INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 1

Introduction: Controlling Reactors and Plasmas

Pages 1–23

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Controlling Reactors and Plasmas," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–23.

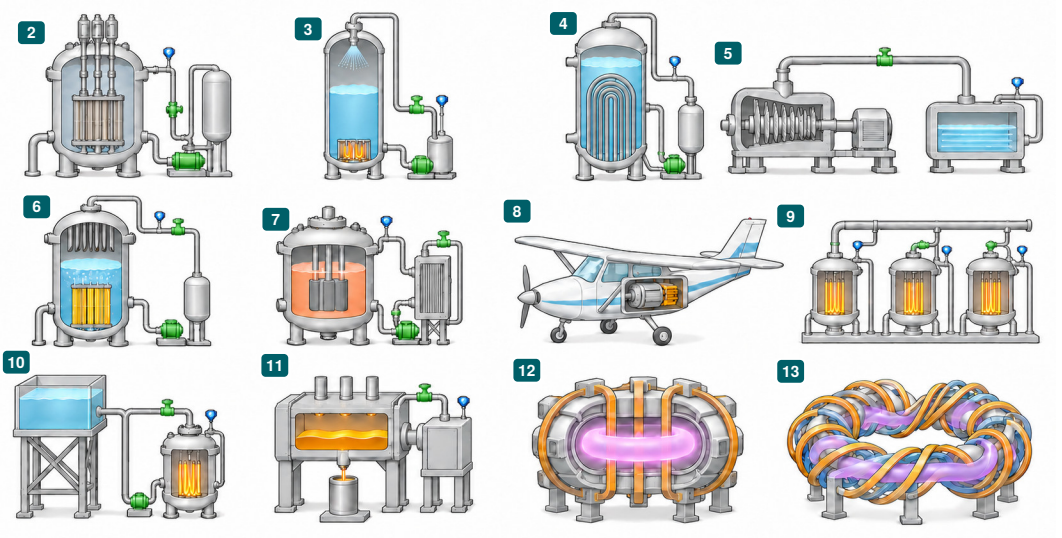
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 1

Introduction: Controlling Reactors and Plasmas



The chapters 2 PWR rod control 3 pressurizer 4 steam generator level 5 turbine trip and dump 6 boiling water reactor 7 molten salt reactor 8 aircraft reactor 9 three-module header 10 decay-heat removal 11 glass melter 12 tokamak position 13 stellarator exhaust

Notice that the top row follows the steam of a pressurized water reactor from core to condenser, the middle row holds the other reactors, and the bottom row shows what happens after a trip, at the end of the fuel cycle and in the two fusion machines. Illustration, not to scale.

IN PLAIN TERMS **Keeping Reactors and Hot Gas Within Their Limits**

The main idea. A nuclear reactor makes heat by splitting atoms. A fusion machine tries to make heat by joining them, inside a gas so hot that magnets must hold it.

The challenge. Much of the control is done by the plant itself, because a hotter core slows its own reaction. What is left is hard. A water level first moves the wrong way, a lost load leaves only seconds before a safety valve opens, and the hot gas falls toward the wall faster than a blink.

Why it matters. Good control keeps the plant running and leaves the safety systems as a last line of defense that is rarely called on.

THREE WAYS TO READ THIS CHAPTER

Explore	8 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.3 and 1.4.6, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the examples R-1 and P-1, then Exercises 1 to 10.
Practice	10 min	This page, Section 1.4.5, Section 1.4.7 with Table 1.1, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A nuclear plant is controlled around physics that acts faster than any controller and around limits that no controller may cross. A pressurized water reactor follows a cut in steam demand with no command at all, because hotter fuel and hotter water remove reactivity. A boiling water reactor uses its steam

voids at once as the missing moderator, as the strongest feedback in the core and as an error in its own level reading. An elongated tokamak plasma falls toward the vessel wall within milliseconds unless a coil supply pushes it back. Around each of these plants sit protection limits: a pressure at which safety valves lift, a temperature, a water level, a heat flux on a target. This volume answers one question: how do a plant's own physics, its protection limits and its layers of control fit together, so that each control layer does only what the physics leaves undone and stays clear of the protection?

Seven control problems define the industry, and every chapter meets several of them.

- *The plant regulates itself, and the controller removes what is left.* Negative temperature coefficients of reactivity bring fission to the load and leave a temperature offset for the rods or the pumps.
- *Power follows reactivity on a clock set by delayed neutrons.* Less than one neutron in a hundred is delayed, and that fraction makes a reactor slow enough for mechanical actuators. Circulating fuel carries part of it out of the core.
- *Slow poisons and residual heat outlast every transient.* Xenon builds for hours after a power cut, and the fission products go on making heat for days after the reactor stops.
- *One measurement counts two things.* A drum or vessel level reads the water and the steam held under its surface together, and the two answer a feedwater change in opposite directions.
- *Some disturbances arrive faster than their measurement.* A turbine that disconnects from the grid leaves seconds before the safety valves lift, against instruments and valves that take seconds each.
- *Protection outranks control.* Trips, safety valves, insertion limits and coping times are fixed first, and every control layer is designed to stay clear of them.
- *A plasma is unstable or constrained on its own clock.* An elongated tokamak plasma must be held within milliseconds by a supply that saturates, and a long-pulse plasma must radiate its exhaust without poisoning its core.

Twelve chapters answer the question on twelve plants, each with an app that the reader runs in a browser. Chapters 2 to 5 follow a pressurized water reactor from its control rods, through its pressurizer and its steam generator, to a turbine that has just disconnected from the grid. Chapter 6 boils the water in the core. Chapters 7 and 8 dissolve the fuel in a circulating salt, the second on board an aircraft. Chapter 9 joins three small modular reactors on one steam header. Chapter 10 removes the heat of a reactor that has tripped, and Chapter 11 turns high-level waste into glass. Chapters 12 and 13 hold the position and the heat exhaust of two fusion plasmas.

A reader who has worked through the volume can do the following on a nuclear or fusion plant that is new to the reader.

- Write the reactivity ledger of a reactor, say what the inherent feedback does after a load change, and find the offset left for the controller.
- Compute a prompt jump, a reactor period, the delayed neutrons a circulating fuel keeps and the xenon that a power change brings.
- Recognize a level that measures two things, locate its right-half-plane zero and say what feedforward removes and what it leaves.
- Count the seconds between a disturbance and a protection limit, and say which signals can act inside them.
- Size the heat a tripped core still makes and the water it boils.
- State the gain floor, the gain ceiling and the recoverable region of an unstable plasma, and the radiated fraction that its exhaust needs.

This chapter develops each of these on small examples and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward and cascade from Volume 3, split range, selectors and overrides from Volume 4, dead time and estimation from Volume 5, pairing and allocation from Volume 6, unstable plants and state-space design from Volume 7, and predictive control from Volume 9. Volume 17 dispatches the heat of a nuclear plant between electricity, storage and hydrogen.

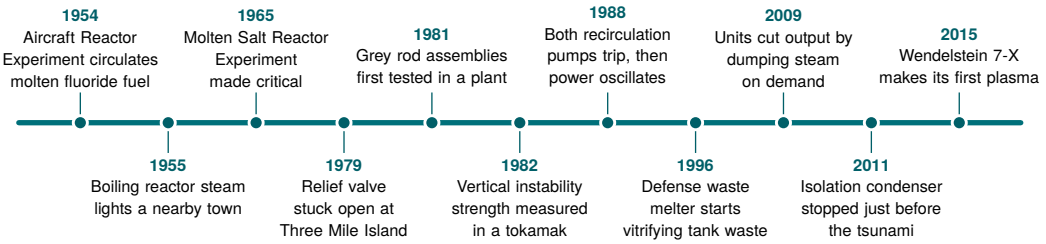


Figure 1.1 Milestones of nuclear and fusion plants and their control. The dates are from [1] (1954), [2] (1955), [3] (1965), [4] (1979), [5] (1981), [6] (1982), [7] (1988), [8] (1996), [9] (2009), [10] (2011) and [11] (2015). Notice that the fission plants span the whole period, while the fusion devices enter only in 1982. Events are evenly spaced, not to a time scale.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first experiments with circulating fuel and boiling water to a stellarator built for long pulses.

> ⧖ HISTORY **1954 to 1956: Reactors That Followed Their Load**
>
> The Aircraft Reactor Experiment at Oak Ridge, part of the aircraft nuclear propulsion project, ran in November 1954. Its molten fluoride fuel was pumped through the core and on to heat exchangers, and while at power it was easily controlled, because its high negative temperature coefficient of reactivity made the reactor a slave to the load placed upon it [1]. In Idaho the first test boiling water reactor proved that such a reactor regulates itself, and on the night of July 17, 1955, a successor with a turbine generator lit a nearby town. The experimental boiling water reactor that followed ran at constant power, and an automatic steam valve took the swings in electric demand [2]. The same self-regulation carries Chapter 7 and Chapter 6, and the same bypass of steam returns in Chapter 5.

The Molten Salt Reactor Experiment was made critical in June 1965 by adding enriched uranium to its carrier salt [3], the kind of circulating core whose delayed neutrons Section 1.4.3 counts. In March 1979 a relief valve on a reactor at Three Mile Island stuck open, and much of the core was uncovered and severely damaged [4], the events behind Chapter 3 and Chapter 4. Grey rod assemblies, several times less efficient than the ordinary rods, were first tested in a French plant in 1981 and, with an average coolant temperature allowed to vary, significantly raised the maneuvering capability [5], the load following of Chapter 2. In 1982 a divertor tokamak measured how unstable its elongated plasma was by holding the input of its vertical feedback at zero and timing the drift [6], the instability of Chapter 12. In March 1988 both recirculation pumps of a boiling water reactor tripped, and within a few minutes its power began to oscillate until an overpower scram [7], the oscillation that Chapter 6 simulates.

A melter in the United States has vitrified high-level waste from tanks since 1996 [8], the process of Chapter 11. From 2009 a surplus of baseload generation in one Canadian province led selected nuclear units to hold their reactors at full power and send steam to the condenser to cut their output [9], the bypass that Chapter 9 uses. In March 2011 the operators of a boiling water reactor at Fukushima Daiichi stopped its isolation condenser just before the tsunami struck [10], the passive heat sink of Chapter 10. In December 2015 the stellarator Wendelstein 7-X made its first plasma after more than 15 years of construction [11], the machine behind Chapter 13. The aircraft reactor of Chapter 8 returns to the project that began the list.

> ⧖ HISTORY **1979: Instruments Not Built for an Accident**
>
> The commission that investigated the accident at Three Mile Island found its control room adequate for normal operation but seriously deficient in an accident. More than 100 alarms went off in the first few minutes, with no way to suppress the unimportant ones. Several instruments went off-scale, because they were not designed to follow an accident, and the alarm printer ran more than two hours behind. Pressure and temperature were both shown, yet nothing said directly that together they meant the cooling water was turning into steam. The commission asked for instruments that read the full range of temperatures in the reactor vessel and the actual position of valves, with the key measurements displayed clearly by computer [4].

1.3 Notation and Conventions

Nuclear power n is a fraction of the rated thermal power P_0, which is in megawatts. Reactivity ρ is the fractional excess of neutron production over loss, in pcm ($1\,\text{pcm} = 10^{-5}$). Divided by the delayed neutron fraction β it is in dollars, and a hundredth of a dollar is a cent. The prompt neutron generation time is Λ, the decay constant of a delayed neutron precursor group is λ_i, and α_F and α_M are the fuel (Doppler) and moderator temperature coefficients. Temperatures are in °C and their differences in kelvin, and pressures are in bar or MPa as each chapter gives them. Most chapters count time in seconds. Chapter 2 counts it in minutes, Chapter 11 in minutes and days, Chapter 10 in minutes and hours and Chapter 12 in milliseconds.

The framework's h is its integration step and its a the controller action, and neither is used for anything else here. Loop action follows Section 1.6.2: $a = +1$ is reverse acting, for a process whose measurement rises with the manipulated variable, and $a = -1$ is direct acting. A rod loop on coolant temperature is reverse acting, because withdrawing the rods makes the coolant hotter. A pump or valve that lowers what it controls is direct acting, as are the fuel pump on the core outlet temperature of a molten salt reactor, the steam dump on average coolant temperature and the turbine valves on the dome pressure of a boiling water reactor. The vertical position loop of a tokamak is reverse acting, because a positive coil voltage moves the plasma up.

The chapters keep their own symbols, and several collide. λ is a precursor decay constant in the fission chapters, the iodine and xenon decay constants λ_I and λ_X in Chapter 2 and the unstable eigenvalue in Chapter 12. γ is the direct xenon yield in Chapter 2, the vertical growth rate in Chapter 12, the ratio of specific heats in Chapter 9 and a loading acceleration in Chapter 11. β is the delayed neutron fraction except in the β drop of Chapter 12, a loss of plasma pressure. Λ is the generation time in the fission chapters and the radiating-layer parameter in Chapter 13. κ is the elongation in Chapter 12, a profile factor κ_b in Chapter 13, a schedule exponent in Chapter 8, a turbine exponent in Chapter 9 and an off-gas restriction in Chapter 11. δ is the control latency in Chapter 12, a pressure ratio in Chapter 8 and an initial delayed share in Chapter 10. W is a rod worth $W(z)$ and a scaled precursor population W_i in Chapter 2, Chapter 7 and Chapter 8, and the plasma stored energy in Chapter 13. z is a bank position in those three chapters and the vertical plasma position in Chapter 12. n is the power fraction in the fission chapters, the decay-heat exponent in Chapter 10 and, as $\bar{n}$, the line-averaged density in Chapter 13. τ_c is the core transit time of a circulating fuel in Chapter 7 and Chapter 8 and a core energy time in Chapter 13. L is the turbine load fraction in Chapter 2, a level in Chapter 4 and Chapter 5, and an inductance L_w in Chapter 12. T_t is the framework's tracking time and the divertor target temperature of Chapter 13. a is the xenon burnup rate in Chapter 2 and the minor radius in Chapter 13, neither of them the action, and φ is the cold-cap coverage of Chapter 11. This chapter writes the xenon burnup rate r_X.

1.4 Physics, Limits and Control Layers

1.4.1 The Flow Scheme

Figure 1.2 places each chapter's plant on the path from fission or fusion to the user. In a pressurized water reactor the core heats water that is kept liquid by a pressurizer, a vessel whose single steam bubble sets the pressure of the whole circuit. The primary water gives its heat to steam generators,

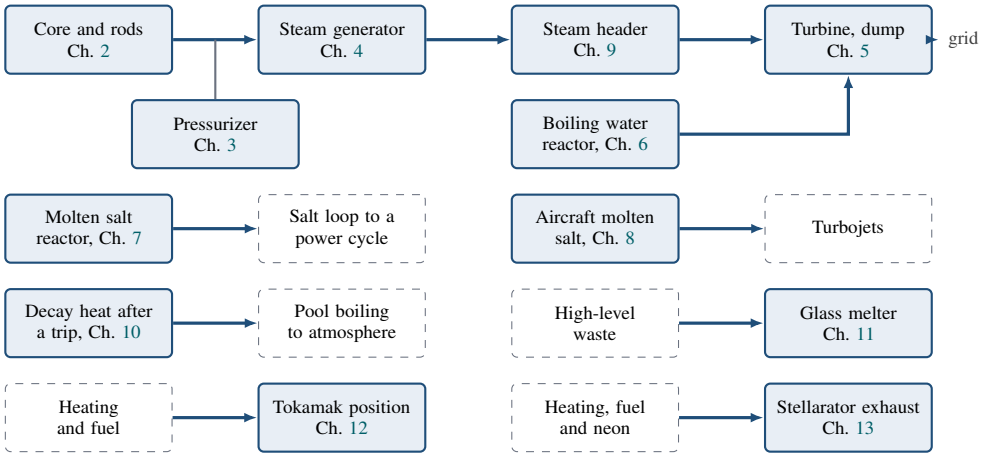


Figure 1.2 The Plants of the Volume. Filled Blocks Are the Plants of Chapters 2 to 13, Dashed Blocks Lie Outside the Volume. Top Two Rows: a Pressurized Water Reactor From Core to Grid, With the Pressurizer on Its Primary Circuit, and a Boiling Water Reactor That Sends Core Steam to the Turbine. Third Row: Fuel Dissolved in a Circulating Salt. Fourth Row: Heat Removal After a Trip and the Glass Melter at the Back End of the Fuel Cycle. Bottom Row: the Two Fusion Plasmas.

the steam passes through a header to the turbine, and a dump to the condenser takes the steam that the turbine refuses [12]. A boiling water reactor makes its steam in the core and sends it straight to the turbine. A molten salt reactor dissolves its fuel in a salt that a pump carries through the core and an external heat exchanger [13]. Several small modules can share one header. When a reactor trips, its decay heat must still be carried to a heat sink, and at the back end of the fuel cycle the high-level waste from reprocessing is melted into glass [14]. Fusion devices confine a hot plasma in a magnetic field. A tokamak drives a current in the plasma and must hold its vertical position [15], and a stellarator makes its field with external coils alone and runs long enough that its exhaust must be handled in a steady state.

The techniques that the fission chapters share are developed below on one running example, a small pressurized water reactor called R-1, rated at 300 MW of heat. Its delayed neutrons are lumped into one group with $\beta = 650\,\text{pcm}$ and $\lambda = 0.08\,\text{s}^{-1}$, and its prompt generation time is $\Lambda = 2 \times 10^{-5}\,\text{s}$. Its Doppler coefficient is $\alpha_F = -3\,\text{pcm/K}$ on the fuel temperature and its moderator coefficient $\alpha_M = -12\,\text{pcm/K}$ on the coolant temperature. At full power the fuel runs 300 K above the coolant with a time constant of 5 s, the primary coolant holds 90 MJ/K, and the control bank is worth 15 pcm per percent of travel at its working position. R-1's xenon, steam generator, trip and decay heat are declared where they are used. No number of R-1 applies to a plasma, so Section 1.4.6 declares a second example, P-1. Every number of both is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Point Kinetics and Reactivity Feedback

Reactor power changes only while reactivity differs from zero. Most neutrons are born within microseconds of a fission, but a fraction β is emitted seconds to a minute later by the decay of fission products called delayed neutron precursors [16]. Point kinetics treats the core as one neutron population and one or more precursor groups [17]. The prompt neutrons come into balance with the precursors in a time $\Lambda/(\beta - \rho)$, 3.4 ms for R-1 at small reactivity, and the chapters set that fast mode to zero, which is the prompt jump approximation [18]. With one precursor group scaled by the

generation time, $W = \Lambda c$, the model is Equation (1.1).

$$\frac{\mathrm{d}W}{\mathrm{d}t} = \beta n - \lambda W, \qquad n = \frac{\lambda W}{\beta - \rho}, \qquad \frac{n(0^+)}{n(0^-)} = \frac{\beta}{\beta - \rho} \tag{1.1}$$

A step of reactivity changes n at once by the prompt jump on the right and then changes W, and so n, at the rate of Equation (1.2), the inverse of the stable period T.

$$\frac{1}{n}\frac{\mathrm{d}n}{\mathrm{d}t} = \frac{1}{W}\frac{\mathrm{d}W}{\mathrm{d}t} = \frac{\lambda\rho}{\beta - \rho} = \frac{1}{T} \tag{1.2}$$

On R-1 a step of 65 pcm, a tenth of a dollar, raises the power by 11.1 % at once and then multiplies it by e every 112.5 s. With nothing to stop it, the power is 1.89 times rated a minute after the step. At $\rho = \beta$ the denominator vanishes, the reactor is prompt critical and the prompt mode can no longer be dropped. Chapter 2, Chapter 5, Chapter 7 and Chapter 8 use six precursor groups with this approximation, and Chapter 2 annunciates prompt criticality instead of integrating through it. Chapter 6 keeps the prompt term with two groups.

Something always stops it, because reactivity depends on temperature. Equation (1.3) is the reactivity ledger of R-1 about its design point, with rod, fuel, coolant and xenon terms.

$$\rho = \Delta\rho_{rod} + \alpha_F \, \Delta T_f + \alpha_M \, \Delta T_c + \Delta\rho_{Xe} \tag{1.3}$$

Every steady state of Equation (1.1) has $\rho = 0$, because a steady $W = \beta n/\lambda$ gives $n = \beta n/(\beta - \rho)$. A settled reactor is therefore a statement about the ledger, and its power is set by the heat that the steam takes away. The fuel runs above the coolant by an amount proportional to power, 300 K at full power on R-1, so with the rods and xenon held Equation (1.3) fixes how far the coolant moves, as Equation (1.4) shows.

$$\Delta T_c = -\frac{\Delta\rho_{rod} + \alpha_F \, \Delta T_{fc,0} \, \Delta n}{\alpha_F + \alpha_M}, \qquad K_p = \frac{\mathrm{d}T_c}{\mathrm{d}z} = -\frac{\mathrm{d}W/\mathrm{d}z}{\alpha_F + \alpha_M} \tag{1.4}$$

where $\Delta T_{fc,0}$ is the fuel-to-coolant difference at full power and K_p is the steady gain from bank position z to coolant temperature at fixed power, with $\mathrm{d}W/\mathrm{d}z$ the differential worth of the bank.

Figure 1.3 shows both effects on R-1. With the heat sink held at full power, a 65 pcm withdrawal makes the prompt jump, but the fuel and then the coolant warm, and the power is back within 0.5 % of rated 48 s after the step with the coolant 4.33 K warmer. Withdrawing the rods made the plant hotter and not more powerful. With the rods parked and the load cut by 10 %, the power follows the load without any command and stays within 0.2 % of the new load from 94 s after the cut. The fuel-to-coolant difference falls by 30 K, the Doppler term returns 90 pcm, and the coolant settles 6.0 K warmer, where its own coefficient takes the 90 pcm back. Removing that offset is the controller's whole task. With $K_p = 1.0$ K per percent the rods must go in by 6 %.

> **+ DID YOU KNOW? Sixty-Nine Rod Tubes in One Core**
>
> The core of the reactor at Three Mile Island held 36,816 fuel rods in 177 fuel assemblies. Tubes in the assemblies held the instruments, for example to measure the temperature inside the core, and the control rods: 52 tubes with instruments and 69 with control rods [4].

Chapter 2 is this example at full scale. Its reactor goes from 100 to 89.9 % power within eight minutes of a load cut with every rod stationary and leaves the coolant 13.9 °C above its programmed average. Its rod gain is 2.58 °C per percent, from Equation (1.4) with a differential worth of 18.09 pcm per percent and coefficients that sum to −7 pcm/°C. Its drive takes a rod speed, so the actuator supplies the integrator and a proportional controller leaves no offset. In Chapter 5 a moderator coefficient of −30 pcm/°C takes the power from 100 to 73 % within 14.8 s of a full-power shut-in without a rod

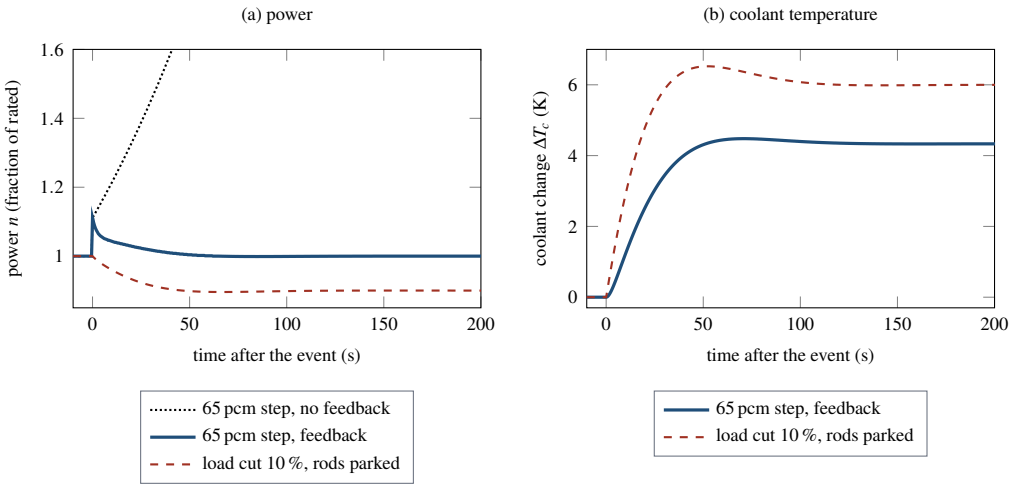


Figure 1.3 R-1 After a Rod Withdrawal of 65 pcm at a Fixed Heat Sink and After a 10 % Load Cut With the Rods Parked. (a) Power Without Feedback (Dotted, Rising on a 112.5 s Period), With Feedback (Solid, Returning to Rated) and After the Load Cut (Dashed, Settling at 90 %). (b) The Coolant Settles 4.33 K Warmer After the Step (Solid) and 6.0 K Warmer After the Load Cut (Dashed).

moving. Chapter 6 measures 30.6 pcm of negative feedback per percent of power, mostly from voids, whose coupling to the neutrons can also drive the density-wave oscillations that it simulates [19], so a rod insertion of 366 pcm lowers the power by only 11.46 % and the recirculation pumps, not the rods, follow the load. In Chapter 7 and Chapter 8 the ledger pins a weighted core temperature, so a load change moves the temperature rise across the core and leaves its mean almost where it was.

1.4.3 Delayed Neutrons and Xenon

The six precursor groups of uranium-235 that Chapter 2 and Chapter 7 use add up to $\beta = 650.2$ pcm, and their mean delay, weighted by each group's share, is 13.0 s, an equivalent single decay constant of 0.0767 s^{-1} that R-1 rounds to 0.08. In a solid core each precursor decays where it was born. In a circulating fuel the pump carries it out of the core, and a neutron emitted in the external loop is lost to the chain. A group of decay constant λ_i that spends τ_c in the core and τ_L outside keeps the fraction of its worth given by Equation (1.5), the steady state of the circulating precursor balance that Chapter 7 and Chapter 8 integrate.

$$\beta_{eff} = \sum_{i=1}^{6} \beta_i \frac{\lambda_i \tau_c}{\lambda_i \tau_c + 1 - \mathrm{e}^{-\lambda_i \tau_L}} \tag{1.5}$$

With 9 s in the core and 16 s outside, the six-group sum is 419.4 pcm, the value of Chapter 7, and with 6 and 10 s it is 390 pcm, the value of Chapter 8. A single group at 0.08 s^{-1} gives 325 pcm for the first case and is too crude here, because the short-lived groups decay inside the core and keep most of their worth while the long-lived ones decay outside. Three consequences follow. A circulating core is critical at positive reactivity, $\rho = \beta - \beta_{eff}$, which its rods must hold, 231 pcm at the design flow of Chapter 7. The dollar moves with the pump: the same 100 pcm is 0.247 dollars at 120 % pump speed and 0.203 at 40 %. A pump trip hands the held reactivity back, and in that chapter a trip at low power takes the power from 9.0 to 16.8 %.

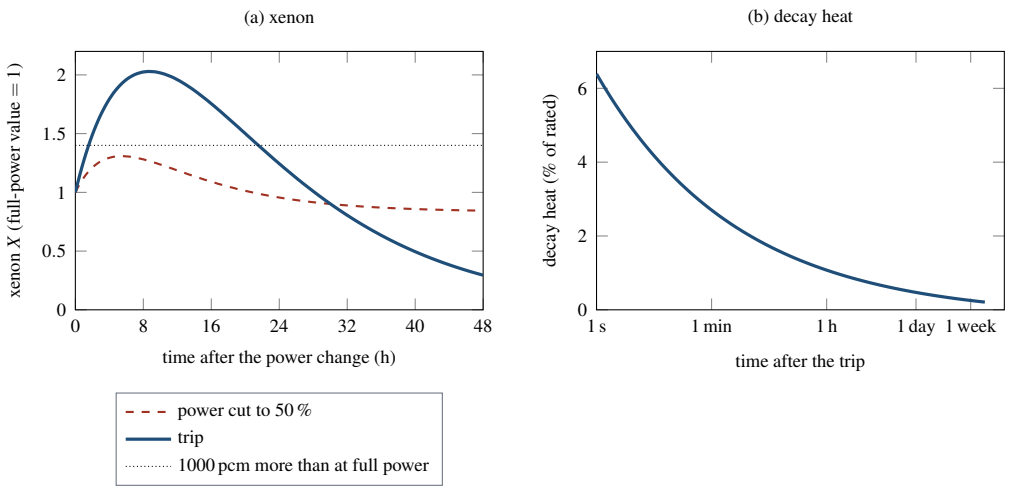


Figure 1.4 The Two Slow Clocks of R-1. (a) Xenon After a Power Cut to 50 % (Dashed) and After a Trip (Solid). Above the Dotted Line the Extra Xenon Exceeds 1000 pcm. (b) Decay Heat After a Year at Full Power, From Equation (1.9).

> **+ DID YOU KNOW? A Reactor Run of 221 Hours**
>
> From first criticality to final shutdown, the Aircraft Reactor Experiment ran for 221 hours. For the last 74 of them its power was between 0.1 and 2.5 MW, and it made about 96 MWh in all [1]. R-1 makes as much heat in about 19 minutes.

Xenon-135 is the strongest neutron absorber in the core. Most of it comes from the decay of iodine-135, with half lives of 6.71 h for iodine and 9.21 h for xenon, and at power it is destroyed mainly by neutron capture [12]. Normalized to one at full-power equilibrium and ignoring the few percent made directly in fission, the balance is Equation (1.6), with r_X the burnup rate at rated flux.

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \lambda_I (n - I), \qquad \frac{\mathrm{d}X}{\mathrm{d}t} = (\lambda_X + r_X) I - (\lambda_X + r_X n) X, \qquad X_\infty = n \frac{\lambda_X + r_X}{\lambda_X + r_X n} \tag{1.6}$$

R-1 burns xenon four times faster than it decays, $r_X = 4\lambda_X$, and its equilibrium xenon is worth 2500 pcm. After a cut from 100 to 50 %, capture falls at once while iodine keeps feeding xenon, so xenon rises for 5.5 h to 1.31 times its full-power value, a further 770 pcm that the rods must withdraw to offset, before it falls toward 0.83. After a trip it peaks at 2.03 times its full-power value 8.75 h later, 2570 pcm more poison, and it does not fall back below its full-power value until 28 h after the trip. A reactor with 1000 pcm of spare reactivity cannot restart from about 1.75 h to 21.5 h after the trip, as Figure 1.4a shows. Chapter 2 uses Equation (1.6) with the direct yield included, and its sixteen-hour maneuver to 60 % load and back walks the bank through 60, 39, 68, 60, 100 and 44 %, exhausting its withdrawal travel.

1.4.4 Shrink and Swell

A drum or vessel level transmitter reads the height of a two-phase mixture, which is the water plus the steam bubbles held below the surface [20]. Subcooled feedwater condenses some of those bubbles, so the level first falls when the feed rises, which is shrink. A pressure drop expands them, so the level first rises when the steam demand rises, which is swell. Written as deviations, the level after a feed change q_f is a water term that integrates the flow and a void term that settles with the bubbles, as

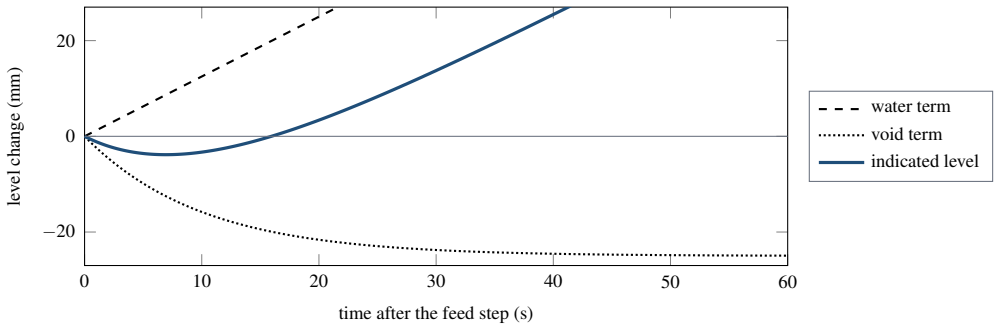


Figure 1.5 Shrink on R-1's Drum After a Feed Step of 10 kg/s. The Water Term (Dashed) Rises at 1.25 mm/s, the Void Term (Dotted) Settles at −25 mm With a Time Constant of 10 s, and the Indicated Level (Solid) Falls to −3.84 mm at 6.93 s Before Rising.

shown in Equation (1.7).

$$\frac{L(s)}{q_f(s)} = \frac{K_w}{s} - \frac{K_v}{\tau_v s + 1} = \frac{K_w\left[1 - (K_v/K_w - \tau_v)\, s\right]}{s\,(\tau_v s + 1)}, \qquad z_{rhp} = \frac{K_w}{K_v - K_w \tau_v} \tag{1.7}$$

where $K_w = 1/(\rho_w A_d)$ for water density ρ_w and surface area A_d, K_v is the settled fall per unit of extra feed and τ_v the time constant of the voids. When $K_v > K_w \tau_v$ the transfer function has a zero in the right half plane at z_{rhp}, and the step response first goes the wrong way. A zero at z_{rhp} bounds the crossover frequency of any feedback loop to about $z_{rhp}/2$ [21].

R-1's steam generator drum has a surface of 10 m^2 over water of 800 kg/m^3, so $K_w = 0.125$ mm per kilogram. Each extra kilogram per second of feed collapses enough steam to lower the reading by 2.5 mm once the bubbles have settled, with $\tau_v = 10$ s. The zero is at 0.1 s^{-1}. A step of 10 kg/s first lowers the level at 1.25 mm/s, reaches a minimum of −3.84 mm at $\tau_v \ln[K_v/(K_w \tau_v)] = 6.93$ s, returns through the starting level at 15.9 s and then rises at 1.25 mm/s, as Figure 1.5 shows. A level loop on this drum should cross over below 0.05 rad/s, a closed-loop time constant of 20 s or longer. At low load the circulation is slower and more steam waits under the surface per kilogram made. With $K_v = 4.0$ mm per kilogram per second and $\tau_v = 14$ s, the zero moves in to 0.056 s^{-1} and the level returns through its start only after 27.5 s.

Three-element control answers the integrating part of the problem. A feedforward from steam flow to feedwater flow, held by a cascade on the feed flow (the structure of Volume 3), keeps the inventory in balance as the load moves, and the level controller only trims. The feedforward removes the mass imbalance. It does not remove the void term, which still appears on the level and which the trim must not chase. Chapter 4 measures the moving zero by its crossing time, 38 s at full load and 65 s at a quarter, and reports no zero location because the zero moves with load. It compares single-, two- and three-element control on a load rejection and a pump trip. In Chapter 6 a 10 % feedwater step drives the level indication −40.9 mm while the vessel gains 193 kg, and on a recirculation ramp three-element control holds the level to −63.0 mm against +287.4 mm for single-element control at identical gains.

1.4.5 Protection Limits

A reactor plant is controlled in layers that act in a fixed order. Regulating control holds setpoints and is tuned for performance. Limiting control acts before a limit is reached: a runback cuts reactor power, an override takes a valve from its controller, an insertion limit stops the rods. Protection acts last and on its own instruments: the reactor trips, safety valves lift, interlocks close. No control layer may defeat a protection action, and each protection setpoint sits between the operating point and the limit it protects [12].

When a heat sink is lost, that time is the margin divided by the rate at which the imbalance consumes it, as shown in Equation (1.8).

$$t_{lim} = \frac{\Delta p_{margin}}{(\mathrm{d}p_{sat}/\mathrm{d}T)\,\Delta P/(Mc)_s} \tag{1.8}$$

where ΔP is the heat no longer taken away, $(Mc)_s$ the heat capacity of the secondary side and $\mathrm{d}p_{sat}/\mathrm{d}T$ the slope of the saturation curve. R-1's secondary holds 270 MJ/K and its saturation pressure rises 0.9 bar/K, so a shut-in at full power raises the steam pressure at 1.0 bar/s, and a 10 bar margin to the first safety valve lasts 10 s. A loop that reads a 4 s temperature detector and moves a valve that strokes in 3 s spends 70 % of the clock before the valve is open, whatever its gain. Only a signal from upstream of the effect, the breaker or the turbine's own pressure, acts inside the clock. That is the case of Chapter 5: its secondary reaches the first safety valve 14.80 s after a full-power shut-in, with only 6.70 bar between the no-load pressure and that valve. A temperature loop peaks at 83.37 bar and discharges 1.53 t, the same loop given turbine impulse pressure peaks at 79.74 bar and lifts nothing, and two low-selected overrides with a reactor runback peak at 76.39 bar.

A trip stops the chain reaction within seconds, and the delayed-neutron tail dies away within minutes, but the fission products go on decaying. Their heat after a time t at shutdown, for a reactor that ran at P_0 for a time T_0, is approximated by Equation (1.9), with t and T_0 in seconds [22, 23].

$$\frac{Q_d}{P_0} = 0.066\left[t^{-0.2} - (t+T_0)^{-0.2}\right], \qquad m_{boil} = \frac{1}{h_{fg}}\int Q_d\,\mathrm{d}t \tag{1.9}$$

After a year at power R-1 still makes 3.96 % of its rating 10 s after the trip, 1.07 % (3.2 MW) after an hour and 0.47 % after a day, as Figure 1.4b shows. Its first day after the trip releases 46.1 MWh, enough to boil 74 t of saturated water at 2257 kJ/kg. Each megawatt boils 1595 kg an hour. A passive heat sink therefore sets a coping time, the hours until its water is spent, and natural circulation carries the heat to it with no pump [24]. The plant's stored heat is a second account on the same water, at an exchange rate of the heat capacity over the latent heat: R-1's primary holds 120 MJ/K, so every kelvin of cooldown boils 53.2 kg of pool water.

Chapter 10 has no setpoint at all. Its core makes 33.9 MW of decay heat at the trip on a 530 MW rating, and its exchange rate is 165.2 kg of pool water per kelvin of avoided cooldown. A latched-open condenser takes the core outlet to 211.9 °C at four hours and leaves 15.72 t in the pool, while throttling to 300 °C leaves 29.92 t and moves the hour at which the tube bundle uncovers from 3.93 to 6.26. The other chapters meet their limits the same way. Chapter 3 holds the primary pressure between heaters and spray whose full actions move it at +1.232 kPa/s and −11.513 kPa/s, and it adds a high-pressure override. In Chapter 9 each module stops at an average coolant temperature limit, not at its licensed power, and that ceiling falls as the header pressure rises. Chapter 8 holds a fuel salt temperature limit, and Chapter 11 protects a pool temperature floor and an off-gas ceiling with overrides on a low selector, the structures of Volume 4.

1.4.6 Plasma Position and Exhaust

A tokamak stretches its plasma vertically to confine more current in the same magnet, and the field that stretches it also pushes it away from the midplane. Eddy currents in the conducting wall slow the fall, so the instability grows at a rate set by the wall time constant τ_w and a stability margin m, $\gamma\tau_w = 1/m$ [15, 25]. Reduced to its unstable mode, with the plasma position z, a coil voltage V limited to $\pm V_{max}$ and a proportional loop acting after a latency δ, the plant and its loop are Equation (1.10).

$$\frac{\mathrm{d}z}{\mathrm{d}t} = \gamma z + bV, \qquad V(t) = -K\,z(t-\delta), \qquad K > \frac{\gamma}{b}, \qquad |z| < \frac{bV_{max}}{\gamma} \tag{1.10}$$

The first inequality is the gain floor: below it the coil pushes back more weakly than the plasma runs away. The second is the recoverable region: beyond it even the full voltage cannot reverse the motion,

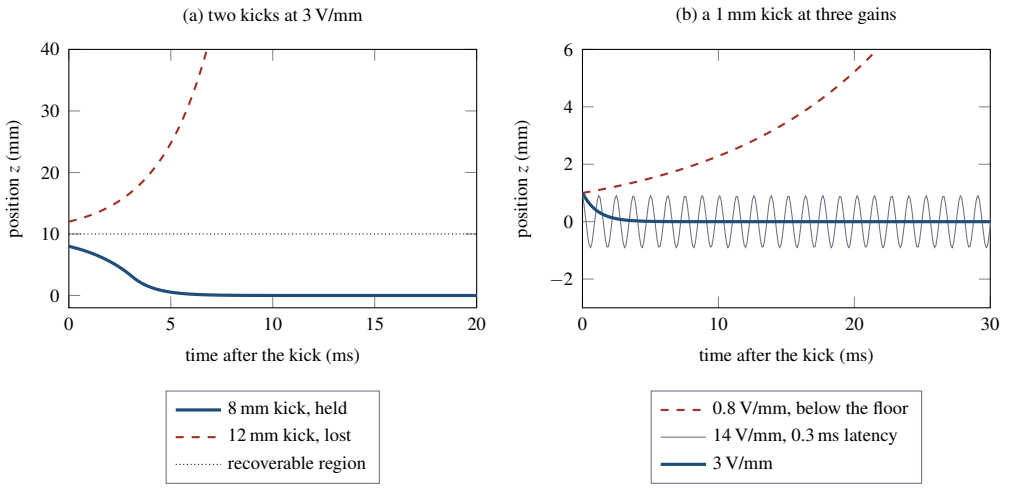


Figure 1.6 Vertical Position of P-1 Under a Proportional Loop With a ±10 V Supply. (a) At 3 V/mm and 0.1 ms Latency an 8 mm Kick Inside the 10 mm Region (Dotted) Is Held (Solid) and a 12 mm Kick Is Lost (Dashed). (b) A 1 mm Kick Drifts to the Wall Below the Floor (Dashed, Leaving the Axis), Is Held at 3 V/mm (Thick Solid) and Cycles Above the Ceiling (Thin Solid).

whatever the gain. The latency sets a ceiling. On the stability boundary $s = j\omega$ the characteristic equation $s = \gamma - bKe^{-s\delta}$ splits into Equation (1.11).

$$bK\cos\omega\delta = \gamma, \qquad bK\sin\omega\delta = \omega \quad \Rightarrow \quad K_{max} = \frac{\sqrt{\gamma^2+\omega^2}}{b}, \qquad \tan\omega\delta = \frac{\omega}{\gamma} \tag{1.11}$$

The window between floor and ceiling closes when $\gamma\delta$ reaches one, so the latency must be shorter than the growth time $1/\gamma$.

The second example, P-1, is a small elongated plasma with $\gamma = 400\,\mathrm{s}^{-1}$, a growth time of 2.5 ms, a coil authority $b = 400\,\mathrm{mm/s}$ per volt, a supply of $\pm 10\,\mathrm{V}$ and a wall 150 mm away. Its gain floor is 1.0 V/mm and its recoverable region 10 mm. Its ceiling is 38.6 V/mm at 0.1 ms of latency and 12.5 V/mm at 0.3 ms, and the window closes at 2.5 ms. Figure 1.6 shows all three limits. At 3 V/mm an 8 mm kick saturates the supply for 3.1 ms and is held, while a 12 mm kick reaches the wall 10.6 ms later at that gain and at ten times it. A 1 mm kick is lost after 40.6 ms at 0.8 V/mm, held at 3 V/mm, and at 14 V/mm with 0.3 ms of latency it settles into a limit cycle of 0.63 mm rms with the supply saturated for 42 % of the time after the first 40 ms.

Chapter 12 measures the same three limits on a plasma, wall and coil model whose unstable mode grows at $356\,\mathrm{s}^{-1}$ at elongation 1.8. Its floor is about 0.45 V/mm there and about 0.70 V/mm at elongation 2.0, its loop settles into a 1.45 mm limit cycle at 10 V/mm with 300 µs of latency, and a 20 mm kick at elongation 1.9 is lost at the default gains and sooner at 10 V/mm, but held with a stronger velocity term, a 15 V supply or at elongation 1.8. Gain decides how the loop behaves inside the region, and the growth rate and the supply decide how large the region is [26].

A plasma held in a steady state must also get rid of its heating. Whatever the plasma does not radiate reaches the divertor targets, concentrated by a peaking factor PF on a wetted area A_{wet}, as shown in Equation (1.12).

$$q_{peak} = \frac{PF\,(1-f_{rad})\,P_{heat}}{A_{wet}}, \qquad W = \tau_E P_{heat} \propto P_{heat}^{0.39} \quad \text{for } \tau_E \propto P_{heat}^{-0.61} \tag{1.12}$$

where f_{rad} is the radiated fraction of the heating P_{heat}, W the stored energy and τ_E the energy confinement time. Let P-1 also carry 5 MW of heating onto $1.5\,\mathrm{m}^2$ with $PF = 3$. Unseeded, at

$f_{rad} = 0.3$, its targets see 7.0 MW/m^2, and a 5 MW/m^2 rating needs $f_{rad} \geq 0.5$, rising to 0.69 at 8 MW of heating.

> ✦ DID YOU KNOW? **Ten Megawatts per Square Meter**
>
> The water-cooled divertor of Wendelstein 7-X has a limit of 10 MW/m^2 in steady operation. Its first divertor results projected a peak of only 6.7 MW/m^2 for 10 MW reaching it if the wetted area stays around 1.5 m^2 [27], the area of P-1.

An impurity such as neon, seeded into the edge, radiates there and cools the divertor plasma until it detaches from the targets [28, 29]. The same neon leaking into the core radiates power the plasma needs, and when the core radiates more than it absorbs the plasma ends in a radiative collapse. The confinement scaling of stellarators makes performance expensive: with $\tau_E \propto P_{heat}^{-0.61}$ [30], 20 % more stored energy takes 1.60 times the heating, and all of the extra heating leaves as exhaust. Chapter 13 measures that cost, about 60 % more heating for 20 % more stored energy, on a stable plant whose neon inventory answers in 46.5 s against an energy confinement time of 130.3 ms. There a held neon flow collapses the plasma 18.0 s after the valve opens under plain PI control, while the core-radiation override of protected PI keeps the core under its limit and completes the pulse.

1.4.7 How Physics, Limits and Layers Fit Together

Each chapter's plant can be read in three parts. The physics either helps, as negative feedback that follows a load, or hurts, as a zero, a delay, an instability or a coupling. The limits are set by protection and by equipment. The control layers sit between them: a regulating loop that does what the physics leaves undone, a limiting layer of feedforward, overrides, runbacks or allocation that keeps the regulating loop off the limits, and, in most chapters, a predictive option that plans against the limits directly. Table 1.1 lists the three parts for each chapter. Chapter 14 adds the measured results and compares how physics, limits and control layers fit together in each system.

1.5 Instruments and Final Elements in This Volume

Figure 1.7 lists what the controllers of each chapter read and what they move. Seven chapters read a temperature, five a pressure, five a level and four the neutron flux. Six read a load, a demand or an output ahead of its effect, the signal that a feedforward needs. A valve on steam or water is the commonest handle, in six chapters, and rods or a module's power demand come next, in five. The two plasmas share only a temperature, a heating power and a gas injection with the other chapters.

Neutron flux is measured by detectors that range from fission and ion chambers to self-powered detectors. Detectors outside the vessel give the flux as a representation of reactor power, and detectors in the core give local values. Three ranges usually span operation: proportional counters at start-up, compensated ion chambers in the intermediate range and uncompensated ion chambers at power, where neutrons so outnumber gamma rays that no compensation is needed [31]. Temperatures come mostly from thermocouples and resistance thermometers, often in thermowells [31] (Figure 1.8), and a thermowell adds a lag that, if severe enough, may compromise the stability of feedback control [32]. That lag is the 4 s detector of Section 1.4.5.

Level is measured most often by a differential pressure in nuclear plants [31] (Figure 1.9). The pressure at the bottom of a vessel gives the height only if the density of the liquid is known and constant [32], which is why Chapter 4 and Chapter 6 convert their readings with ranging densities. The impulse lines can clog, trap gas or foul, which may cause a considerable inaccuracy or a poor dynamic response [31].

Control rod drives are mounted on top of the vessel head and position the rods on automatic or manual signals. Two independent systems indicate where a rod is, one from a detector on the drive housing and one from the pulses sent to the drive, and after a dropped or stuck rod the two disagree [33].

Table 1.1 For Each Chapter, the Physics That Helps or Hurts, the Limit That Binds and the Control Layers Between Them

Ch.	Physics	Limit	Control layers
2	Negative feedback follows the load, xenon over hours	Rod speed, deadband, insertion limit	Proportional rod loop on a temperature program, feedforward or flux cascade
3	One steam bubble, spray 9.35 times stronger than heaters	Relief and high pressure	Split range with deadband, level program, high-pressure override, MPC
4	Shrink and swell, a zero that moves with load	Level trips	One-, two- and three-element control, MPC forecasting swell
5	Reactor keeps making heat, a 14.8 s clock	Safety valves, condenser, level	Dump loop, impulse-pressure feedforward, two overrides, runback
6	Voids are moderator, feedback and level error	Level trips, power-flow map	Pressure on turbine valves, three-element feed, pumps for power, scheduling, MPC
7	Pump sets the delayed neutron fraction, transport delay	Delayed neutron margin	Pump loop with load feedforward and flow cascade, MPC with a margin limit
8	Inherent load following, a published mission	Fuel temperature, rod rate	Rod and pump loops on a temperature program, preview feedforward, DMC
9	Header holds 7.55 s of steam, modules 52.5 s	Coolant temperature ceiling that moves	Margin allocation, bypass as fast actuator, runback, MPC
10	Decay heat as a power law, natural circulation	Outlet temperature, pool level, coping time	Valve loop with a margin override, auxiliary train
11	Cold cap, joule heating unstable on voltage, four time scales	Pool floor, off-gas, coverage	Power regulation, feed overrides on a low selector, feedforward, MPC
12	Unstable mode at 356 s^{-1}	Supply voltage, latency, wall	Position and velocity loop, priority allocation, explicit MPC
13	Neon protects the targets and poisons the core	Heat flux, surface temperature, core radiation	Three PI loops, heat-flux cascade, low selectors, MPC

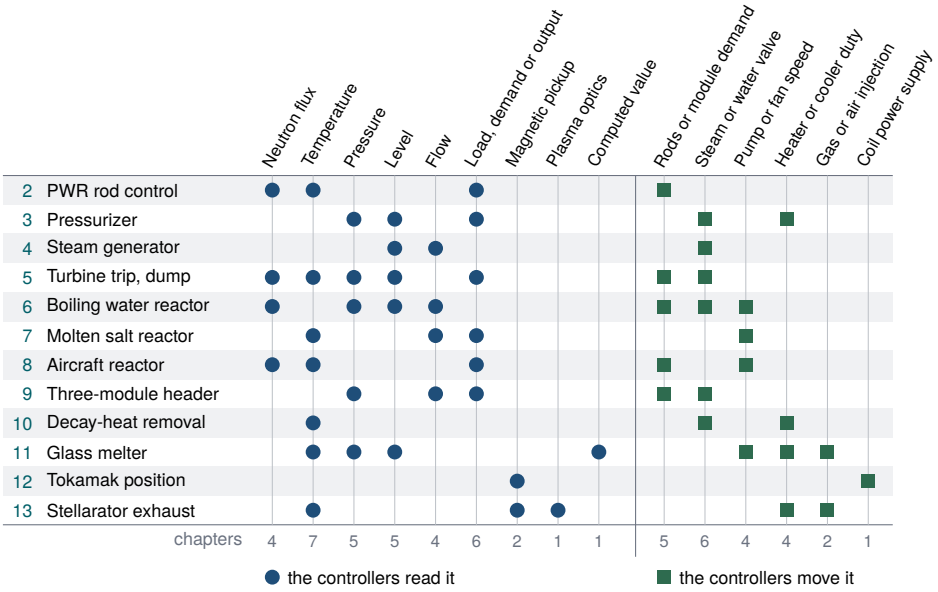


Figure 1.7 Sensors and final elements by chapter: a round mark where the chapter's controllers read the device and a square mark where they move it, as its variables table, control structure or technology section names them. A flow set through a flow loop counts as a valve, and a computed value is not a sensor. Notice that the two fusion rows share almost nothing with the rows above them.

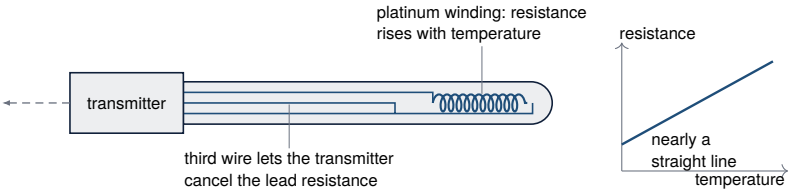


Figure 1.8 Principle of a resistance thermometer. Notice that the winding reads its own temperature, so the well around it sets how late the coolant temperature arrives [32], 3 s in Chapter 2 and 4 s in Chapter 5.

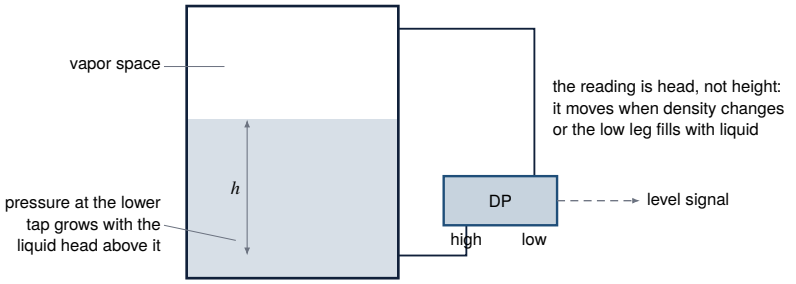


Figure 1.9 Principle of a level measured by differential pressure. Notice that the transmitter reads a head, so a lighter boiling mixture reads lower than cold water of the same height.

Layer	What it does in this volume	Where it runs	How often
Supervisory and predictive control	plans against the limits, follows a flight plan, shares a load among modules	**Supervisory software** that writes setpoints to the loops below	0.2 s to 1 min in the apps
Regulatory control	PI loops on temperature, pressure, level, flow and power; split range and cascade	**Distributed control system** or programmable controller	milliseconds to minutes; 0.05 s to 3 s in the apps
Plasma control	holds a plasma's position, heating, density and radiation	**Plasma control system** on real-time computers	10 kHz on one research tokamak; 100 µs and 20 ms in the apps
Limiting control	runbacks, overrides on selectors, insertion limits, interlocks	function blocks in the controller of the loop they limit	with the loop it limits
Protection	reactor trip, safety valves, passive heat removal on demand	**Protection system** apart from control, with priority over it	when a threshold is crossed
Field devices	neutron detectors, thermometers, transmitters, magnetic loops, rods, valves, pumps, supplies	on wires or a bus to the controller	continuously

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. Loop periods are from [34], the 10 kHz from the research tokamak that Chapter 12 cites and the app periods from the chapters' technology sections. Notice that the protection layer acts on its own, whatever the layers above it command.

Figure 1.10 places the volume's loops on their platforms. Supervisory and lower loops run with cycle times from milliseconds to minutes [34]. Between safety and non-safety equipment a priority logic ensures that safety actuations always override non-safety demands [31]. The apps run their fastest loop every 100 µs on the plasma of Chapter 12 and their slowest every 3 s on the tripped core of Chapter 10.

Table 1.2 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

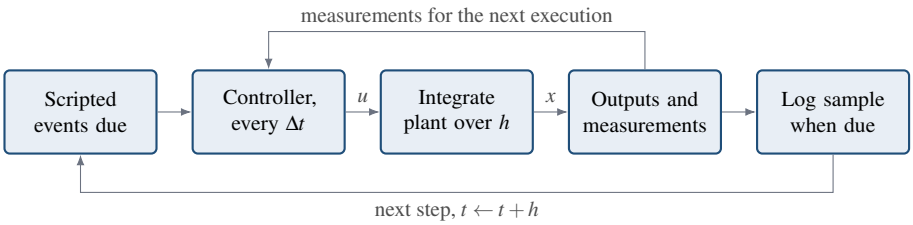


Figure 1.11 Work of the Engine in One Integration Step

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.2 lists what a reader meets.

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{dx}{dt} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings

are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.13) to (1.16) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.13}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.14}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.15}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.16}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.16) drives I toward the value that reproduces the output applied, which is back-calculation [35]. With no integral action, Equation (1.16) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.17).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.17}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.18) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.18}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [36]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [37]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [38].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according

to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed (the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 14.

Table 1.3 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Pressurized water reactor	The plant follows the load, and the rods remove the temperature offset.	introductory	3 h
3	Pressurizer	One output drives two actuators of very different strength.	intermediate	3 h
4	Steam generator	The wrong-way level response moves with load.	intermediate	3.5 h
5	Steam dump	Fifteen seconds in which structure beats gain.	intermediate	3 h
6	Boiling water reactor	Every actuator moves the voids, and the plant controls itself.	advanced	3 h
7	Molten salt reactor	The pump that holds the temperature also sets the margin.	intermediate	3 h
8	Aircraft molten salt reactor	A published flight profile is worth more than retuning.	advanced	2.5 h
9	Three modules, one header	A pressure on setpoint is not a plant under control.	advanced	3 h
10	Decay-heat removal	A hard limit and a water supply that does not come back.	introductory	3 h
11	Glass melter	Loops on four time scales share one feed.	intermediate	3 h
12	Tokamak position	An unstable plant has a floor, a ceiling and a region.	introductory	2.5 h
13	Stellarator exhaust	Limits, not stability, bound a stable plasma.	advanced	2.5 h

1.8 Chapter Roadmap

Table 1.3 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 2, 10 and 12 form a first course: a plant that regulates itself, a plant with no setpoint and a hard limit, and a plant that is unstable. They rest on Sections 1.4.2, 1.4.5 and 1.4.6. Chapters 3, 4, 5, 7 and 11 add split range, a moving zero, a disturbance faster than its measurement, an actuator that moves the margin and loops on four time scales. Chapters 6, 8, 9 and 13 combine coupling, preview or allocation with constraints and are best left until Section 1.4.7 is familiar. A reader who follows the steam can read Chapters 2 to 5 in order, then Chapter 9, before turning to the other reactors and the plasmas.

KEY POINTS

- A nuclear plant is controlled around physics that acts faster than any controller and around limits that no controller may cross.
- A settled reactor is a statement about the reactivity ledger, and removing the temperature offset that feedback leaves is the controller's whole task.
- A level that counts water and steam together first moves the wrong way, and its right-half-plane zero bounds the crossover of any loop on it.
- Only a signal from upstream of the effect acts inside the clock to a protection limit.
- On an unstable plasma, gain decides how the loop behaves inside the recoverable region, and the growth rate and supply decide its size.

Where else this applies. Overrides and selectors are the subject of Volume 4, and unstable plants of Volume 7.

References

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

[2] L. E. Boing, E. A. Wimunc, and G. A. Whittington, "Design-development and operation of the experimental boiling-water reactor (EBWR) facility, 1955–1967," Argonne National Laboratory, Argonne, IL, Tech. Rep. ANL-91/13, Nov. 1990.

[3] Oak Ridge National Laboratory, “The development status of molten-salt breeder reactors,” Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL-4812, Aug. 1972.

[4] President’s Commission on the Accident at Three Mile Island, “Report of the President’s commission on the accident at Three Mile Island. The need for change: The legacy of TMI,” President’s Commission on the Accident at Three Mile Island, Washington, DC, Tech. Rep. NP-24200, Oct. 1979.

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

[6] H. Takahashi, K. Bol, H. Maeda, M. Okabayashi, and M. Reusch, “Experimental determination of the vertical instability strength in PDX tokamak,” Princeton Plasma Physics Laboratory, Princeton, NJ, Tech. Rep. PPPL-1911, Jul. 1982.

[7] J. R. Larson, “Comparison of TRAC-BF1 calculations with the LaSalle 2 instability event,” Idaho National Engineering Laboratory, Idaho Falls, ID, Tech. Rep. EGG-NRE-10803, May 1993.

[8] B. C. Hodges, “Replacement of melter 2 at the defense waste processing facility (DWPF),” in *WM2018 Conference*, Phoenix, AZ, 2018, paper 18608; report SRR-STI-2018-00060, OSTI 1431108. [Online]. Available: https://www.osti.gov/biblio/1431108

[9] International Atomic Energy Agency, *Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2018, no. NP-T-3.23.

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

[11] T. Klinger, A. Alonso, S. Bozhenkov, R. Burhenn, A. Dinklage, G. Fuchert *et al.*, “Performance and properties of the first plasmas of Wendelstein 7-X,” *Plasma Physics and Controlled Fusion*, vol. 59, no. 1, p. 014018, 2017.

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

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

[14] J. D. Vienna, “Nuclear waste vitrification in the United States: Recent developments and future options,” *International Journal of Applied Glass Science*, vol. 1, no. 3, pp. 309–321, 2010.

[15] E. A. Lazarus, J. B. Lister, and G. H. Neilson, “Control of the vertical instability in tokamaks,” *Nuclear Fusion*, vol. 30, no. 1, pp. 111–141, 1990.

[16] G. R. Keepin, T. F. Wimett, and R. K. Zeigler, “Delayed neutrons from fissionable isotopes of uranium, plutonium, and thorium,” *Physical Review*, vol. 107, no. 4, pp. 1044–1049, 1957.

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

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

[19] J. March-Leuba and J. M. Rey, “Coupled thermohydraulic-neutronic instabilities in boiling water nuclear reactors: A review of the state of the art,” *Nuclear Engineering and Design*, vol. 145, no. 1–2, pp. 97–111, 1993.

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

[21] R. H. Middleton, “Trade-offs in linear control system design,” *Automatica*, vol. 27, no. 2, pp. 281–292, 1991.

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

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

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

[25] M. L. Walker and D. A. Humphreys, “On feedback stabilization of the tokamak plasma vertical instability,” *Automatica*, vol. 45, no. 3, pp. 665–674, 2009.

[26] D. A. Humphreys, T. A. Casper, N. Eidietis, M. Ferrara, D. A. Gates, I. H. Hutchinson, G. L. Jackson, E. Kolemen, J. A. Leuer, J. Lister, L. L. LoDestro, W. H. Meyer, L. D. Pearlstein, A. Portone, F. Sartori, M. L. Walker, A. S. Welander, and S. M. Wolfe, “Experimental vertical stability studies for ITER performance and design guidance,” *Nuclear Fusion*, vol. 49, no. 11, p. 115003, 2009.

[27] T. S. Pedersen, R. König, M. Krychowiak, M. Jakubowski, J. Baldzuhn, S. Bozhenkov *et al.*, "First results from divertor operation in Wendelstein 7-X," *Plasma Physics and Controlled Fusion*, vol. 61, no. 1, p. 014035, 2019.

[28] A. Kallenbach, M. Bernert, R. Dux, L. Casali, T. Eich, L. Giannone, A. Herrmann, R. McDermott, A. Mlynek, H. W. Müller, F. Reimold, J. Schweinzer, M. Sertoli, G. Tardini, W. Treutterer, E. Viezzer, R. Wenninger, M. Wischmeier *et al.*, "Impurity seeding for tokamak power exhaust: From present devices via ITER to DEMO," *Plasma Physics and Controlled Fusion*, vol. 55, no. 12, p. 124041, 2013.

[29] P. C. Stangeby, "Basic physical processes and reduced models for plasma detachment," *Plasma Physics and Controlled Fusion*, vol. 60, no. 4, p. 044022, 2018.

[30] H. Yamada, J. H. Harris, A. Dinklage, E. Ascasibar, F. Sano, S. Okamura, J. Talmadge, U. Stroth, A. Kus, S. Murakami, M. Yokoyama, C. D. Beidler, V. Tribaldos, K. Y. Watanabe, and Y. Suzuki, "Characterization of energy confinement in net-current free plasmas using the extended international stellarator database," *Nuclear Fusion*, vol. 45, no. 12, pp. 1684–1693, 2005.

[31] 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.

[32] 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.

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

[34] 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.

[35] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

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

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

[38] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Theory and calculation · 4 min

A reactor has one delayed neutron group with $\beta = 650\,\text{pcm}$ and $\lambda = 0.08\,\text{s}^{-1}$, and a prompt generation time of 2×10^{-5} s. At steady power a rod withdrawal inserts a step of 130 pcm, and no feedback acts.

(a) Find the prompt jump and the stable period. *Checkpoint:* a jump to 1.25 times the power, a period of 50 s.

(b) Find the time constant of the prompt mode that the prompt jump approximation drops, and the power a minute after the step. *Checkpoint:* 3.85 ms, 4.15 times the initial power.

1.2 Theory and calculation · 5 min

A pressurized water reactor has a Doppler coefficient of −2.5 pcm/K on the fuel temperature and a moderator coefficient of −17.5 pcm/K on the coolant temperature. At full power its fuel runs 400 K above the coolant, and the fuel-to-coolant difference is proportional to power. The turbine load falls from 100 to 80 % with the rods parked.

(a) Find the reactivity the Doppler term returns and the steady rise of the coolant temperature. *Checkpoint:* 200 pcm, a rise of 10.0 K.

(b) The bank is worth 12 pcm per percent of travel. Find the steady gain from bank position to coolant temperature and the insertion that removes the rise. *Checkpoint:* 0.60 K per percent, 16.7 % inserted.

1.3 Concept check · 2 min

The rod drive of a pressurized water reactor takes a speed command, and the average-temperature controller in front of it is proportional only. After a load cut the temperature settles exactly on its program. Why is there no steady offset?

(a) the negative temperature coefficients remove the offset by themselves
(b) the drive integrates the speed, so the bank moves until the error is zero
(c) the temperature program moves the setpoint to the measured temperature
(d) the reactor power settles at the load, which removes the error

1.4 Theory and calculation · 3 min

A molten salt reactor's fuel produces a delayed fraction of 650.2 pcm. At its design pump speed the effective delayed fraction is 419.4 pcm, and at 40 % pump speed it is 493.1 pcm.

(a) Express a 70 pcm rod insertion in dollars at each pump speed, and in a solid core of the same fuel. *Checkpoint:* 0.167, 0.142 and 0.108 dollars.

(b) Find the excess reactivity the rods must hold at each pump speed. *Checkpoint:* 230.8 pcm and 157.1 pcm.

1.5 Theory and calculation · 4 min

A reactor burns xenon three times faster than xenon decays at full power, $r_X = 3\lambda_X$, and its full-power equilibrium xenon is worth 2700 pcm. Ignore direct xenon production in fission.

(a) Find the equilibrium xenon, relative to full power, at 60 % and at 25 % power, and the reactivity each frees. *Checkpoint:* 0.857 and 0.571, freeing 386 pcm and 1157 pcm.

(b) The power is cut from 100 to 60 %. With the xenon half life of 9.21 h, find the rate at which the normalized xenon changes just after the cut, and its sign. *Checkpoint:* a rise of 0.090 per hour.

1.6 Theory and calculation · 5 min

A drum has a surface of 8 m^2 over water of 750 kg/m^3. Each extra kilogram per second of feedwater collapses enough steam to lower the level by 3 mm once the bubbles settle, with a time constant of 12 s.

(a) Find the water-term gain K_w and decide whether the level has a right-half-plane zero. If it has, locate it. *Checkpoint:* $K_w = 0.167$ mm per kilogram, zero at $0.167\ s^{-1}$.

(b) For a feed step of 5 kg/s, find when the level reaches its minimum and how deep it goes, and the crossover frequency a level loop should stay below. *Checkpoint:* 4.87 s, –0.945 mm, about 0.083 rad/s.

1.7 Concept check · 2 min

Three-element drum level control adds a feedforward from steam flow to the feedwater flow setpoint, with a cascade on the feed flow. What does the feedforward remove?

(a) both the shrink and the swell that the level shows
(b) the right-half-plane zero of the level response to feedwater
(c) the lag of the level transmitter
(d) the mass imbalance between steam and feed, leaving the void response

1.8 Theory and calculation · 4 min

A 2000 MW reactor trips after a year at full power. Use the decay-heat approximation of this chapter.

(a) Find the decay heat one hour and one day after the trip. *Checkpoint:* 21.5 MW and 9.42 MW.

(b) Find the decay energy released between one hour and one day, and the mass of saturated water it boils at 2257 kJ/kg. *Checkpoint:* 280 MWh, about 446 t.

1.9 Theory and calculation · 4 min

An elongated plasma's vertical mode grows at $\gamma = 600\,\mathrm{s}^{-1}$, and its coil moves it at $b = 300\,\mathrm{mm/s}$ per volt from a supply limited to $\pm 12\,\mathrm{V}$. Neglect latency.

(a) Find the smallest proportional gain that holds the plasma and the largest displacement the supply can recover. *Checkpoint:* 2.0 V/mm, 6.0 mm.

(b) A higher elongation doubles the growth rate. Find the new floor and region, and the region at the original growth rate with a 15 V supply. *Checkpoint:* 4.0 V/mm and 3.0 mm, then 7.5 mm.

1.10 Concept check · 2 min

The confinement time of a stellarator plasma scales as the heating power to the -0.61. To raise the stored energy by 20 % at fixed density, about how much more heating is needed?

(a) about 60 % more, all of it leaving as exhaust

(b) 20 % more, because stored energy follows heating

(c) about 12 % more, because confinement improves with power

(d) 39 % more, the exponent of the stored energy