and tested against a wide range of pools, including spent-fuel pools, only recently [11, 12].

Cooling and makeup requirements evolved from the licensing basis rather than from control practice. Two trains, each a pump and a plate exchanger to a component cooling water system, remove decay heat and hold the bulk temperature below a technical-specification limit. Makeup replaces evaporation, is very often started by hand, and is sized generously because the case it has to cover is a boiling pool rather than an evaporating one. The depth of water above the fuel is not there to cool anything at all. Seven meters of water over the tops of the assemblies is shielding, and the technical-specification minimum in this application is 6.40 m, so the level constraint that governs the inventory loop is a dose-rate constraint and not a heat-transfer constraint.

> ✦ **DID YOU KNOW? Real Pools Are Built to This Scale**
>
> Pools in light water reactors are typically about 12 m deep, with approximately 7 m of water over the stored fuel. A 2006 French assessment found that during a refueling outage a total loss of cooling could bring a pool to boiling within five to eight hours [13].

March 2011 changed what was measured. Loss of cooling at Fukushima Daiichi left operators without a reliable indication of pool level or pool temperature at exactly the point where those were the only quantities that mattered, and the response in the years that followed was directed at instrumentation and at inventory rather than at cooling capacity. The National Academies review of the accident treats pool instrumentation and the ability to add water by alternate means as central findings [14]. Analysis of pool behavior during a station blackout became a question about the joint history of temperature and level rather than about a peak temperature [15], and the evaporation rate from the damaged pools was reconstructed from a mass-transfer model to check the decay heat estimates against observation [16]. That reordering of priorities is the reordering this application asks a student to make. The temperature trend goes flat at 100 °C and stops carrying information, at which point the level is the only measurement left and the useful number is how many hours remain rather than how many degrees the pool is off setpoint.

9.3 Pool Model

The pool is 11.40 m deep with 108 m^2 of free surface, holding 1231 m^3 of borated water at 35.00 °C over racks whose tops are at 4.40 m, so 7.00 m of water covers the fuel. The design inventory is 448 assemblies, six batches of 64 at 3.0, 4.5, 6.0, 7.5, 9.0 and 10.5 y producing 682.7 kW between them, plus a batch of 64 thirty days out of the reactor producing 1465.0 kW, for a total decay heat of 2147.7 kW. Two cooling trains take suction below the surface, pass the water through a plate exchanger to component cooling water at 30.0 °C and return it to the pool. A makeup line replaces evaporation. Figure 9.1 shows the arrangement and every symbol that appears in the equations.

Table 9.1 classifies what the two loops interact with. Four limits bound normal operation, a high pool temperature alarm at 52 °C, a technical-specification temperature limit at 60 °C, a low-level alarm at 6.70 m of shielding and the technical-specification shielding minimum at 6.40 m. None of them stops a run. The two alarms color the badges and annunciate on the status bar while the condition holds, and the two technical-specification limits also mark the trend when they are crossed and total the hours spent outside them. The plant is expressed in the units an operator reads, temperatures in °C, masses in kg, mass flows in kg/h, volumes in m^3, volume flows in m^3/h, duties in kW, elevations in m and valve positions in percent. **Time is in hours throughout.** Seven states are integrated.

🕮 **Pause.** Try Exercise 9.2 before reading on.

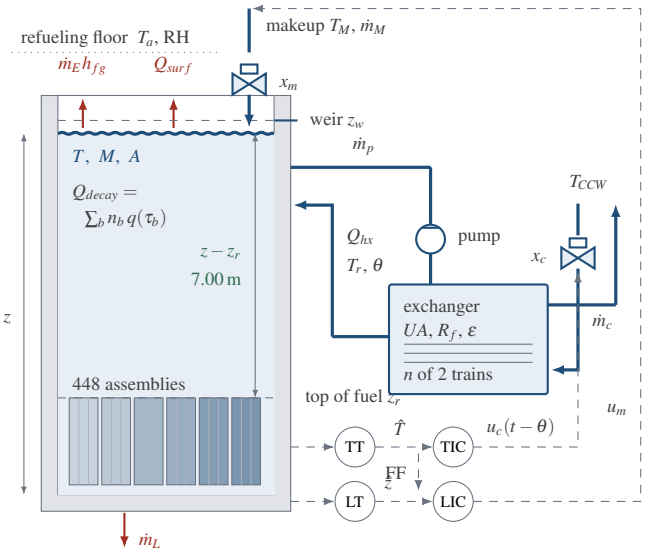


Figure 9.1 Spent-fuel pool with its cooling train and makeup line. Water of mass M at bulk temperature T stands at elevation z over racks whose tops are at z_r, so $z - z_r$ is the shielding depth the license is written around. Decay heat Q_{decay} is summed over batches of n_b assemblies at age τ_b. Rack shading darkens as batch age falls, with the freshest batch darkest and at the right. One of n trains draws $\dot{m}_p$ through a plate exchanger of conductance UA and fouling resistance R_f, returning water at T_r after a transport delay θ and rejecting Q_{hx} to component cooling water at T_{CCW} through butterfly valve x_c. Evaporation $\dot{m}_E$ and surface loss Q_{surf} leave through the free surface of area A against refueling floor air at T_a and relative humidity RH, makeup at T_M enters through equal-percentage valve x_m, and any liner leak $\dot{m}_L$ leaves at the floor. Water above the weir z_w spills. The makeup command u_m carries the evaporation feedforward computed from $\hat{T}$, T_a and RH.

9.3.1 Energy and Mass Balances

The two accumulations are Equations (9.1) and (9.2), written as Section 1.4.2 writes them for a heated tank.

$$Mc_p \frac{\mathrm{d}T}{\mathrm{d}t} = Q_{decay} - Q_{hx} - Q_{surf} - \dot{m}_E h_{fg} + \dot{m}_M c_p (T_M - T) \tag{9.1}$$

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \dot{m}_M - \dot{m}_E - \dot{m}_L - \dot{m}_{boil} \tag{9.2}$$

where M is the pool water mass in kg, $c_p = 4.182\,\mathrm{kJ/(kg\,K)}$, T is the bulk temperature, Q_{decay} is decay heat, Q_{hx} is the duty of the running trains, Q_{surf} is convection plus radiation from the free surface, $\dot{m}_E$ is evaporation, h_{fg} is the latent heat, $\dot{m}_M$ is makeup at temperature T_M, $\dot{m}_L$ is a liner leak and $\dot{m}_{boil}$ is boil-off. Both sides of (9.1) carry kJ/h once the duties are multiplied by 3600 s/h, and both sides of (9.2) carry kg/h. At the design point $Mc_p = 5.120\,\mathrm{GJ/K}$.

Nothing on the right-hand side of (9.2) is a function of M, and that absence is the cleanest integrating structure in the plant. A makeup valve one percent from the balancing position produces a level ramp that never stops and never settles anywhere. The application verifies the property directly by evaluating $\mathrm{d}M/\mathrm{d}t$ at two elevations 1.9 m apart and requiring the two values to agree to within 1×10^{-9} kg/h. Equation (9.1) is different, because Q_{hx}, Q_{surf} and $\dot{m}_E$ all rise with T, and Section 9.4 quantifies how much of that self-regulation survives when the pumps stop.

Table 9.1 What the Two Loops Interact With. The Nominal Column Is the 35 °C Design Point the Pool Opens At.

Symbol	Role	Nominal	Range
T (TT)	CV, measured	35.00 °C	20–60 °C SP
$z - z_r$ (LT)	CV, measured	7.00 m	6.40–7.20 m SP
x_c (CCW)	MV	44.20 %	0–100 %
x_m (makeup)	MV	2.37 %	0–100 %
n trains	MV, discrete	2	0, 1, 2
offload	DV, measured	64 assemblies	0–193
trains available	DV, measured	2	0–2
T_{CCW}	DV, measured	30.0 °C	15–45 °C
T_a	DV, measured	30.0 °C	10–45 °C
RH	DV, measured	45 %	10–95 %
T_M	DV, measured	20.0 °C	5–35 °C
ventilation	DV, unmeasured	1×	0.4–2×
R_f fraction	DV, unmeasured	0.15×	0–1×
$\dot{m}_L$ leak	DV, unmeasured	0 m³/h	0–5 m³/h
T	state	35.00 °C	$\le T_{sat}$
M	state	1.2243×10^6 kg	$\ge \rho A z_r$
T_r	state	32.33 °C	—
x_c, x_m	states	44.20, 2.37 %	0–100 %
R_f	state	9.0×10^{-5} K/kW	$0–6 \times 10^{-4}$ K/kW
T_{cw}	state	30.0 °C	—

Level and shielding follow from the mass through the water density, as shown in Equation (9.3).

$$z = \frac{M}{\rho(T)A}, \qquad \rho(T) = 1001.6 - 0.0872\,T - 0.00337\,T^2 \tag{9.3}$$

where z is the surface elevation above the pool floor in m, $A = 108\,\text{m}^2$ is the free surface area and ρ is in kg/m³ from a quadratic fit good to 1.5 kg/m³ between 20 and 100 °C. The shielding depth is $z - z_r$ with $z_r = 4.40\,\text{m}$. Because ρ depends on T, differentiating (9.3) gives the level rate in Equation (9.4).

$$\frac{\mathrm{d}z}{\mathrm{d}t} = \frac{1}{\rho A}\frac{\mathrm{d}M}{\mathrm{d}t} - \frac{z}{\rho}\frac{\mathrm{d}\rho}{\mathrm{d}T}\frac{\mathrm{d}T}{\mathrm{d}t} \tag{9.4}$$

The second term is positive for a warming pool because $\mathrm{d}\rho/\mathrm{d}T$ is negative, so a heating pool rises while it is losing water. That term is the reason a level transmitter cannot distinguish thermal expansion from inventory and the reason a present-rate extrapolation of the level gets the sign wrong early in a loss of cooling. A level that rises while the vessel loses water is the swell that Chapter 6 meets in a flash vessel, there from vapor held below the surface and here from the density of warming water.

Above the skimmer weir at $z_w = 11.60\,\text{m}$ the surface is pinned and the excess leaves as spill, as shown in Equation (9.5).

$$\dot{m}_{spill} = \max\left(0,\ \frac{\mathrm{d}M}{\mathrm{d}t} - A z_w \frac{\mathrm{d}\rho}{\mathrm{d}T}\frac{\mathrm{d}T}{\mathrm{d}t}\right) \quad \text{when } z \ge z_w \tag{9.5}$$

The subtracted term is the mass the pinned surface must shed as the water expands, so heating alone sends water over the weir while the inventory is falling. Equation (9.5) is also what makes feed and spill work as an emergency heat sink, because makeup can be run far above the loss rate and the surplus leaves over the weir.

At saturation the pool cannot get hotter, so any remaining surplus leaves as steam, as shown in Equation (9.6).

$$\dot{m}_{boil} = \frac{M c_p}{h_{fg}(T)} \left.\frac{\mathrm{d}T}{\mathrm{d}t}\right|_{\text{unclamped}}, \qquad \frac{\mathrm{d}T}{\mathrm{d}t} = 0 \quad \text{when } T \geq T_{sat} \tag{9.6}$$

with $T_{sat} = 100\,^{\circ}\mathrm{C}$. Past this point the temperature trend is flat and the level trend carries the whole of the information.

9.3.2 Decay Heat Over a Mixed-Age Inventory

Decay power per assembly is a three-term fit, as shown in Equation (9.7).

$$q(\tau) = 88.08\,\tau^{-0.5277} + 7.975\,\mathrm{e}^{-\tau/188.6} + 1.491\,\mathrm{e}^{-\tau/1160} \tag{9.7}$$

where q is in kW per 17×17 pressurized water reactor assembly and τ is days since shutdown. The power law carries the fission products in aggregate and the two exponentials carry the short-lived and medium-lived groups a single power law cannot. Equation (9.7) returns 40.71 kW at 7 d, 22.89 kW at 30 d, 6.15 kW at 1 y, 1.98 kW at 5 y and 0.81 kW at 20 y, and it is a least-squares fit to decay powers of the kind the ANSI/ANS-5.1 standard produces for a burnup of 45 GWd/tU, with a worst residual near 5 % over that range. On logarithmic axes it falls with a mean slope of -0.56 between one week and twenty years, steepening from -0.40 over the first month to -0.70 between one and five years, which is why a pool remains a heat source for decades. The same power-law clock drives the shut-down reactor of Volume 18, Chapter 10, whose only heat sink is a pool that boils to atmosphere. The total is summed over batches and each batch ages as the simulation advances, as shown in Equation (9.8).

$$Q_{decay}(t) = \sum_{b=0}^{n_{aged}-1} n_b\, q(\tau_b + t) + n_{fresh}\, q\big(\tau_{fresh} + t\big) \tag{9.8}$$

A full core of 157 assemblies eight days out of the reactor raises the total from 2148 kW to 6731 kW, and the same core twelve days out gives 5816 kW.

9.3.3 Evaporation and Surface Loss

Evaporation is the natural-convection correlation for undisturbed water surfaces, in which the driving force is the buoyancy of the humid layer standing over the pool and the potential is a humidity-ratio difference, as shown in Equation (9.9) [11, 12].

$$E'' = C_E\, \rho_w\, (\rho_a - \rho_w)^{1/3}\, (W_w - W_a), \qquad C_E = 35 \tag{9.9}$$

where E'' is in kg/(m^2 h), ρ_w is the air density at the water surface temperature and saturated, ρ_a is that of the room air, and W_w and W_a are the corresponding humidity ratios, $W = 0.622\, p_v/(p_{atm} - p_v)$. Both densities are the mass of dry air per unit volume of moist air, which is the density a psychrometric chart carries and the definition the published model uses, so each is evaluated as $\rho = (p_{atm} - p_v)/(0.287\, T_K)$ and the vapor partial density is not added to it. The coefficient $C_E = 35$ belongs to this set of units and to no other. Saturation pressure comes from Buck's correlation, accurate to 0.05 % below 100 °C [17], as shown in Equation (9.10).

$$p_{sat}(T) = 0.61121\, \exp\left[\left(18.678 - \frac{T}{234.5}\right)\frac{T}{257.14 + T}\right] \tag{9.10}$$

with p_{sat} in kPa and T in °C. Total evaporation is $\dot{m}_E = E'' A f_v$ with f_v the ventilation multiplier, and the correlation is evaluated at a surface temperature capped 5 °C below saturation because W_w diverges there. The underlying model has been tested against 24 studies covering pool areas from 0.022 to 525 m^2, water temperatures from 7 to 94 °C, air temperatures from 6 to 51 °C, relative humidities from 16 to 98 % and air velocities to 1.9 m/s, with a mean absolute deviation of 13.9 % [12], and it

has been applied to spent-fuel pools specifically [16]. The pool of this application sits inside that range on every variable except that the boiling branch of (9.6) takes over above 95 °C, one degree above the highest water temperature in the validation set.

One departure from the published model has to be stated, because the evaporation term drives the whole of the loss-of-cooling result. The published model returns the larger of (9.9) and a vapor-pressure floor of $5 \times 10^{-5}\,(p_w - p_a)$ in the same units with the pressures in Pa, and the application omits the floor. Evaluated over the whole operating envelope of this pool the floor is always the smaller of the two, at 0.19 kg/(m^2 h) against 0.37 kg/(m^2 h) at the design point, 0.10 kg/(m^2 h) against 0.15 kg/(m^2 h) in a hot and humid building and 4.13 kg/(m^2 h) against 17.26 kg/(m^2 h) at 95 °C, so omitting it changes no result in this chapter. The density definition, which is the one place an implementation of this correlation most easily goes wrong, was checked term by term against the published nomenclature and agrees with it, so the rates below are the rates the correlation gives for this pool rather than an implementation of something near it.

Equation (9.9) as implemented returns 39.5 kg/h at the design point and 104.2 kg/h at 45 °C, 217.2 kg/h at 55 °C, 705.1 kg/h at 75 °C and 1864 kg/h at 95 °C, which is a doubling for every 9.6 °C averaged between 35 and 75 °C. The latent heat is a linear fit, $h_{fg} = 2500.8 - 2.36\,T$ in kJ/kg, within 9 kJ/kg over 20 to 100 °C.

Sensible loss from the surface is convection plus gray-body radiation, as shown in Equation (9.11).

$$Q_{surf} = A\left[\frac{h_c(T - T_a)}{1000} + \varepsilon\sigma\left(T_K^4 - T_{a,K}^4\right)\right] \tag{9.11}$$

where $h_c = 3.5\,\mathrm{W/(m^2\,K)}$, $\varepsilon = 0.95$, $\sigma = 5.67 \times 10^{-11}\,\mathrm{kW/(m^2\,K^4)}$ and T_K denotes absolute temperature. At the design point (9.11) returns 5.21 kW against 26.56 kW of evaporative duty, so the surface removes 31.8 kW of the 2148 kW total.

9.3.4 Cooling Train and Makeup Valve

Each train is an effectiveness-NTU counter-flow exchanger with the fouling resistance in series and a turbulent film coefficient on the component cooling water side, as shown in Equations (9.12) and (9.13).

$$UA = \left(\frac{1}{UA_p} + \frac{1}{UA_c\,f^{0.8}} + R_f\right)^{-1}, \qquad f = \frac{x_c}{100} \tag{9.12}$$

$$k = \varepsilon\,C_{\min}, \qquad Q_{hx} = n\,\dot{m}_p c_p (T - T_r) \tag{9.13}$$

where $UA_p = UA_c = 1800\,\mathrm{kW/K}$, $C_{\min}$ and $C_{\max}$ are the smaller and larger of the pool-side capacity rate $\dot{m}_p c_p = 395.6\,\mathrm{kW/K}$ and the component cooling water capacity rate $\dot{m}_c c_p f$, $C_r = C_{\min}/C_{\max}$, $N = UA/C_{\min}$, and ε is the counter-flow effectiveness $(1 - \mathrm{e}^{-N(1-C_r)})/(1 - C_r\mathrm{e}^{-N(1-C_r)})$, reduced to $N/(1+N)$ when $C_r > 0.999$. The duty coefficient k per train is strongly nonlinear in valve position, giving 63.4 kW/K at 10 % open, 141.9 kW/K at 25 %, 227.1 kW/K at 50 % and 303.7 kW/K wide open at the design fouling of 9.0×10^{-5} K/kW, against 311.7 kW/K clean and 264.4 kW/K at the end-of-season resistance of 6.0×10^{-4} K/kW. The return temperature is a state with its own thermal lag, as shown in Equation (9.14).

$$\tau_{hx}\frac{\mathrm{d}T_r}{\mathrm{d}t} = \left[T - \frac{k\,(T - T_{cw})}{\dot{m}_p c_p}\right] - T_r \tag{9.14}$$

with $\tau_{hx} = 0.08\,\mathrm{h}$. Substituting the bracketed steady value into (9.13) gives $Q_{hx} = nk(T - T_{cw})$, so the exchanger duty is a conductance multiplied by a temperature difference and nk is exactly the negative feedback the pool has. With two trains at 44.20 % that conductance is 423.0 kW/K, and wide open it is 607.5 kW/K.

Pause. Try Exercise 9.1 before reading on.

The makeup valve has equal-percentage trim, as shown in Equation (9.15).

$$q_M = q_{M,\max}\,\frac{R^{x_m/100}-1}{R-1}, \qquad R = 50 \tag{9.15}$$

where $q_{M,\max} = 20\,\mathrm{m^3/h}$ and R is the rangeability. The same trim passes 0.0396 m³/h at 2.37 % of travel, which is the design evaporation rate, and 20 m³/h wide open, so its duty spans a factor of 505 and equal-percentage trim is the reason one valve covers it. The linear gain at the design position is 0.0175 m³/h per percent of travel, which by (9.3) is 1.622×10^{-4} m/h of level per percent.

9.3.5 Instruments

Neither controller sees the plant directly. A Gaussian sequence is passed through a first-order filter and rescaled so that what comes out carries the standard deviation the instrument is specified at, that is added to the true value, and the sum then goes through the transmitter's own filter, as shown in Equation (9.16).

$$\begin{aligned} \tau_{Tm}\frac{\mathrm{d}\hat{T}}{\mathrm{d}t} &= T + n_T(t) - \hat{T} \\ \tau_{Lm}\frac{\mathrm{d}\widehat{(z-z_r)}}{\mathrm{d}t} &= (z-z_r) + n_L(t) - \widehat{(z-z_r)} \end{aligned} \tag{9.16}$$

with $\tau_{Tm} = 0.05\,\mathrm{h}$, $\tau_{Lm} = 0.10\,\mathrm{h}$, noise correlation times of 0.1 and 0.3 h, and default standard deviations of 0.05 °C and 0.003 m. Both filters step forward on the closed-form hold solution $\hat{x}_k = \alpha\hat{x}_{k-1} + (1-\alpha)x_k$ with $\alpha = \exp(-\Delta t/\tau)$, so no choice of sample period can destabilize them. Filtered rates of change of both measurements are reported separately in °C/h and m/h using filters of 0.6 and 1.5 h, and those are the instruments a step test is read from.

9.3.6 Constants, Opening State and Integration

Table 9.2 carries the 31 constants. Twenty-four of them are on a slider somewhere in the interface and the remaining seven are never shown at all, and it is the seven invisible ones that fix how the pool answers a step.

No scenario starts from a guess. Bisection on the net heat of (9.1), monotone decreasing in T, returns the pool temperature that the requested train count and valve position actually hold. The mass is set to put the surface at z_r plus the shielding setpoint, the return temperature to its own steady value from (9.14), the fouling resistance to its commanded fraction of $R_{f,\max}$ and the makeup valve to the position that (9.15) makes equal to evaporation plus any leak. In automatic modes the valve position is instead found by a second bisection for the opening that holds the temperature setpoint. The one exception is on/off train control, which has a limit cycle rather than a steady state, and so starts at the setpoint with the train stopped. The application checks the resulting state by requiring every scaled derivative to be below 1×10^{-7} and the temperature to be within 0.01 °C of 35 °C.

States advance on fourth-order Runge-Kutta with the step held at 0.01 h, which is 36 s, and both controllers scan on that same 36 s boundary. The historian records ten points per simulated hour. The engine, its seeding and its Reset are those of Section 1.6.1. Noise is drawn from a generator seeded at 20 260 906, so a run with the RTD noise switched on repeats to the last digit, and the five speed settings of 0.5, 1, 2, 5 and 20 simulated hours per wall-clock second change only how long a reader waits. A finished run leaves the browser as a comma-separated file of 98 columns, holding the true temperature beside the measured one, the commanded valve beside the actual stem, the four heat-balance terms beside the four mass-balance terms, and both time-to-limit calculations beside each other.

Table 9.2 The 31 Model Constants. The Seven Marked Hidden Never Reach the Screen, and a Step Test Is the Only Route to Them.

Symbol	Meaning	Value	Unit	Visibility
A	free surface area	108	m^2	adjustable
z_r	top of the assemblies	4.40	m	adjustable
z_w	skimmer weir	11.60	m	adjustable
—	shielding minimum	6.40	m	adjustable
—	low-level alarm	6.70	m	adjustable
—	temperature alarm	52	°C	adjustable
—	temperature limit	60	°C	adjustable
T_{sat}	saturation temperature	100	°C	adjustable
n_b	assemblies per batch	64	—	adjustable
n_{aged}	aged batches	6	—	adjustable
τ_0	youngest aged batch	3.0	y	adjustable
—	years between discharges	1.5	y	adjustable
τ_{fresh}	fresh batch cooling time	30	d	adjustable
$\dot{m}_p$	pool-side flow per train	94.6	kg/s	adjustable
$\dot{m}_c$	CCW flow per train, open	158	kg/s	adjustable
UA_p	pool-side conductance	1800	kW/K	adjustable
UA_c	CCW-side conductance	1800	kW/K	adjustable
$q_{M,\max}$	makeup line capacity	20	m^3/h	adjustable
R	makeup rangeability	50	—	adjustable
C_E	evaporation coefficient	35	—	adjustable
h_c	surface convection	3.5	$W/(m^2\,K)$	adjustable
ε	surface emissivity	0.95	—	adjustable
τ_{Tm}	pool RTD filter	0.05	h	adjustable
τ_{Lm}	level transmitter filter	0.10	h	adjustable
$R_{f,\max}$	end-of-season fouling	6.0×10^{-4}	K/kW	hidden
τ_{hx}	exchanger thermal lag	0.08	h	hidden
τ_v	CCW valve lag	0.03	h	hidden
θ	return transport delay	0.06	h	hidden
τ_{foul}	fouling development	10	h	hidden
τ_{cw}	CCW supply lag	0.5	h	hidden
τ_{vm}	makeup valve lag	0.05	h	hidden

9.3.7 Assumptions

Two balances carry all of the detail and everything else in the pool is thin, and each of the simplifications below has a place where it gives out.

One well-mixed bulk temperature. A real pool is stratified, warmer over the freshest racks and cooler at the surface, and local water temperature beside a recently discharged assembly can be several degrees above the bulk. That distribution is exactly what a pool thermal-hydraulic analysis is built to compute [9, 15], and it is the reason real pools carry more than one temperature element. The assumption fails hardest immediately after a discharge and during natural circulation with the pumps stopped.

Decay heat is a curve, not a calculation. Equation (9.7) is a least-squares fit with a worst residual near 5 % between one week and twenty years, and it carries no dependence on enrichment, burnup history or actinide content. A licensing calculation uses a depletion code validated against measured decay heat [10].

Evaporation has no forced-convection term, and is extrapolated one degree past its validation range. Equation (9.9) is the natural-convection branch, and the ventilation multiplier is a blunt stand in for air movement that on some refueling floors doubles the rate. The full model carries a velocity term that this application omits, and it is that term which extends the correlation to moving

air [12]. The surface temperature is capped at 95 °C, one degree above the 94 °C top of the validated water-temperature range, and the 1864 kg/h that the capped correlation returns there sets the 1173 kW of latent absorption that fixes both clocks during a loss of cooling. That single evaluation therefore carries more weight than any other in Section 9.6, it sits outside the range the correlation was tested over, and the correlation's own mean absolute deviation of 13.9 % is the smallest uncertainty that can honestly be attached to it.

Boiling is a clamp. Equation (9.6) pins the pool at 100 °C and sends the surplus out as steam, with no building pressure rise, no venting resistance, no condensation on the structure and no local boiling on hot assemblies before the bulk saturates. All of those matter to an accident analysis and none of them changes the control lesson.

One exchanger per train. There is no tube-side maldistribution, no internal pinch, and no distinction between fouling on the pool side and the component cooling water side, and R_f develops as a first-order lag rather than from a growth model.

Dead time is one lumped constant. The delay applies to the CCW valve command and covers piping transport only, with the valve actuator and the exchanger inventory represented as the separate lags τ_v and τ_{hx}. A real transport delay falls as pool-side flow rises, and a step test cannot separate the three contributions anyway, which is the subject of Section 9.4.3.

Nothing below the top of the racks exists. When the surface reaches $z_r + 0.05\,\mathrm{m}$ the level is clamped and the run is flagged, because cladding temperature, zirconium oxidation, hydrogen generation and the dose rate on the refueling floor are not represented in any form.

No chemistry. Boron concentration, water chemistry, criticality and rack neutronics are absent, and makeup is assumed borated to the pool concentration. Dilution is a real and separate safety concern this application says nothing about.

The numerics are more repeatable than a control system. Hold the step and the scan and every machine produces the identical trajectory, which is not a property any distributed control system has.

9.4 Dynamic Analysis

9.4.1 Two Plants, One Measurement

Linearizing (9.2) and (9.15) about any operating point, as Section 1.4.3 linearizes a tank, gives the level plant of Equation (9.17), because no loss term contributes a derivative with respect to M.

$$\frac{Z(s)}{U_m(s)} = \frac{K_L\,\mathrm{e}^{-\theta_m s}}{s}, \qquad K_L = \frac{1}{\rho A}\frac{\partial \dot m_M}{\partial x_m} \tag{9.17}$$

where $K_L = 1.622 \times 10^{-4}\,\mathrm{m/h}$ of level per percent of makeup valve at the design position and $\theta_m = 0.05\,\mathrm{h}$ is the makeup valve lag, a first-order lag in the model that the integrating form approximates as a dead time. Because the steady-state gain of (9.17) is infinite, a step response has no plateau to measure a gain against, and the pool does not regulate itself at any elevation whatever. Linearizing (9.1) instead gives Equation (9.18).

$$\frac{T(s)}{U_c(s)} = \frac{K\,\mathrm{e}^{-\theta_{\mathrm{eff}} s}}{\tau s + 1}, \qquad \tau = \frac{M c_p}{3600\,k_{self}}, \qquad K = K_I \tau \tag{9.18}$$

where $k_{self} = nk + k_{surf}$ is the total negative feedback in kW/K and K_I is the short-horizon integrating gain in °C/h per percent of valve. The whole of the difficulty is in k_{self}. With two trains at the design point $nk = 423.0\,\mathrm{kW/K}$ and the surface contributes $k_{surf} = 4.20\,\mathrm{kW/K}$, so $\tau = 3.329\,\mathrm{h}$ and the pool is an easy self-regulating loop. With both trains stopped the same expression gives $\tau = 339\,\mathrm{h}$ against decay heat of 2148 kW, and a time constant a hundred times longer than any experiment anybody will run is an integrator, the limit of a long lag that Section 1.4.4 describes. The pool becomes an integrator exactly when it begins to matter. The tank of Chapter 2 turns into an integrator when its outlet valve closes, and this pool's temperature turns into one when its pumps stop.

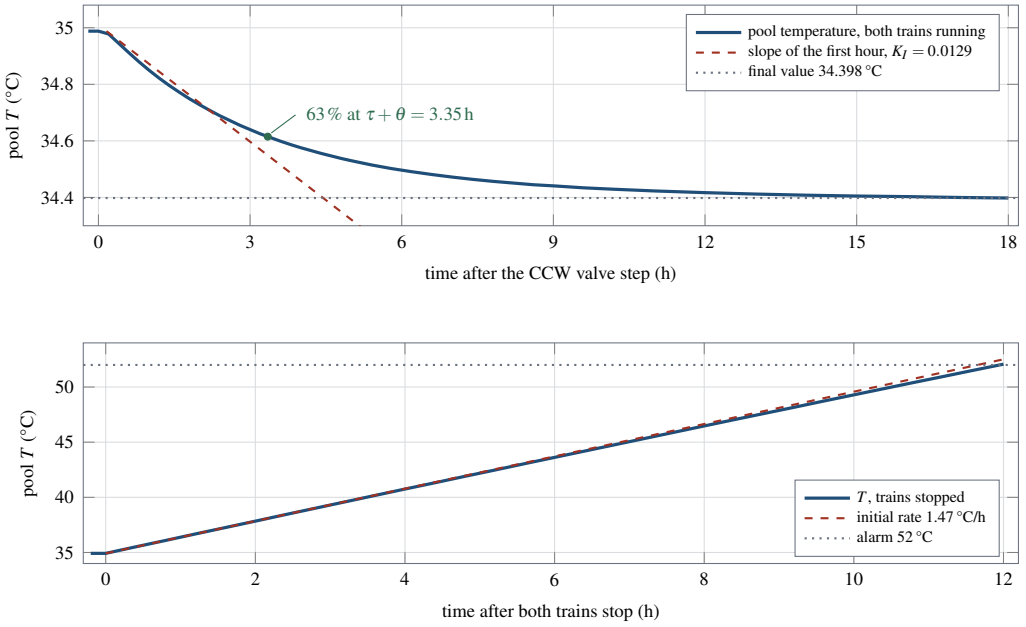


Figure 9.2 The same pool identified as two different plants. Above, a 10.70 % open step on the CCW valve with both trains running settles 0.590 °C lower, giving $K = -0.0551$ °C per percent, $\tau + \theta = 3.35$ h and $\theta = 0.16$ h. Reading only the first hour gives the integrating gain 0.0129 °C/h per percent shown as the dashed tangent, which is a correct measurement of the wrong model. Below, both trains stop 40 h into the same scenario and the pool ramps at 1.47 °C/h initially, averaging 1.43 °C/h over the twelve hours to 52.07 °C because evaporation grows on the way up.

9.4.2 Identification From One Step Test

The identification experiment steps the CCW valve from 44.3 % to 55.0 % with both trains running, waits, and then stops both trains, so that one scenario produces both models. The measured responses are shown in Figure 9.2.

With the trains running the step settles 0.590 °C lower, which gives $K = -0.0551$ °C per percent of valve. The response reaches 63 % of that change 3.35 h after the step and first departs from the pre-step value by 0.005 °C at 0.160 h, so $\theta_{\text{eff}} = 0.16$ h and $\tau = 3.19$ h, an easy loop with $\tau/\theta = 19.9$. Both times are read at the 0.01 h integration step, and the 0.1 h historian places them at its 0.2 and 3.4 h samples. The step test recovers 3.19 h where (9.18) at the design point gives 3.329 h, because the step traverses valve positions at which k is larger and the average gain over the traverse, 0.0551 °C per percent, is smaller than the incremental gain at 44.2 %, which is 0.0674 °C per percent. That difference is the nonlinearity of (9.12) and it is why the identified plant depends on where it is identified.

One variable does not hold still during the experiment. Decay heat falls from 2143 to 2129 kW over the eighteen hours of the test, and a second run in which the valve is returned to 44.3 % on the scan after the scripted step shows the pool cooling 0.032 °C on that account alone. Differencing the two runs gives 0.558 °C of step response and $K = -0.0522$ °C per percent, so the raw reading overstates the gain by 5.7 %. A student who wants the plant rather than the load has to run the unstepped case as well, which is the ordinary price of identifying anything on a plant with a decaying forcing function.

The instructive failure is to stop the test early. Read over the first hour alone, the same response has a slope of –0.1377 °C/h and gives $K_I = 0.0129$ °C/h per percent, a perfectly good integrating gain

and a perfectly good number for a rule that assumes the pool never settles. Every self-regulating plant looks like an integrator over a small fraction of its time constant, and one hour is 31 % of 3.19 h. The response reaches 98 % of its change only 13.65 h after the step, so a test long enough to distinguish the two models costs most of a day. Section 9.5 measures what the distinction is worth.

Stopping both trains produces the lower panel of Figure 9.2. The pool ramps at 1.466 °C/h at the instant of the stop and averages 1.429 °C/h over the twelve hours it takes to reach 52.07 °C, the difference being evaporation growing from 39 to 178 kg/h on the way. The stop occurs 40 h into the scenario, by which point decay heat has fallen to 2117 kW, and the same stop at the untouched design point of 2148 kW gives 1.487 °C/h. With a full core eight days out of the reactor the design-point rate is 4.710 °C/h. There is no gain, no time constant and no final value in any of these cases.

> **TRY IT IN THE APP** **3 · Step Test: One Pool, Two Plants (MANUAL)**
>
> Load the case and press Run. Watch the pool temperature trend after the CCW valve step: it bends over and settles 0.590 °C lower, a plant that regulates itself. Then, 40 h in, both trains stop. The same trend turns into a straight ramp at 1.47 °C/h, an integrator with no final value.

9.4.3 Why the Apparent Dead Time Exceeds the Transport Delay

The plant of (9.18) is not first order plus dead time. It carries a pure transport delay of 0.06 h on the valve command, a valve actuator lag of 0.03 h and the exchanger and train inventory as a further lag of 0.08 h, and a step test cannot separate the three, because a small lag in series with a dominant one displaces the response in the same way a pure delay does. The sum is approximately the quantity an identification recovers, as shown in Equation (9.19).

$$\theta_{\text{eff}} = \theta + \tau_v + \tau_{hx} = 0.06 + 0.03 + 0.08 = 0.17\,\text{h} \tag{9.19}$$

The value read from Figure 9.2 is 0.160 h, one integration step short of (9.19), and it depends on the threshold, 0.13 h for a departure of 0.002 °C and 0.27 h for 0.02 °C, because a chain of lags has no sharp onset to read. Section 1.4.6 makes the same point on a heater: an effective dead time depends on the method that reads it. The consequence for derivative action is the opposite of the one an integrating level loop reaches. On a pure integrator the controlled variable is already a ramp and derivative action amplifies measurement noise without adding anticipation, but here 0.08 h of the 0.17 h is a real first-order lag, so derivative time up to 0.6 h is offered as an adjustable and is worth something on a quiet measurement.

9.5 Controller Design

9.5.1 Four Architectures

The application implements the four structures in the order a plant grows them. Manual mode gives the operator both valve stations and nothing holds anything. On/off train control starts a train when the measured temperature reaches $T_{\text{SP}} + \Delta/2$ and stops it at $T_{\text{SP}} - \Delta/2$, running the CCW valve wide open while the train runs and shut while it does not, which is how most pools are actually operated and which costs nothing and needs no tuning. Proportional-integral control modulates the CCW valve on measured temperature and the makeup valve on measured shielding. The fourth adds an evaporation feedforward to the makeup valve. Both manual valve stations remain live whenever their own loop is not in automatic, which is not the same condition as the global mode, so selecting the manual architecture in automatic leaves both sliders active.

9.5.2 PID Algorithm as Implemented

One incremental, or velocity, algorithm serves both loops, differentiating the measurement rather than the error, as set out in Equations (9.20) and (9.21). Each scan forms the proportional and derivative

increments and accumulates them in a running sum.

$$
\begin{aligned}
e_k &= a\,(x_{\mathrm{SP}} - \hat{x}_k)\,, \qquad P_k = K_c e_k \\
D_k &= -a\,K_c \tau_D \,\frac{\hat{x}^f_k - \hat{x}^f_{k-1}}{\Delta t_c} \\
s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
c_k &= u_{bias} + s_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k,\, 0,\, 100\,\%)
\end{aligned}
\tag{9.20}
$$

$$
I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right] \Delta t_c \tag{9.21}
$$

These are Equations (1.24) to (1.27) of Section 1.6.2 in this chapter's symbols, with $s_k = S_k$, $c_k = v_k$, $u_{bias} = u_b$ and the measurement $\hat{x}$ as the PV. The derivative filter has the time constant $\max(\tau_{Tm}, 0.02\,\mathrm{h}) = 0.05\,\mathrm{h}$, the tracking time T_t defaults to τ_I, and $\Delta t_c = 0.01\,\mathrm{h}$. The temperature loop uses $a = -1$ and is therefore direct acting, so a pool above setpoint opens the CCW valve. The level loop uses $a = +1$ and is reverse acting, so a falling level opens the makeup valve. Anti-windup is back-calculation through the second term of (9.21) [18], and transfer between manual and automatic is bumpless in both directions, both as Section 1.6.2 describes. Setting $\tau_I = \infty$ disables integral action, and the top detent of the level integral-time slider does exactly that, because proportional-only averaging control on a pool this size is an ordinary deliberate choice. With no integral there is nothing to unwind, so the back-calculation is skipped and a proportional-only output is simply clipped. Both outputs are clamped to 0–100 % and neither carries a rate limit, the actuator dynamics being the lags $\tau_v = 0.03\,\mathrm{h}$ and $\tau_{vm} = 0.05\,\mathrm{h}$ of Table 9.2. The CCW valve command is delayed by θ before it reaches the plant rather than delaying the flow, which is equivalent and keeps the delay line out of a right-hand side the integrator evaluates four times per step. Execution is every 0.01 h, against $\theta_{\mathrm{eff}} \approx 0.17\,\mathrm{h}$ of effective dead time.

9.5.3 Which Tuning Rule the Identification Commits To

Three rules are offered and all three read only the numbers the student identified. Equation (9.22) is the SIMC setting for an integrator with dead time and Equation (9.23) the setting for a first-order plant with dead time [5].

$$
K_c = \frac{1}{K_I(\tau_c + \theta)}, \qquad \tau_I = 4(\tau_c + \theta) \tag{9.22}
$$

$$
K_c = \frac{\tau}{K(\tau_c + \theta)}, \qquad \tau_I = \min(\tau,\, 4(\tau_c + \theta)) \tag{9.23}
$$

Which numbers go into these rules matters, and the application's own tuning buttons use the model's incremental gains at the present operating point rather than a student's fit. At the design point those are $K = -0.0674\,^{\circ}\mathrm{C}$ per percent, $\tau = 3.329\,\mathrm{h}$ and $K_I = K/\tau = 0.0202\,^{\circ}\mathrm{C/h}$ per percent, and at $\tau_c = 2\,\mathrm{h}$ Equation (9.23) returns $K_c = 22.77\,\%/^{\circ}\mathrm{C}$ and $\tau_I = 3.33\,\mathrm{h}$, which the application ships rounded to 23 and 3.3 h. On those same gains Equation (9.22) returns the identical gain and $\tau_I = 8.68\,\mathrm{h}$, because $K = K_I\tau$ makes the two expressions for K_c algebraically the same and only the integral time differs, the integrating rule having no plant time constant to cap it with. The whole penalty of the shorter step test is therefore an integral time 2.6 times too long, and Section 9.6 measures what that costs on a setpoint move.

The student's own fit is not the same pair of numbers, and it does not carry that identity. Feeding the measured $K = -0.0551\,^{\circ}\mathrm{C}$ per percent, $\tau = 3.19\,\mathrm{h}$ and $\theta = 0.16\,\mathrm{h}$ into (9.23) gives $K_c = 26.8\,\%/^{\circ}\mathrm{C}$ and $\tau_I = 3.19\,\mathrm{h}$, while the measured $K_I = 0.0129\,^{\circ}\mathrm{C/h}$ per percent in (9.22) gives $K_c = 36.0\,\%/^{\circ}\mathrm{C}$ and $\tau_I = 8.64\,\mathrm{h}$. The gains differ by 34 % because $K_I\tau = 0.0411\,^{\circ}\mathrm{C}$ per percent against a measured K of 0.0551 °C per percent, an inconsistency the step test cannot see and the identity above conceals.

The classical reaction-curve settings for the same integrating model follow from $K_u = \pi/(2K_I\theta)$ and $P_u = 4\theta$, giving Equation (9.24).

$$K_c = \frac{0.94}{K_I\theta}, \qquad \tau_I = 2\theta, \qquad \tau_D = 0.5\theta \tag{9.24}$$

On the model's K_I and $\theta = 0.17\,\mathrm{h}$, Equation (9.24) returns $K_c = 273.2\,\%/^\circ\mathrm{C}$, $\tau_I = 0.34\,\mathrm{h}$ and $\tau_D = 0.085\,\mathrm{h}$, and on the measured pair it returns 456.5 %/°C, 0.32 h and 0.08 h. The two gains are twelve and seventeen times the SIMC gain from the same identification respectively, because the rule places the closed-loop time constant at the dead time instead of at two hours. Neither can be entered, because the temperature gain is bounded at 200 %/°C, which is itself a statement about how aggressive the rule is on this plant.

The level loop is tuned by the averaging rule instead, stated the way an operator states it. For an integrator under proportional-integral control the closed loop is exactly second order with damping ratio $\zeta = \frac{1}{2}\sqrt{K_c K_L \tau_I}$, so critical damping fixes the integral time once the gain is chosen, as shown in Equation (9.25).

$$K_c = \frac{\Delta u}{\Delta h}, \qquad \tau_I = \frac{4}{K_c K_L} \tag{9.25}$$

Choosing $\Delta u = 26\,\%$ of valve travel for $\Delta h = 0.10\,\mathrm{m}$ of level gives $K_c = 260\,\%/\mathrm{m}$ and $\tau_I = 94.8\,\mathrm{h}$, against the shipped 96 h on a slider whose detents are 4 h apart, which is a damping ratio of 1.006 rather than exactly critical. Equation (9.25) is (9.22) with $\tau_c + \theta = \Delta h/(\Delta u K_L)$, so stating how much deviation a disturbance may cause is the same act as choosing a closed-loop time constant. Four days of integral time is a design decision and not a lapse, because the nearest level constraint is 0.60 m away and costs nothing until it is reached, while makeup water is metered, treated and borated. The argument reverses the moment the disturbance is a leak rather than an evaporation change, and the graded challenge is built on exactly that reversal.

9.5.4 The Evaporation Feedforward

The feedforward computes the makeup flow that closes the mass balance now, from quantities a plant measures, and positions the valve there through the inverse of (9.15), as shown in Equation (9.26).

$$\begin{aligned} q_{FF} &= \frac{k_{FF}\,\dot{m}_E(\hat{T}, T_a, \mathrm{RH})}{\rho_M} \\ u_{bias} &= \frac{100}{\ln R}\ln\left[1 + (R-1)\,\frac{q_{FF}}{q_{M,\max}}\right] \end{aligned} \tag{9.26}$$

where k_{FF} is an adjustable trim defaulting to 1 and ρ_M is the density of the makeup water. Above 99.5 °C the correlation stops meaning anything and the branch instead returns the boil-off that the energy balance cannot put anywhere else, $\dot{m}_{boil} = (Q_{decay} - Q_{hx} - Q_{surf} - \dot{m}_E h_{fg})/h_{fg}$. Both branches are arithmetic an operator can do on paper, which is the point. Equation (9.26) enters the level loop as the bias term of (9.20), so the feedback loop is left trimming the error in the calculation and catching the leak the calculation cannot know about, which is the structure simple feedforward tuning rules assume [19]. A change of architecture moves the bias by a whole valve position, so the level integral is re-seeded against the new bias on the change, which keeps the transfer bumpless.

9.5.5 Time to the Limits

The quantity an operator needs during a loss of cooling is not an error but a time. The application reports both the marched answer and the present-rate extrapolation, as shown in Equations (9.27) and (9.28).

$$t_{boil} : \text{Euler march of (9.1) and (9.2) at } 0.25\,\mathrm{h} \tag{9.27}$$

$$\hat{t}_{boil} = \frac{T_{sat} - T}{\mathrm{d}T/\mathrm{d}t}, \qquad \hat{t}_{low} = \frac{z - z_r - 6.40\,\mathrm{m}}{-\mathrm{d}z/\mathrm{d}t} \tag{9.28}$$

Table 9.3 The Twelve Scenarios and What Each One Is For

#	Scenario	What it exercises
1	Normal operation	the scale of everything, no self-regulation on level
2	Hold it by hand for a week	exact flow matching over days
3	Step test, both plants	FOPDT and IPDT from one experiment
4	On/off train control	deadband against starts, capacity margin
5	PI from the identified model	SIMC, reset time, a load it cannot reject
6	Full core offload	actuator saturation, thermal expansion
7	A season of fouling	degradation visible only in valve travel
8	Hot, humid building	one disturbance, two very different balances
9	A liner leak	when averaging tuning becomes the wrong choice
10	Cooling lost during the outage	marched against extrapolated time to limit
11	Evaporation feedforward	computed disturbance against feedback
12	Graded challenge	one train, a fresh core and a leak

The march freezes the present cooling, makeup and leak, advances the two balances to a horizon of 120 h, and reports the first crossing of saturation and of the shielding minimum. It is refreshed every 0.50 h of simulated time, which is far more often than anybody would redo it by hand, and it costs 480 evaluations of the balance. The present-rate form of (9.28) is what an operator computes in their head and what most plant computers display, and it uses the level rate of (9.4) including the expansion term, so it returns no answer at all while the pool is warming faster than it is evaporating.

9.6 Results

Table 9.3 lists the twelve scenarios in the order a laboratory would work through them. Each one arrives with its mode, its architecture, both sets of tuning constants, every disturbance setting and every timed event already fixed, and its Learn card opens by asking for a written prediction before the run. Section 1.6.4 describes the briefing, the completion rules and the scorecard that every scenario shares. The app lists the twelfth as its Challenge, without a number. Figure 9.3 shows the tenth in progress.

9.6.1 The Scale of Normal Operation

The arithmetic of scenario 1 settles the question before a controller appears anywhere in it. The design inventory produces 2147.7 kW, two trains at 44.20 % remove 2115.2 kW of it, the surface takes 31.8 kW and the cold makeup water the remaining 0.7 kW, so the pool sits at 35.00 °C. Evaporation of 39.55 kg/h is 8.84 mm of level a day against 3 mm of transmitter noise, and the makeup valve balances it at 2.367 % of its travel. One cubic meter an hour of imbalance is 9.26 mm/h of level. Scenario 2 puts both valves in manual for a simulated week and asks for the level to be held, which cannot be done, because holding it requires matching a flow of 0.0396 m^3/h exactly and indefinitely on a valve whose smallest adjustment is 0.05 % of travel.

9.6.2 Reset Time and the Load the Loop Cannot Reject

Scenario 5 verifies the tuning of (9.23) and then finds the end of the plant. A setpoint step from 35 to 38 °C at $t = 6$ h with $K_c = 23\,\%/°\mathrm{C}$ and $\tau_I = 3.3$ h settles within ±0.15 °C in 2.72 h and rises 0.051 °C above setpoint 27.6 h after the step, as a late noise-driven excursion (2.66 to 2.95 h and 0.032 to 0.056 °C over five other noise seeds). That excursion is not a dynamic overshoot. Repeating the run with the RTD noise set to zero leaves the response monotone, settling inside ±0.15 °C in 2.76 h and inside ±0.06 °C in 3.34 h with no crossing of setpoint at all, so the 0.05 °C is the true temperature being pushed past setpoint by a controller acting on a measurement that carries 0.05 °C of noise, and not a property of the tuning. The integrating rule's $\tau_I = 8.7$ h at the same gain, run the same way with the RTD noise at zero, takes 13.9 h to settle inside ±0.15 °C, which is what the shorter step test costs.

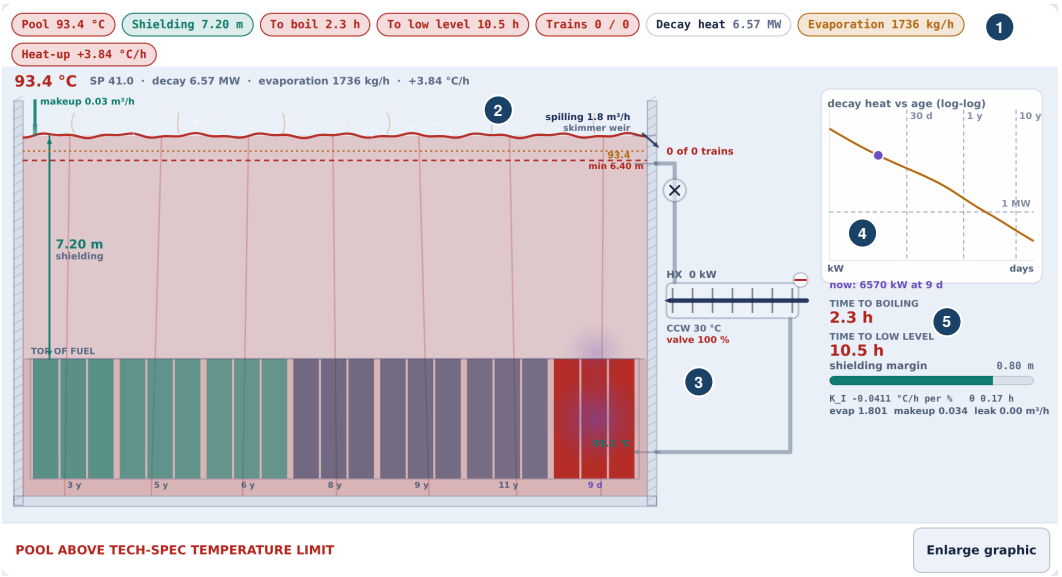


Figure 9.3 Loss of pool cooling in scenario 10 at $t = 13.00$ h, 12.00 h after both trains were lost with a full core in the pool and 1.74 h before the pool reaches saturation at $t = 14.74$ h, cropped to the plant view. (1) The badges read 93.4 °C, 7.20 m of shielding, 2.3 h to boiling, 10.5 h to the low-level limit, zero of zero trains, 6.57 MW of decay heat, 1736 kg/h of vapor and a heat-up rate of +3.84 °C/h. The vapor badge is total vapor, evaporation plus boil-off, and at this instant boil-off is still zero, so all 1736 kg/h of it is surface evaporation. (2) The surface is pinned on the skimmer weir with 1.8 m³/h going over it, and the line under the clocks reads 1.801 m³/h of evaporation and 0.034 m³/h of makeup. (3) The exchanger passes 0 kW, with the CCW valve wide open into a train that is not running. The pool stands more than thirty degrees over its 60 °C technical-specification limit while the shielding badge holds at the weir, which is the whole difficulty of the scenario: a level that has not moved at all. (4) The decay-heat inset places the freshest batch at nine days and 6570 kW on the log-log curve. (5) The two clocks disagree by design. The marched integral, last refreshed at 12.50 h and held until the next controller scan, is 2.25 h to boiling and 10.50 h to the low-level limit, shown rounded, while the straight-line extrapolation of the present rate carried in the CSV row at $t = 13.0$ h reads 1.71 h to boiling and no answer at all for the level, because the level is not falling. The status line reads POOL ABOVE TECH-SPEC TEMPERATURE LIMIT.

The component cooling water supply then rises from 30 to 38 °C and the loop answers by opening the valve until there is no valve left, settling at 41.31 °C with x_c on its stop. Two trains wide open have a duty coefficient of 607.5 kW/K, and 2148 kW across that conductance is a 3.5 °C approach to whatever the cooling water happens to be, so the entire disturbance-rejection range of this loop against T_{CCW} is about 1.5 °C. That is the plant and not the tuning, and no controller improves on it.

9.6.3 On/Off Control and the Cost of a Narrow Band

Scenario 4 runs one train on/off against the aged inventory alone. All five deadbands in Table 9.4 were measured over the same 168 h window, because a start count is meaningless without one. Figure 9.4 draws the starts and the swing side by side.

Narrowing the deadband from 3 to 1.5 °C nearly doubles the starts from 14 to 27 and nearly halves the swing from 3.19 to 1.72 °C, and buys nothing at all, because the temperature limit is more than 23 °C above the top of either band. Over five other noise seeds the narrow band gives 27 or 28 starts and every other count is unchanged, and removing the RTD noise leaves every count unchanged while narrowing each swing by 0.06 to 0.13 °C. Widening it has a hard end. One train wide open holds the

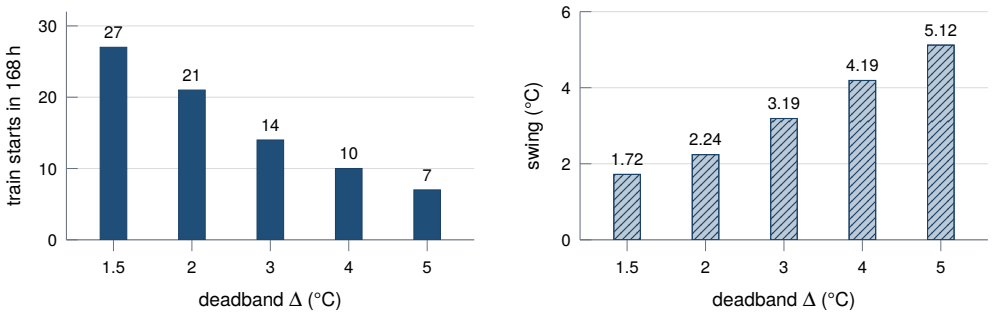


Figure 9.4 On/off train control against the aged inventory alone, from Table 9.4. Left, the train starts in 168 h (solid bars); right, the peak-to-peak swing of the true temperature (hatched bars). Notice that narrowing the band from 3 to 1.5 °C nearly doubles the starts and nearly halves the swing, while the temperature limit stays more than 23 °C above the top of either band, so the narrow band buys nothing.

Table 9.4 On/Off Train Control Against the Aged Inventory Alone, $T_{SP} = 35\,°C$ and the Shipped 0.05 °C RTD Noise, Measured Over 168 h With the Peak-to-Peak Swing of the True Temperature Read After $t = 24\,h$.

Δ (°C)	1.5	2	3	4	5
Starts	27	21	14	10	7
Swing (°C)	1.72	2.24	3.19	4.19	5.12
Mean period (h)	6.2	8.0	12.0	16.8	24.0
Duty (%)	45.9	46.1	46.3	49.5	53.0

aged inventory at 32.18 °C, so a stop temperature below that can never be reached, and any deadband wider than $2(35 - 32.18) = 5.64\,°C$ is a standing order to run. At $\Delta = 6\,°C$ the measured result is one start, 96.0 % duty and a swing of 0.15 °C, which is not on/off control at all. On/off therefore requires a capacity margin that nothing in the tuning displays.

9.6.4 Degradation That Nothing Alarms On

Scenario 7 fouls the exchanger over a season with the temperature loop in service, and the pool temperature does not move, staying within 0.05 °C of setpoint for 130 h (34.98 to 35.05 °C with the shipped RTD noise and 35.00 to 35.02 °C without it). What moves is the actuator. With the RTD noise removed so that the valve reports the plant rather than the instrument, the CCW valve walks from 42.83 % at $t = 4\,h$ to 48.83 % at $t = 130\,h$. The comparison that isolates fouling is the one the scenario asks for, a second run in which the exchanger fouling is set back to zero on the scan after the scripted event at $t = 4\,h$, and there the same valve walks from 42.83 to 40.03 % over the same window because decay heat falls from 2145 to 2055 kW. Fouling is therefore worth 8.81 % of valve travel at the same load and the same temperature. The fouled run on its own shows only 6.00 % of travel, because decay heat is closing the valve by 2.81 % at the same time, so the uncontrolled reading understates the degradation by a third. The duty coefficient measured at the operating point is not the number that reveals it, since it is 404.7 kW/K in both runs to four significant figures, both simply reporting the load the loop is carrying. The number that reveals it is the duty coefficient at a fixed valve position, which for two trains wide open is 532.4 kW/K fouled at $t = 130\,h$ against 623.5 kW/K clean, a loss of 14.6 % that reaches 15.2 % at the end-of-season resistance. When two batches are then discharged into the pool at $t = 130\,h$ both valves saturate and the fouled plant settles at 37.74 °C at $t = 180\,h$ where the clean plant gives 36.59 °C, a difference of 1.15 °C. Fouling is a margin problem rather than a temperature problem, and margin is only ever measured in actuator travel.

9.6.5 One Disturbance, Two Very Different Balances

Scenario 8 holds the pool at 42 °C with a full core twelve days out and moves the refueling floor. Decay heat is falling throughout, so every number below is read with the instrument noise removed and differenced against a companion run in which the building is set back to 30 °C and 45 % relative humidity on the scan after each scripted step and nothing else is changed. Taking the building to 40 °C and 85 % drops evaporation from 80.6 to 16.2 kg/h, and the latent duty with it from 53.8 to 10.8 kW. The CCW valve absorbs that in 0.98 % of travel, and the pool departs from the companion run by at most 0.026 °C, two hours after the step, and is back on it within a day rather than anywhere near its alarm. Taking the building to 22 °C and 20 % at $t = 72$ h raises evaporation to 108.9 kg/h and gives back 0.40 % of valve against the same reference, for at most 0.036 °C of pool temperature in the other direction. The whole factor of 6.7 in evaporation is therefore 1.38 % of valve travel, the pool temperature stays within 0.048 °C of its setpoint, and the level deviation never exceeds 10.3 mm. On the mass balance the second step is 0.099 m^3/h of unbalanced flow, which is 0.92 mm/h of level and clears 3 mm of transmitter noise in 3.3 h, so this disturbance is one the level can eventually see. The raw valve travel from the first step at $t = 8$ h to $t = 160$ h is 54.3 to 41.8 %, of which 12.2 of the 12.6 points is decay heat falling from 5761 to 4997 kW rather than the weather.

9.6.6 Saturation, and a Makeup Valve That Shuts

Scenario 6 discharges a full core eight days out of the reactor into a pool holding only the aged inventory, taking decay heat from 683 to 6680 kW, and starts the second train at the same instant. The CCW valve command reaches its 100 % stop at the scan at $t = 5.18$ h, 1.18 h after the discharge, and the temperature goes on rising for another 9.85 h, peaking at 40.68 °C at $t = 15.03$ h and falling only slowly afterward, to 40.49 °C a day after the discharge and 40.09 °C two days after. With the actuator pinned the pool is an integrator again, and an integrator does not care that its input has stopped moving. Watching the controlled variable reports the trouble about ten hours after watching the actuator would have.

The makeup valve does the opposite of what the mass balance asks. Evaporation nearly doubles from 39.5 to 71.5 kg/h, and with the instrument noise removed the level loop closes the valve to 0 % at the scan at $t = 4.98$ h and holds it there for 21.9 h, because 5.68 °C of warming on 1231 m^3 raises the surface by 18.1 mm through the second term of (9.4), and at the prevailing 0.66 mm/h of evaporation the surplus takes 29.4 h to leave, from its peak at $t = 9.81$ h to the return to setpoint at $t = 39.24$ h. With the shipped noise the valve chatters off its stop, but the level follows the same course, peaking 18.2 mm high at $t = 9.82$ h and back on setpoint at $t = 39.47$ h. A level transmitter cannot tell thermal expansion from inventory, and that is a property of every level measurement on a vessel whose temperature moves.

🕮 **Pause.** Try Exercise 9.3 before reading on.

9.6.7 A Leak Against an Averaging Loop

Scenario 9 opens a leak-chase drain at 0.05 m^3/h with the level loop deliberately loose at $K_c = 30\,\%/\mathrm{m}$ and proportional only. With the instrument noise removed the pool temperature does not move at all, staying between 34.996 and 35.000 °C over the first four days, because the whole makeup flow, 0.040 to 0.057 m^3/h of 20 °C water, is at most 1.0 kW against 2148 kW and the level is not in the energy balance in any case. On the level the deviation is 1.84 mm after four hours, 10.5 mm after a day and 33.0 mm after the eighty-eight hours before the leak grows, so it takes most of a day to clear 3 mm of transmitter noise. The mass balance is not so slow. At $t = 96$ h the metered makeup is 4.63 m^3 against an evaporation total of 3.82 m^3, and that difference plus the 3.57 m^3 the inventory has lost is within 0.02 m^3 of the 4.40 m^3 the drain has taken. When the drain then opens to 0.6 m^3/h, the same loop gives up 690 mm by $t = 400$ h and parks the pool at 6.310 m of shielding, below the 6.40 m minimum, with the makeup valve only 23.1 % open and still drifting downward at 0.36 mm/h. It is not diverging. Carried to $t = 900$ h the same run settles at 6.278 m, so the eventual offset is 722 mm

Table 9.5 Level Control Against a Swinging Evaporation Rate, Scenario 11. The Two Columns Differ Only in the Architecture. Worst Deviation Is Measured Over $6 \le t \le 240$ h With the Instrument Noise Removed.

	PI only	PI + evaporation FF
Worst level deviation (mm)	170.0	0.8
Time of the worst deviation (h)	132.0	10.7
Makeup valve travel (%)	5.36–10.46	4.52–14.19
Feedforward contribution (%)	0	4.52–14.18
Makeup used to $t = 240$ h (m^3)	38.5	46.8

and the 690 mm at 400 h is 96 % of it. An averaging design that is correct against an evaporation change is the wrong design against a leak, and no alarm distinguishes them.

9.6.8 Feedforward Against Feedback

Scenario 11 holds a hot pool at 50 °C with one train and swings the building twice. The two runs in Table 9.5 differ in the architecture and in nothing else, both starting from the same state with the same level tuning.

Evaporation goes from 94.8 to 301.8 and back to 78.9 kg/h, and the feedforward moves the makeup valve from 5.35 to 14.17 and back to 4.52 % without the level having moved at all, holding the worst deviation to 0.8 mm against 170.0 mm on feedback alone. The reason feedback loses is not the tuning, and the 170.0 mm figure needs its window stated, because the deviation is still growing when the disturbance reverses at $t = 132$ h. Holding the disturbance instead, by setting the building back to its hot, dry, ventilated state on the scan after the reversal, the proportional-only loop settles 294.2 mm low by $t = 900$ h, so the 170.0 mm observed at $t = 132$ h is 58 % of the eventual offset and the loop is neither diverging nor finished. What controls the first day is Equation (9.29).

$$\Delta z = \frac{\Delta q\, t}{A} \tag{9.29}$$

which contains no controller gain. An extra 0.21 m^3/h of evaporation over 108 m^2 is 1.9 mm/h whatever K_c is, and raising the gain shortens the tail while doing nothing to the beginning, because reading a change in evaporation off the level means waiting for its integral to clear the transmitter noise. The evaporation rate, by contrast, is computable now from a pool resistance temperature detector, a building thermometer, a hygrometer and Equation (9.9). The general criterion is that feedforward wins whenever the disturbance can be computed faster than its effect on the controlled variable can be measured, and on this pool the comparison is not close.

> ✦ DID YOU KNOW? **Better Level Control Costs More Water**
>
> From the simulator: in scenario 11 the evaporation feedforward holds the worst level deviation to 0.8 mm against 170.0 mm for PI alone, yet it draws more makeup water, 46.8 m^3 against 38.5 m^3 by $t = 240$ h (Table 9.5).

9.6.9 Time to the Limits, Marched and Extrapolated

Scenario 10 is the memorable one. A full core eight days out of the reactor produces 6731 kW, both trains are lost at $t = 1$ h, and both valve stations are left in manual with the CCW valve wide open and the makeup valve at the 2.05 % the scenario ships, which passes 0.034 m^3/h against the 0.0396 m^3/h that balances the design point. The pool boils at $t = 14.74$ h and reaches the shielding minimum at $t = 22.90$ h. Figure 9.5 compares what the two calculations of (9.27) and (9.28) said along the way.

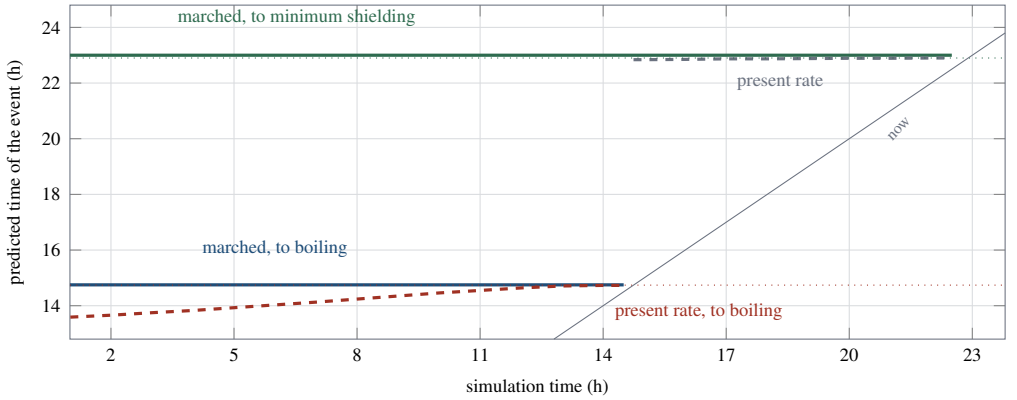


Figure 9.5 Predicted time of each limit against elapsed time during the loss of cooling of scenario 10, sampled at each half-hour refresh of the march, and hourly for the present-rate form. Solid lines are the marched predictions and dashed lines the present-rate ones, the lower pair for boiling and the upper pair for the shielding minimum. Dotted horizontals mark the outcomes, boiling at 14.74 h and the shielding minimum at 22.90 h, and the thin diagonal is the present instant. The marched prediction is available at $t = 1.00$ h, the controller scan on which the trains are lost, and is never more than 0.10 h from the outcome. The present-rate prediction of boiling starts 1.15 h early and reaches the truth only as the event arrives, and the present-rate prediction of the shielding minimum does not exist at all before $t = 14.74$ h because the level is rising until $t = 9.78$ h and pinned on the overflow weir from there until the pool boils.

> **TRY IT IN THE APP** **10 · Cooling Lost During the Outage**
>
> Load the case and press Run. Both trains are lost at 1 h. Watch the shielding badge: the level does not fall. Thermal expansion carries the surface onto the skimmer weir at 9.78 h and holds it there while the pool heats, and the pool boils at 14.74 h, from when the level is the only measurement that still says anything.

The marched calculation is available on the same controller scan the trains are lost on and it is right. At $t = 1.00$ h it reports 13.75 h to boiling and 22.00 h to the shielding minimum, which place the two events at 14.75 and 23.00 h against outcomes of 14.74 and 22.90 h, errors of 0.01 and 0.10 h on forecasts made thirteen and twenty-two hours ahead, and over the whole approach no refresh of either march is more than 0.10 h out. The present-rate extrapolation is wrong in both directions at once. On temperature it reports 12.11 h half an hour after the trip, placing boiling at 13.62 h and running 1.12 h early, because it takes no account of evaporation growing from 90 to 1864 kg/h and absorbing 1173 kW on the way up. On level it returns no answer at all, because the pool is warming at 4.68 °C/h and expanding upward at 20.4 mm/h through (9.4) while evaporation takes 0.85 mm/h away, a net rise of 19.9 mm/h, so the arithmetically correct present-rate answer is that the level is rising and will never reach the minimum.

Two consequences of (9.4) and (9.5) appear on the way. Thermal expansion carries the surface onto the skimmer weir at $t = 9.78$ h and holds it there until the pool boils at $t = 14.74$ h, spilling 10.15 m^3 of inventory at up to 2.48 m^3/h while the pool is losing water, so heating alone costs shielding. Then the pool saturates, and the vapor leaving the surface becomes 1864 kg/h of evaporation plus 8399 kg/h of boil-off, a total of 10.3 t/h, which is 99 mm/h of level against 0.40 mm/h before the trip. By $t = 30$ h the pool has boiled off 131.6 m^3. The temperature trend is flat from 14.74 h onward and the level is the only measurement that still says anything, which is the reordering of priorities that Section 9.2 describes.

9.6.10 Graded Challenge

The challenge puts a full core twelve days cooled, 5816 kW, into a pool with a fouling fraction of 0.6 and a makeup line rated at only 8 m³/h. One of the two trains trips at $t = 8$ h and returns forty hours later, and two hours after the trip a leak-chase drain opens at 2.5 m³/h. The target is the pool below 52 °C and the shielding above 6.60 m throughout, with the pool back on its 44 °C setpoint within ±0.6 °C by $t = 132$ h. Starting from the cautious default tuning the run fails all three tests, peaking at 52.03 °C at $t = 24.2$ h, falling to 5.165 m of shielding, spending 94.1 h below the minimum and ending 1.77 °C off setpoint. The suggested configuration, PI with the evaporation feedforward, $K_c = 23\,\%/°\mathrm{C}$ and $\tau_I = 3.3$ h on temperature and $K_c = 400\,\%/\mathrm{m}$ and $\tau_I = 32$ h on level, meets every part of it, with a peak of 51.18 °C, a minimum shielding of 6.886 m, no time below the minimum and a final error of 0.02 °C, at a cost of 317 m³ of makeup water against 121 m³. That leaves 0.286 m of shielding and 0.82 °C of temperature to spare. The verdicts are the same when the configuration is carried through Reset, and five other noise seeds give the same verdicts for both runs. The two bounds that must hold throughout keep their extremes from every physics step, including the starting state, so an excursion between logged rows still fails. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

The two excursions look identical on a trend and have opposite causes. A slow ramp away from setpoint that eventually stops somewhere unhelpful describes both, and the actuator is what separates them. A CCW valve pinned at 100 %, as it is under the suggested tuning from the scan at $t = 8.91$ h until the train returns, is a plant that has run out of capacity, while a makeup valve sitting near 12 % at $t = 12$ h with the level already falling, as it does under the shipped tuning, is a controller that was never asked hard enough. The suggested configuration therefore does three different things. It tunes the temperature loop from the identified first-order model by (9.23), so the recovery after the train returns takes hours instead of days. It raises the level gain from 30 to 400 %/m and gives it a real integral time of 32 h, because this week the level loop is holding a shielding limit against a leak rather than doing averaging control against evaporation, and the right closed-loop time constant of (9.25) has changed. And it puts the evaporation feedforward in service, which does nothing about the leak but takes the evaporation term off the feedback loop, leaving all of the loop's authority for the disturbance nobody predicted. The makeup valve ends at 73.7 % of an eight-cubic-meter line, which is what balancing a 2.5 m³/h leak on equal-percentage trim requires.

9.7 Technology in Practice

The model of Section 9.3 reads one temperature and one level and moves a cooling water valve and a makeup valve. Figure 9.6 draws both loops with the devices of a real pool in place of the schematic's symbols.

What Measures It

Pools in light water reactors monitor level, temperature, liner leakage, pump discharge pressure and system flow. Most of it is shown on a local panel, only the important parameters reach the control room, and in general the instruments have no redundancy [13]. Temperature is read by thermocouples or resistance temperature detectors, mounted in contact with the water or in thermowells [20]. One pool design places them in thermal wells at the inlets and outlets of its cooling system and alarms at typically ±3 °C on the local panel, and a typical maximum pool temperature in normal operation is 60 °C [13], the technical-specification limit the chapter uses. The chapter's detector carries 0.05 °C of noise behind a 0.05 h filter, and its level transmitter 3 mm behind 0.10 h.

TECHNOLOGY Sensor: Pool Level on a Narrow Span

The commonest level measurement in nuclear plants reads differential pressure, the hydrostatic head between a tap at the bottom and one at the top [20]. Some pools use an air bubbler, a purge gas fed slowly down a dip tube whose back pressure equals the water head at its tip [13, 21]. Too much purge flow adds friction and reads high [21]. A pool level sensor typically spans only about 2 m, around the high and low alarms and the technical-specification minimum, so after a large loss of water it is off scale and useless. In the pools of the 2011 accident the level was read in the skimmer surge tank, which follows the pool only while the pool overflows into it [13].

What Moves It

A pool cooling system pumps the water through heat exchangers to an intermediate cooling system, which rejects the heat to the plant's ultimate heat sink. It takes suction through a skimmer or strainer at an elevation where a falling level makes the pumps lose suction, and its pipes enter the pool high on the wall, so that a broken pipe cannot drain the pool. Pipes that reach deep into the pool carry anti-siphon devices that let air in once the falling water uncovers them [13]. The chapter's two trains, each a pump and a plate exchanger to component cooling water, and its skimmer weir stand for this arrangement, with the valve x_c on the cooling water side as the final element of the temperature loop.

TECHNOLOGY Final Element: Makeup and Cooling Water Valves

Makeup comes from a high-purity source, usually the refueling water storage tank in a pressurized water reactor, through a connection to the suction of the pool cooling pumps, and local valve operations are needed to start it. Rates differ widely among plants, and service water or fire water can serve as alternates. The cooling water valves to the pool exchangers may be opened by hand [13]. Equal-percentage trim is built to offset the droop of a pressure drop that changes with flow, but its textbook formula never reaches zero flow and cannot be trusted at small openings [21], which is where this chapter's makeup valve balances evaporation, at 2.37 % of travel.

Where the Controller Runs

The safety systems of a pool are much simpler than the emergency core cooling systems, and manual operator actions are needed to cope with pool accidents [13]. The loops of a plant run on controllers with cycle times from milliseconds to minutes [22]. The application executes its two velocity-form PI blocks every 0.01 h, against an effective dead time near 0.17 h on temperature, and the evaporation feedforward enters the level block as its bias. The time-to-limit march is a calculation rather than a control block: it advances the two balances every 0.50 h of simulated time, beside the present-rate form that most plant computers display.

IN PRACTICE Valves, Not Pumps, Drain Pools

A 1997 regulatory study in the United States found that pools had lost more than one foot of water 10 times in 12 years, about once every 100 reactor-years, and had lost cooling with a rise of more than 11 °C 4 times. The main cause was human error, most often a mispositioned valve. Loss of cooling usually comes from an inoperable pump, and because these events are slow, operators have always restored cooling long before the pool reached its limit. The lessons drawn are reliable level, temperature and area radiation instruments, procedures and training, and checks of gates, seals and anti-siphon devices [13].

What Is Modernizing It

After the 2011 accident the regulator in the United States ordered reliable means of remotely monitoring wide-range pool level. European stress tests asked for better level, temperature and area radiation instruments, together with fixed and portable means to feed water to the pool [13]. Nuclear instrumentation and control is moving from analog to fully digital systems, apart from the transducer

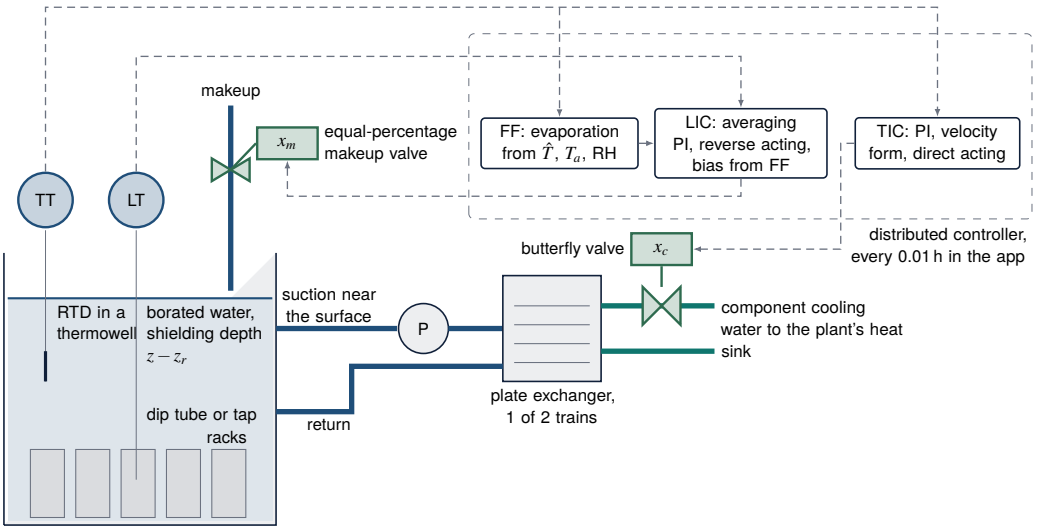


Figure 9.6 The chapter's two loops with the devices of a real pool in place of the schematic's symbols. A resistance temperature detector in a thermowell (TT) feeds the temperature controller (TIC), which moves the butterfly valve x_c on the component cooling water side of the running trains. A level transmitter on a dip tube or pressure taps (LT) feeds the level controller (LIC), whose bias is the evaporation feedforward (FF), and moves the equal-percentage makeup valve x_m. Notice that the only signal between the two loops is the measured temperature that the feedforward reads.

element, through individual upgrades constrained by licensing. Newer smart transmitters send digital serial bus signals, and fiber optics are being applied to temperature measurement [20].

9.8 Limitations

The model is not a spent-fuel pool safety analysis and cannot be made into one. With one well-mixed bulk temperature there is no stratification, no local temperature beside a fresh assembly and no natural-circulation pattern through the racks, all of which a pool thermal-hydraulic calculation exists to produce [9]. Decay heat is a fitted curve with a worst residual near 5 % rather than a depletion calculation [10]. Nothing below the top of the racks is represented in any form, so cladding temperature, oxidation, hydrogen generation and refueling floor dose rate are all absent, and the application clamps the level and raises a flag rather than integrating into nonsense. The boiling branch is a clamp with no building pressure, no venting resistance and no condensation. Evaporation carries no forced-convection term, so the ventilation multiplier stands in for air movement the correlation would otherwise compute [12], and the surface temperature it is evaluated at is capped at 95 °C, one degree past the top of the validated range, so the 1864 kg/h that fixes both clocks during a loss of cooling carries the correlation's 13.9 % mean absolute deviation at best. Chemistry is absent entirely, which puts boron dilution, criticality and rack neutronics out of reach. Disturbances step and events fire on a schedule, where a real refueling floor drifts, and a fixed step with a fixed scan repeats itself more faithfully than any plant does. Three of the scripted events are stylized on purpose. Both cooling trains disappearing at one instant is not how a plant loses cooling, but it isolates the change of plant character. A leak-chase drain that appears instantly at 2.5 m^3/h is not something a stainless liner does, but it separates a demand that exceeds the actuator from a controller that was tuned too loosely. The refueling floor moving from 40 °C and 85 % humidity to 22 °C and 20 % in one step is a caricature of weather, but it is the one building change large enough to leave a readable signature on a trend. The features that survived the simplification are the ones a controller has to be designed around, which are the pure integrator on level, the negative feedback path that can be removed, a transport delay in

series with two lags, the equal-percentage valve characteristic, the nonlinear exchanger conductance and the coupling of the two balances through evaporation. Where each individual assumption breaks is set out in Section 9.3.7.

> ⚠ **OPEN CHALLENGE** **What Happens Below the Top of the Racks**
>
> The chapter's model stops at the top of the racks, and the research does not stop there. A 2015 international status report finds the main phenomena of a pool cooling accident largely understood but names open questions: the flow and neutron behavior of partly uncovered assemblies in rack structures, hydrogen production, and the oxidation and bursting of zirconium alloy cladding in mixed air and steam. Most data come from tests made for reactor cores, and the codes were written for reactors. Progress would be dedicated pool experiments, benchmarks of the codes against them, and a ranking of the phenomena that are both important and uncertain [13].

9.9 Discussion

Pool inventory carries a direct cost and a direct constraint. Makeup water is treated and borated, and the challenge run in Section 9.6 consumes 317 m^3 of it against 121 m^3 for the run that fails, so the tuning that holds the shielding limit is also the tuning that buys the most water. That is the trade averaging level control exists to make, and it is why Table 9.5 reports valve travel and makeup volume beside the peak deviation rather than underneath it. What changes the trade is the identity of the disturbance. Against evaporation the deviation is free until the constraint is reached and a loose loop is correct, while against a leak the deviation ends at 6.310 m of shielding with the valve 23.1 % open, and the same loop is wrong. Distinguishing the two requires a mass balance rather than a level, and the arithmetic that makes the feedforward of (9.26) work is the same arithmetic that detects the leak, namely metered makeup plus reported inventory change against computed evaporation.

The monitoring calculations matter more here than the tuning does. Fouling consumed 15.2 % of the wide-open duty coefficient over a season while holding the pool at 35.00 °C, and the only evidence available at the time was 8.81 % of valve travel at the same load. No alarm exists for an actuator further open than it used to be, and on a plant whose load decays continuously nobody separates the two contributions by eye. The two calculations the exported record supports are the duty coefficient at a known valve position, which reveals exchanger condition, and metered makeup against computed evaporation, which reveals a leak within hours rather than the days a level transmitter needs. The cabin of Volume 6, Chapter 10 uses the same idea on a sealed inventory, an hourly nitrogen balance with a cumulative-sum test as its leak detector.

Methods the application does not itself implement are within reach of the model as stated in this chapter. The seven-state balance with two hard actuator limits and an explicit shielding constraint suits model predictive control, which would enforce that constraint directly and coordinate both valves as one two-input problem. Run as a moving horizon estimator, the same equations would back out the unmeasured leak and the ventilation multiplier from nothing but the two measurements, which promotes the mass-balance monitor to a soft sensor. The march of (9.27) is already an open-loop prediction over a 120 h horizon, so replacing it with an optimization that chooses the makeup profile and the train alignment to maximize the time to the shielding minimum is a small step from what is implemented.

9.10 Conclusion

This chapter studied a spent-fuel pool, a plant whose apparent dynamics depend on which equipment is running. The simulator integrates seven states, the pool temperature and mass, the exchanger return temperature, the two valve positions, the fouling resistance and the component cooling water supply temperature, with decay heat summed over a mixed-age inventory of 448 assemblies producing 2147.7 kW, evaporation from a natural-convection mass-transfer correlation and an effectiveness-NTU exchanger per train.

One step test identifies two plants. With both trains running the response settles and gives $K = -0.0551\,^{\circ}\mathrm{C}$ per percent, $\tau = 3.19\,\mathrm{h}$ and $\theta = 0.16\,\mathrm{h}$, while its first hour alone gives the integrating gain 0.0129 °C/h per percent and an integral time 2.6 times too long. With both trains stopped the same pool ramps at 1.49 °C/h on the design inventory and 4.71 °C/h with a full core in it, because the 423.0 kW/K the exchanger supplied has been replaced by 4.20 kW/K of surface loss.

Feeding the computed evaporation rate forward holds the level within 0.8 mm where the same feedback loop gives up 170.0 mm over 132 h and would eventually give up 294 mm, and the first-day deviation contains no controller gain at all. Marching the two balances forward predicts boiling and the shielding minimum within 0.10 h from thirteen hours out, where extrapolating the present rate is 1.12 h early on temperature and returns no answer on level. On/off control at a 3 °C deadband costs 14 pump starts a week, and widening the band past 5.64 °C stops the train from ever stopping.

A fuller model would add stratification as a second temperature node, a forced-convection evaporation term, boron concentration as a third inventory, and a model predictive controller that maximizes the time to the shielding minimum. As it stands, the pool is a place to identify, tune and predict an integrating process that is slow enough to be ignored.

KEY POINTS

- The pool level only adds up what flows in and out, so a makeup valve slightly off balance makes it ramp without end.
- With two trains running the pool temperature is first order with a time constant of 3.33 h. With both stopped it is 339 h, an integrator.
- Marching the two balances predicts boiling and the shielding minimum within 0.10 h from thirteen hours out, while the present rate is 1.12 h early on temperature and gives no answer on level.
- Feedforward from computed evaporation holds the level within 0.8 mm where feedback alone gives up 170.0 mm, because the first-day deviation contains no controller gain.

Where else this applies. The tank of Chapter 2 turns into an integrator when its outlet valve closes, Chapter 6 has a level that rises while the vessel loses water, and Volume 6, Chapter 10 finds a leak by mass balance.

References

[1] B. W. Bequette, “Process control practice and education: Past, present and future,” *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[2] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, “Benchmark temperature microcontroller for process dynamics and control,” *Computers & Chemical Engineering*, vol. 135, p. 106736, 2020.

[3] W. L. Luyben, “Liquid level control: Simplicity and complexity,” *Journal of Process Control*, vol. 86, pp. 57–64, 2020.

[4] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, “Use of model predictive control and weather forecasts for energy efficient building climate control,” *Energy and Buildings*, vol. 45, pp. 15–27, 2012.

[5] S. Skogestad, “Simple analytic rules for model reduction and PID controller tuning,” *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[6] T.-K. Wang, S.-L. Chang, and S.-P. Teng, “Criticality safety evaluation for high-density spent-fuel storage racks,” *Nuclear Technology*, vol. 83, no. 1, pp. 5–15, 1988.

[7] *Safety and Security of Commercial Spent Nuclear Fuel Storage: Public Report*. Washington, DC: National Academies Press, 2006.

[8] R. Alvarez, J. Beyea, K. Janberg, J. Kang, E. Lyman, A. Macfarlane, G. Thompson, and F. N. von Hippel, “Reducing the hazards from stored spent power-reactor fuel in the United States,” *Science & Global Security*, vol. 11, no. 1, pp. 1–51, 2003.

[9] R. R. Gay, “Spent nuclear fuel storage pool thermal hydraulic analysis,” *Progress in Nuclear Energy*, vol. 14, no. 3, pp. 199–225, 1984.

[10] T. Yamamoto and D. Iwahashi, "Validation of decay heat calculation results of ORIGEN2.2 and CASMO5 for light water reactor fuel," *Journal of Nuclear Science and Technology*, vol. 53, no. 12, pp. 2108–2118, 2016.

[11] M. M. Shah, "Improved method for calculating evaporation from indoor water pools," *Energy and Buildings*, vol. 49, pp. 306–309, 2012.

[12] ——, "Improved model for calculation of evaporation from water pools," *Science and Technology for the Built Environment*, vol. 24, no. 10, pp. 1064–1074, 2018.

[13] OECD Nuclear Energy Agency, Committee on the Safety of Nuclear Installations, "Status report on spent fuel pools under loss-of-cooling and loss-of-coolant accident conditions: Final report," OECD Nuclear Energy Agency, Paris, Report NEA/CSNI/R(2015)2, 2015. [Online]. Available: https://www.oecd-nea.org/upload/docs/application/pdf/2021-02/csni-r2015-2.pdf

[14] *Lessons Learned from the Fukushima Nuclear Accident for Improving Safety of U.S. Nuclear Plants*. Washington, DC: National Academies Press, 2014.

[15] H. Mochizuki, "Evaluation of spent fuel pool temperature and water level during SBO," *Annals of Nuclear Energy*, vol. 109, pp. 548–556, 2017.

[16] M. M. Shah, "Calculation of evaporation from Fukushima nuclear power plant spent fuel pools," *Journal of Nuclear Engineering and Radiation Science*, vol. 5, no. 4, p. 041602, 2019.

[17] A. L. Buck, "New equations for computing vapor pressure and enhancement factor," *Journal of Applied Meteorology*, vol. 20, no. 12, pp. 1527–1532, 1981.

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

[19] J. L. Guzmán and T. Hägglund, "Simple tuning rules for feedforward compensators," *Journal of Process Control*, vol. 21, no. 1, pp. 92–102, 2011.

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

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

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

9.11 Exercises

9.1 ⌕ Concept check · 2 min

Equation (9.1) is the energy balance on the pool water. Which of its terms supplies almost all of the pool's negative feedback while a cooling train runs, and disappears when both trains stop?

(a) Q_{decay}, the decay heat of the 448 assemblies in the racks
(b) $\dot{m}_E h_{fg}$, the latent heat that evaporation carries away
(c) Q_{hx}, the duty of the running trains through their exchangers
(d) $\dot{m}_M c_p (T_M - T)$, the heat taken up by the cold makeup water

9.2 ⌕ Concept check · 1 min

Table 9.1 classifies what the two loops interact with. Which of the following can a controller of the pool move?

(a) The shielding depth $z - z_r$ above the tops of the assemblies
(b) The component cooling water supply temperature T_{CCW}
(c) The liner leak $\dot{m}_L$ that drains from the pool floor
(d) The position x_m of the equal-percentage makeup valve

9.3 Concept check · 2 min

In scenario 6 a full core comes into the pool and evaporation nearly doubles, from 39.5 to 71.5 kg/h. With the instrument noise removed, the level loop closes the makeup valve to 0 % and holds it there for 21.9 h. Why?

(a) The warming water expands and lifts the surface above its setpoint, so the loop sees a high level

(b) The evaporation feedforward subtracts the extra evaporation from the bias of the makeup valve

(c) The level loop is direct acting, so a larger evaporation rate drives its output toward zero

(d) Back-calculation copies the saturated CCW valve into the integral of the level loop

9.4 Theory and calculation · 6 min

At the design point the pool holds $Mc_p = 5.120\,\mathrm{GJ/K}$. Two trains at 44.20 % give $nk = 423.0\,\mathrm{kW/K}$ and the surface gives $k_{surf} = 4.20\,\mathrm{kW/K}$.

(a) Compute the time constant of Equation (9.18) with both trains running and with both stopped. *Checkpoint:* 3.33 h and 339 h.

(b) Both trains stop at the design point. Decay heat is 2147.7 kW, the surface removes 31.8 kW and the cold makeup 0.7 kW. Compute the initial heat-up rate, and the time a straight line at that rate takes from 35 to the 60 °C limit. *Checkpoint:* 1.49 °C/h and 16.8 h.

(c) With a full core eight days out of the reactor the decay heat is 6731 kW. Repeat the initial rate with the same surface and makeup terms, and say why the real pool reaches boiling later than a straight line at that rate predicts. *Checkpoint:* 4.71 °C/h.

9.5 Theory and calculation · 6 min

The step test of scenario 3 steps the CCW valve 10.70 % open with both trains running. Read to the end it gives $K = -0.0551\,^{\circ}\mathrm{C}$ per percent, $\tau = 3.19\,\mathrm{h}$ and $\theta = 0.16\,\mathrm{h}$. Read over its first hour only, the pool falls at 0.1377 °C/h. Use magnitudes, since the controller action carries the sign, and $\tau_c = 2\,\mathrm{h}$.

(a) Compute the integrating gain K_I from the first hour, and the gain and integral time of the first-order rule, Equation (9.23). *Checkpoint:* $K_I = 0.0129\,^{\circ}\mathrm{C/h}$ per percent. $K_c = 26.8\,\%/^{\circ}\mathrm{C}$, $\tau_I = 3.19\,\mathrm{h}$.

(b) Compute the gain and integral time of the integrating rule, Equation (9.22), and the ratio of the two gains. Explain the ratio with $K_I\tau$. *Checkpoint:* $K_c = 36.0\,\%/^{\circ}\mathrm{C}$, $\tau_I = 8.64\,\mathrm{h}$, ratio 1.34.

(c) Compute the settings of Equation (9.24) on the same K_I and θ. Can they be entered in the app? *Checkpoint:* $K_c = 456.5\,\%/^{\circ}\mathrm{C}$, $\tau_I = 0.32\,\mathrm{h}$, $\tau_D = 0.08\,\mathrm{h}$. No.

9.6 Theory and calculation · 8 min

The level plant of Equation (9.17) has $K_L = 1.622 \times 10^{-4}\,\mathrm{m/h}$ of level per percent of makeup valve at the design position.

(a) Apply the averaging rule, Equation (9.25), for 26 % of valve travel on 0.10 m of level, and compute the damping ratio at the shipped integral time of 96 h. *Checkpoint:* 260 %/m, 94.8 h, $\zeta = 1.006$.

(b) Scenario 9 runs the level loop proportional only at 30 %/m, from a makeup valve at 2.367 % that passes the 0.0396 $\mathrm{m^3/h}$ of evaporation. A leak of 0.6 $\mathrm{m^3/h}$ opens and the pool temperature does not move. Use Equation (9.15) to find the valve position that passes the leak and the evaporation, and the level offset that holds it there. *Checkpoint:* 24.1 % and 0.72 m.

(c) Before the loop has acted, how fast does a leak of 0.05 $\mathrm{m^3/h}$ lower the level, and how many hours pass before the fall equals the 3 mm of transmitter noise? *Checkpoint:* 0.46 mm/h and 6.5 h.

9.7 App challenge · 8 min

Load scenario 3, *Step test: one pool, two plants (MANUAL)*, as shipped at the default seed. Run to 70 h at 20× and export the CSV file. Use the true pool temperature, column `pool_C`.

(a) Read `pool_C` at 6.0 and 24.0 h. Give the total change and the gain per percent of the 10.7 % CCW valve step. *Checkpoint:* −0.590 °C and −0.0551 °C per percent, within 0.003 °C and 0.0003 °C per percent.

(b) Find the first row after 6 h at which the pool has made 63.2 % of that change, and the first row at which it is 0.005 °C below its value at 6.0 h. Give both as times after the step, and estimate θ and τ. *Checkpoint:* 3.4 h and 0.2 h, so $\theta \approx 0.2$ h and $\tau \approx 3.2$ h.

(c) Both trains stop at 40 h. Compute the mean slope of `pool_C` from 40 to 52 h, and say why it is smaller than the 1.466 °C/h at the instant of the stop. *Checkpoint:* 1.429 °C/h, within 0.005 °C/h.

9.8 App challenge · 8 min

Load scenario 5, *PI from the identified model*, as shipped at the default seed, with $K_c = 23\,\%/°\text{C}$ and $\tau_I = 3.3$ h. The setpoint steps from 35 to 38 °C at 6 h and the cooling water supply rises from 30 to 38 °C at 40 h.

(a) Run to 100 h at 20× and export the CSV file. In column `pool_C`, find the first row after 6 h from which the pool stays within ±0.15 °C of 38 °C until 40 h, as a time after the step. Read the pool temperature and `ccw_pct` at 100 h. *Checkpoint:* 2.8 h within 0.1 h. 41.31 °C within 0.02 °C, valve at 100 %.

(b) Press Reset, open Advanced, set Pool RTD noise to 0 and Temperature integral time to 8.7 h, run to 40 h and export again. Make the same reading of the time to settle. *Checkpoint:* 13.9 h within 0.1 h.

9.9 App challenge · 8 min

Load scenario 11, *Evaporation feedforward, and why the level loop cannot do it*, at the default seed. Before Run, open Advanced and set Pool RTD noise and Level transmitter noise to 0. Run each case to 240 h at 20× and export the CSV file. The building turns hot and dry at 8 h and cool and damp at 132 h.

(a) On the PI architecture as loaded, read Makeup used on the scorecard, and the lowest value of `shield_m` with its time. *Checkpoint:* 38.5 m^3, and 6.830 m at 132.0 h, within 0.002 m.

(b) Press Reset, set Control architecture to PI + evap FF, and run again. Read the largest departure of `shield_m` from 7.000 m after 6 h, and Makeup used. *Checkpoint:* 0.8 mm within 0.3 mm, 46.8 m^3.

(c) From `q_evap_m3h` at 7.9 and 12 h, find the step in evaporation and the rate at which it lowers the level before any loop acts. Why can no level gain prevent this first fall? *Checkpoint:* about 0.21 m^3/h and 1.9 mm/h.

9.10 Application area · 4 min

Scenario 4 runs one train on/off at a 35 °C setpoint against the aged inventory, which one train wide open holds at 32.18 °C. Use Table 9.4.

(a) What is the widest deadband that still lets the train stop, and what happens at 6 °C? *Checkpoint:* 5.64 °C.

(b) A shift supervisor wants the band narrowed from 3 to 1.5 °C so the trend looks tidier. What does it cost and what does it buy? *Checkpoint:* 14 to 27 starts a week, swing 3.19 to 1.72 °C.

9.11 Application area · 4 min

In scenario 9 a leak-chase drain opens at 0.05 m^3/h at 8 h. At 96 h the metered makeup totals 4.63 m^3, the computed evaporation totals 3.82 m^3, and the inventory has fallen by 3.57 m^3.

(a) Estimate the volume lost to the leak and its mean rate since 8 h. *Checkpoint:* 4.38 m^3 and 0.050 m^3/h.

(b) The level deviation at 96 h is 33.0 mm. Which of the two, the level or the mass balance, is the better leak detector, and what does the mass balance need that the level does not? *Checkpoint:* The mass balance.

9.12 Review the agent · 5 min

An AI assistant drafts an operator aid for a loss of pool cooling: "Compute the time to boiling as $(100 - T)/(\mathrm{d}T/\mathrm{d}t)$ and the time to the shielding minimum as $(z - z_r - 6.40)/(-\mathrm{d}z/\mathrm{d}t)$ from the present trends. Both answers are conservative, because evaporation only grows as the pool heats, so the pool always reaches each limit later than the straight line says." Find the flaw, and say how you would check the claim in the app.