INTERACTIVE CONTROL ENGINEERING

Control of Nuclear and Fusion Systems

Interactive Control Engineering: Principles and Applications

Edited by John D. Hedengren

Control technology

1 Process Dynamics, Modeling, and Identification
2 Practical PID Control
3 Feedforward, Cascade, and Ratio Control
4 Override, Split-Range, and Constraint Control
5 Delays, Slow Measurements, and State Estimation
6 Multivariable and Plantwide Control
7 Nonlinear Dynamics, Oscillations, and Stabilization
8 Batch, Run-to-Run, and Hybrid Control
9 Model Predictive Control and Economic Optimization

Industry applications

10 Control in Upstream and Midstream Oil and Gas
11 Control in Petroleum Refining
12 Control in Petrochemicals and Polymers
13 Control in Chemical Plants
14 Control in Semiconductor Manufacturing
15 Control in Mining and Mineral Processing
16 Control in Iron, Steel, and Metals Production
17 Control of Energy Conversion and Storage
18 Control of Nuclear and Fusion Systems
19 Control of Biological Systems and Human Performance
20 Motion Control, Robotics, and Autonomous Systems

INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18

Control of Nuclear and Fusion Systems

Reactor physics, protection limits and coordinated control

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 18: Control of Nuclear and Fusion Systems
Edited by John D. Hedengren

First edition, 2026

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

Produced by Advanced Process Solutions, LLC.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in a review or scholarly work.

The simulators described in this book are teaching models. Their equations, parameters and results are simplified for instruction and are not validated for the design, operation or safety assessment of any real plant, vehicle or medical device. The editor and the publisher accept no liability for use of this material outside that purpose.

Names of products and companies are used for identification only and may be trademarks of their owners.

Every chapter has a companion simulator that runs in a web browser at https://apmonitor.com/apps/. This volume online, with chapter files and citation records: https://apmonitor.com/apps/books/v18-nuclear-fusion/

How to cite this volume. John D. Hedengren, Ed., *Control of Nuclear and Fusion Systems*, vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026.

Typeset in Times and Helvetica-family type with LaTeX and the memoir class.

Preface

This book is one of 20 volumes of *Interactive Control Engineering: Principles and Applications*. Each chapter is about one plant, machine or system and one control lesson, and each has a simulator that runs in a web browser with nothing to install. The simulators are collected at https://apmonitor.com/apps/. Nine volumes are organized by control technology and eleven by industry or application area; the last pages of this volume list them all.

How to read this book with the apps

Chapter 1 teaches the theory this volume relies on and explains, once, the simulator framework that the apps share. Each chapter after it opens with the control question it answers, its learning objectives, what to read first, and the address of its app, also given as a code to scan. The chapter then develops the model, the controllers and the measured results. Every number reported was measured on the app, so a reader can open the app, load the named scenario and reproduce it.

Every chapter ends with exercises. Each is labeled with its kind and the time it should take, and none takes more than ten minutes. Concept checks and calculations can be done on paper. App challenges name a scenario, the settings and the run length. One exercise in each chapter, labeled *Review the agent*, presents a tuning, a model or a claim proposed by an AI assistant and asks what is wrong with it. Short answers to the odd-numbered exercises are at the back of the book, and worked solutions are available to instructors.

How AI agents were used

AI agents did much of the development work behind this series. Under the editor's direction they wrote and tested the simulators, drafted the papers from which the chapters are drawn, measured the reported results on the running apps, converted the papers to book form, and drafted the introductory and closing chapters and the exercises. Other agents then checked that work: they re-measured reported numbers, audited what each app displays against the state of its model, reviewed the papers as referees, and solved the exercises without seeing the answers. The editor set the scope, chose the plants and the lessons, reviewed the results and is responsible for the content.

That way of working is also a subject of the book. The last chapter of each volume describes what the agents got wrong while this volume was made and how the errors were caught, and it sets out what the editor believes this means for engineers: more of the routine work can be delegated, and more of the engineer's value lies in checking, supervising, choosing the problem and knowing the process.

About this volume

This volume is about nuclear and fusion systems, which are controlled inside strict limits by physics that both helps and hinders. Chapters 2 to 6 cover light water reactors and their steam systems. Chapters 7 to 9 cover molten salt reactors and a plant of several small modules. Chapters 10 and 11 cover decay heat removal and waste vitrification. Chapters 12 and 13 cover the plasma control problems of a tokamak and a stellarator.

The simulators are teaching models and are not validated for the analysis of any real plant. Each chapter names the technology volume that teaches its main technique.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Controlling Reactors and Plasmas 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 3
- 1.3 Notation and Conventions . . . 4
- 1.4 Physics, Limits and Control Layers . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 12
- 1.6 The Simulator Framework . . . 15
- 1.7 Working With AI Agents While Studying . . . 18
- 1.8 Chapter Roadmap . . . 19
- 1.9 Exercises . . . 21

2 Pressurized Water Reactor Rod Control Simulator for Inherent Feedback and Velocity Actuators 24
- 2.1 Introduction . . . 25
- 2.2 Background . . . 26
- 2.3 Reactor Model . . . 28
- 2.4 Dynamic Analysis and Controller Design . . . 33
- 2.5 Results . . . 37
- 2.6 Technology in Practice . . . 43
- 2.7 Limitations . . . 46
- 2.8 Discussion . . . 47
- 2.9 Conclusion . . . 47
- 2.10 Exercises . . . 49

3 Pressurizer Simulator for Split-Range Control Across Asymmetric Actuators 53
- 3.1 Introduction . . . 54
- 3.2 Primary Pressure Control and Its Instruments . . . 55
- 3.3 Process Model . . . 56
- 3.4 The Split Range and Its Two Plants . . . 62
- 3.5 Controller Implementation . . . 64
- 3.6 Two Loops on One Vessel . . . 66
- 3.7 Results . . . 67
- 3.8 Technology in Practice . . . 73
- 3.9 Limitations . . . 74
- 3.10 Discussion . . . 75
- 3.11 Conclusion . . . 76
- 3.12 Exercises . . . 78

4 Three-Element Drum Level Simulator for an Inverse Response That Moves With Load 81
- 4.1 Introduction . . . 82
- 4.2 Background . . . 84
- 4.3 Steam Generator Model . . . 85
- 4.4 Inverse Response and Its Dependence on Load . . . 90
- 4.5 Controller Design . . . 91

4.6 Results . . . 95
4.7 Discussion . . . 103
4.8 Technology in Practice . . . 105
4.9 Limitations . . . 106
4.10 Conclusion . . . 107
4.11 Exercises . . . 109

5 Nuclear Steam Dump Simulator for Feedforward and Override on a Load Rejection 112
5.1 Introduction . . . 113
5.2 Turbine Bypass and Steam Dump Control . . . 114
5.3 Plant Model . . . 116
5.4 The Clock . . . 123
5.5 Controller Design . . . 124
5.6 Results . . . 127
5.7 Technology in Practice . . . 133
5.8 Limitations . . . 136
5.9 Discussion . . . 136
5.10 Conclusion . . . 137
5.11 Exercises . . . 139

6 Boiling Water Reactor Simulator for Coupled Loops With Sign-Reversing Dynamics 142
6.1 Introduction . . . 143
6.2 Control of Boiling Water Reactors . . . 144
6.3 Reactor Model . . . 145
6.4 Coupling, Sign Reversal and the Operating Map . . . 152
6.5 Controller Design . . . 156
6.6 Results . . . 159
6.7 Technology in Practice . . . 165
6.8 Limitations . . . 168
6.9 Discussion . . . 168
6.10 Conclusion . . . 169
6.11 Exercises . . . 171

7 Molten Salt Reactor Simulator Where the Pump Sets the Stability Margin 174
7.1 Introduction . . . 175
7.2 Background . . . 176
7.3 Reactor Model . . . 177
7.4 Selection of the Manipulated Variable . . . 183
7.5 Controller Design . . . 185
7.6 Results . . . 188
7.7 Discussion . . . 196
7.8 Technology in Practice . . . 196
7.9 Limitations . . . 198
7.10 Conclusion . . . 199
7.11 Exercises . . . 201

8 Mission-Profile Load Following and Preview Control of an Aircraft Molten Salt Reactor 204
8.1 Introduction . . . 205
8.2 Background . . . 206
8.3 Aircraft Reactor Model . . . 208
8.4 Control Architectures . . . 212
8.5 Interactive Application . . . 215

8.6 Results . . . 216
8.7 Technology in Practice . . . 221
8.8 Limitations . . . 223
8.9 Conclusion . . . 224
8.10 Exercises . . . 225

9 Multi-Module Steam Header Simulator for Load Allocation and Plantwide Coordination 229
9.1 Introduction . . . 230
9.2 Background . . . 231
9.3 Process Model . . . 232
9.4 Coupling and Capability . . . 239
9.5 Control Architectures . . . 241
9.6 Results . . . 245
9.7 Technology in Practice . . . 252
9.8 Limitations . . . 254
9.9 Discussion . . . 255
9.10 Conclusion . . . 256
9.11 Exercises . . . 258

10 Decay-Heat Removal Simulator for Constraint Control of a Depleting Heat Sink 261
10.1 Introduction . . . 262
10.2 Decay Heat and the Passive Answer to It . . . 263
10.3 Plant Model . . . 265
10.4 Static Behavior of the Loop and the Valve . . . 272
10.5 Four Architectures and Two Controllers . . . 273
10.6 Measured Results . . . 275
10.7 Technology in Practice . . . 282
10.8 Limitations . . . 284
10.9 Discussion . . . 284
10.10 Conclusion . . . 285
10.11 Exercises . . . 286

11 Waste Vitrification Melter Simulator for Control Across Four Time Scales 290
11.1 Introduction . . . 291
11.2 Vitrification and Its Control Problem . . . 292
11.3 Process Model . . . 293
11.4 Dynamic Analysis . . . 300
11.5 Control Structure . . . 302
11.6 Results . . . 305
11.7 Technology in Practice . . . 311
11.8 Limitations . . . 313
11.9 Discussion . . . 314
11.10 Conclusion . . . 315
11.11 Exercises . . . 317

12 Tokamak Vertical Stabilization Simulator for Gain Limits and the Recoverable Region of an Unstable Plant 320
12.1 Introduction . . . 321
12.2 Vertical Control of Elongated Tokamaks . . . 322
12.3 Plasma, Wall and Coil Model . . . 323
12.4 Implemented Controller . . . 325
12.5 Results . . . 328

12.6 Technology in Practice . . . 335
12.7 Limitations . . . 337
12.8 Conclusion . . . 337
12.9 Exercises . . . 339

13 Stellarator Heat Exhaust Simulator in Which One Impurity Both Protects and Poisons 342
13.1 Introduction . . . 343
13.2 Stellarators, Island Divertors and Radiated Power . . . 344
13.3 The Plasma and Exhaust Model . . . 345
13.4 Dynamic Character . . . 350
13.5 The Control System as Implemented . . . 352
13.6 Ten Cases and a Challenge, Measured . . . 353
13.7 Technology in Practice . . . 357
13.8 Limitations . . . 359
13.9 Conclusion . . . 359
13.10 Exercises . . . 361

14 Summary and Outlook 365
14.1 What the Chapters Showed . . . 365
14.2 Choosing a Method . . . 368
14.3 Where Practice Is Heading . . . 369
14.4 Engineering With AI Agents . . . 371
14.5 Gaps in the Control of Nuclear and Fusion Systems . . . 374
14.6 Exercises . . . 378

Answers to Selected Exercises 379

Glossary 384

Index of Control Topics 387

Index of Applications 390

About the Editor 393

Volumes in This Series 394

Notation

The chapters share the symbols below. A symbol used by one chapter only is defined where it first appears, and each chapter states its own units.

Symbol	Meaning
SP	setpoint, the value a controller is asked to hold
PV	process variable, the measurement a controller acts on
CV	controlled variable
MV	manipulated variable; also OP, the controller output
DV	disturbance variable
e	control error, SP $-$ PV
K_c	controller gain
τ_I, τ_D	integral (reset) time and derivative time
K_p	process gain, the steady-state change in PV per unit change in MV
τ_p	process time constant
θ_p	process dead time
τ_c	closed-loop time constant chosen in a tuning rule
Δt	controller sample time
s	Laplace variable
x, u, y, d	state, input, output and disturbance of a state-space model
A, B, C	matrices of the linear model $\dot{x} = Ax + Bu$, $y = Cx$
Q, R	weights on outputs (or states) and on input moves in an optimal controller
$\hat{x}$	an estimate of x

Abbreviations used throughout: PID, proportional-integral-derivative; PI, proportional-integral; FOPDT, first-order plus dead time; MPC, model predictive control; MHE, moving horizon estimation; IAE, integral of the absolute error; RGA, relative gain array. Times are in seconds unless a chapter says otherwise. Controller output and valve position are in percent of span. A scenario or case number refers to the numbered list in the chapter's app.

CHAPTER 1

Introduction: Controlling Reactors and Plasmas

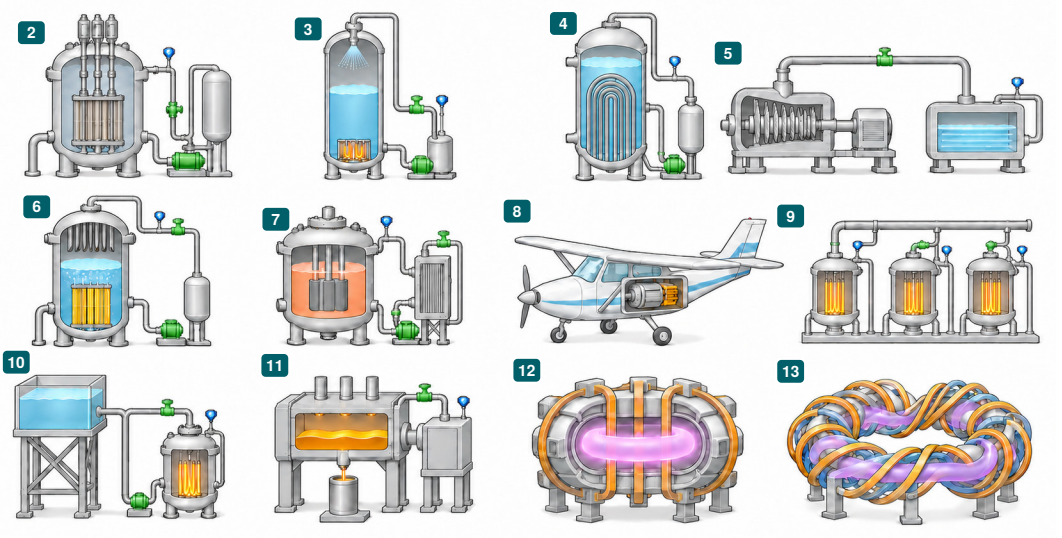

The chapters 2 PWR rod control 3 pressurizer 4 steam generator level 5 turbine trip and dump 6 boiling water reactor 7 molten salt reactor 8 aircraft reactor 9 three-module header 10 decay-heat removal 11 glass melter 12 tokamak position 13 stellarator exhaust

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.

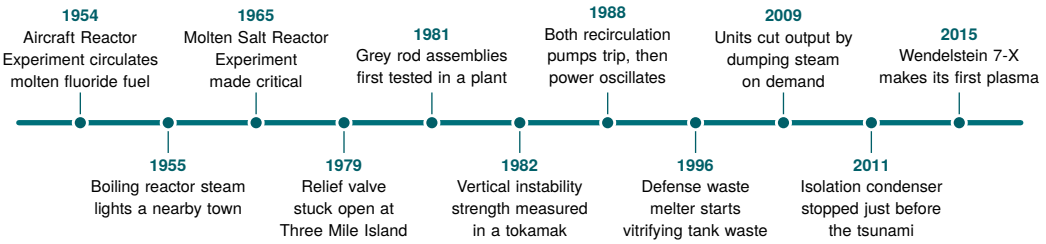

Figure 1.1 Milestones of nuclear and fusion plants and their control. The dates are from [1] (1954), [2] (1955), [3] (1965), [4] (1979), [5] (1981), [6] (1982), [7] (1988), [8] (1996), [9] (2009), [10] (2011) and [11] (2015). Notice that the fission plants span the whole period, while the fusion devices 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,

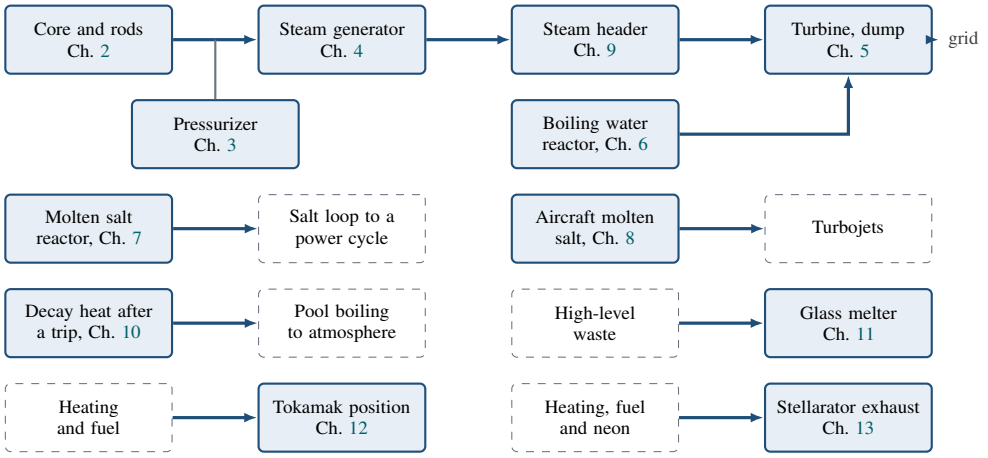

Figure 1.2 The Plants of the Volume. Filled Blocks Are the Plants of Chapters 2 to 13, Dashed Blocks Lie Outside the Volume. Top Two Rows: a Pressurized Water Reactor From Core to Grid, With the Pressurizer on Its Primary Circuit, and a Boiling Water Reactor That Sends Core Steam to the Turbine. Third 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

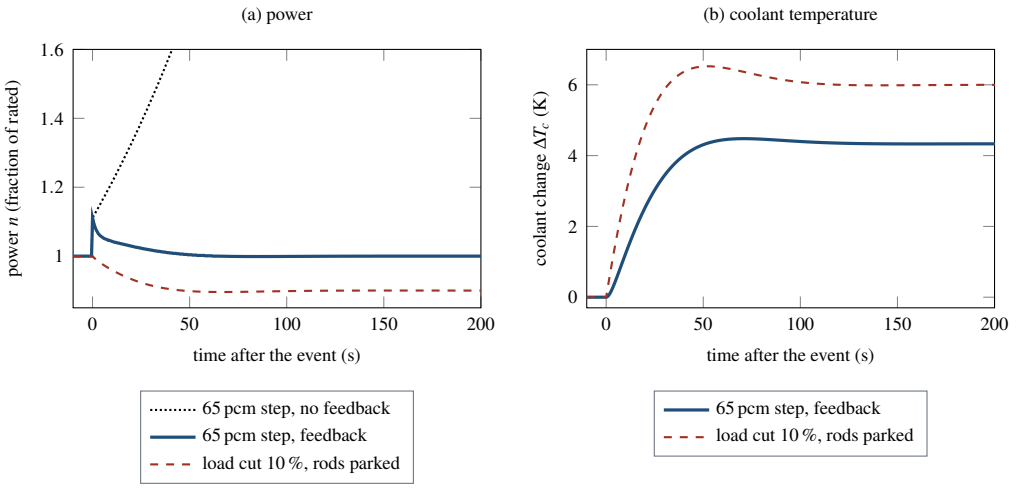

Figure 1.3 R-1 After a Rod Withdrawal of 65 pcm at a Fixed Heat Sink and After a 10 % Load Cut With the Rods Parked. (a) Power Without Feedback (Dotted, Rising on a 112.5 s Period), With Feedback (Solid, Returning to Rated) and After the Load Cut (Dashed, Settling at 90 %). (b) The Coolant Settles 4.33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

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

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 %.

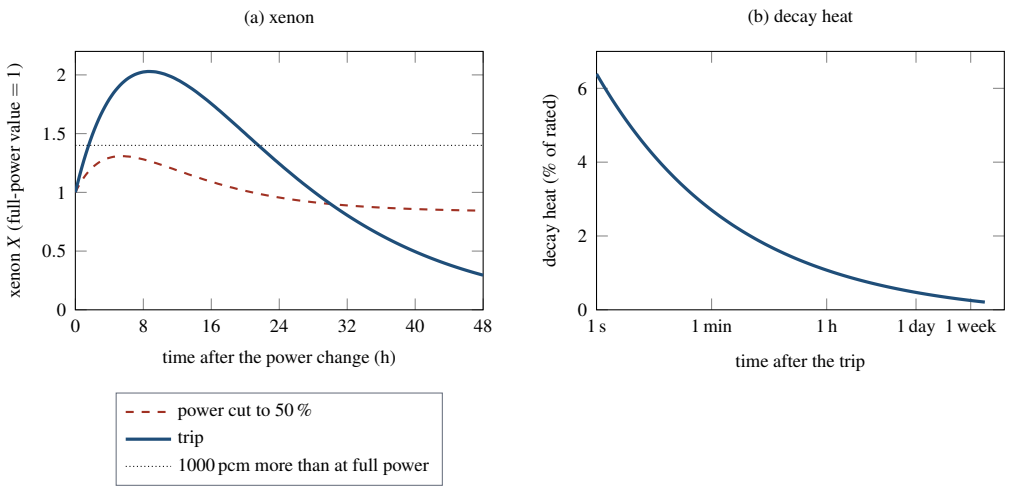

Figure 1.4 The Two Slow Clocks of R-1. (a) Xenon After a Power Cut to 50 % (Dashed) and After a Trip (Solid). Above the Dotted Line the Extra Xenon Exceeds 1000 pcm. (b) Decay Heat After a Year at Full Power, From Equation (1.9).

> **+ 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

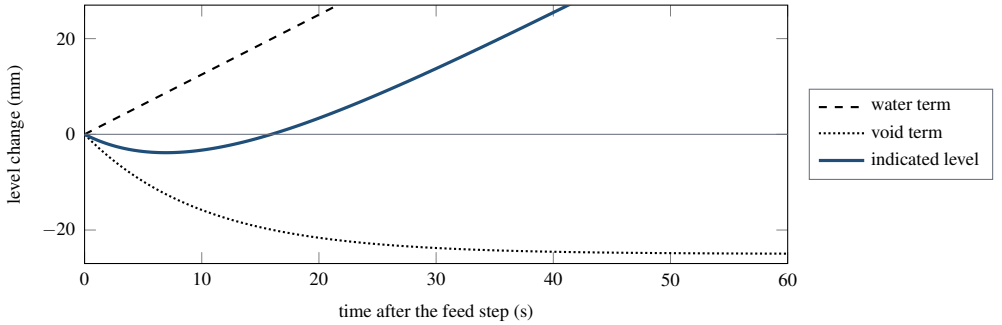

Figure 1.5 Shrink on R-1's Drum After a Feed Step of 10 kg/s. The Water Term (Dashed) Rises at 1.25 mm/s, the Void Term (Dotted) Settles at −25 mm With a Time Constant of 10 s, and the Indicated Level (Solid) Falls to −3.84 mm at 6.93 s Before Rising.

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,

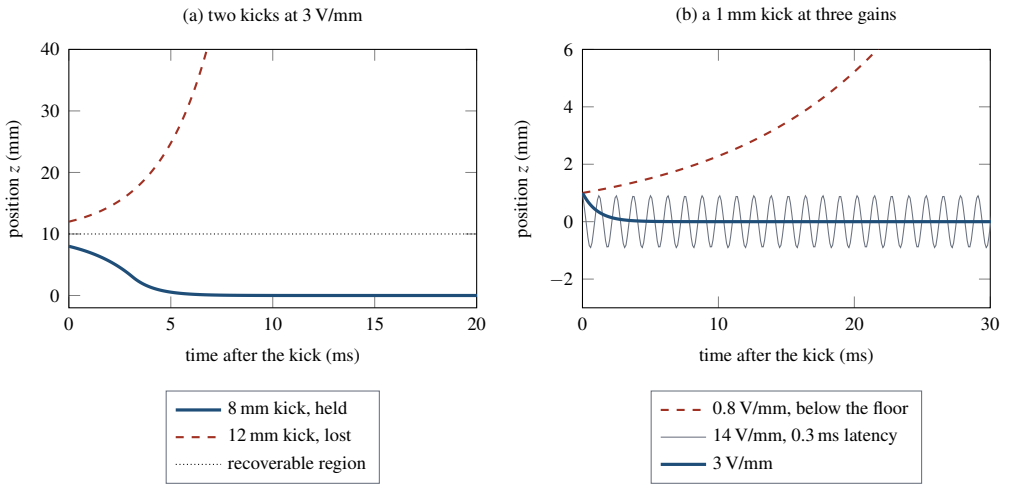

Figure 1.6 Vertical Position of P-1 Under a Proportional Loop With a ±10 V Supply. (a) At 3 V/mm and 0.1 ms Latency an 8 mm Kick Inside the 10 mm Region (Dotted) Is Held (Solid) and a 12 mm Kick Is Lost (Dashed). (b) A 1 mm Kick Drifts to the Wall Below the Floor (Dashed, Leaving the Axis), Is Held at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

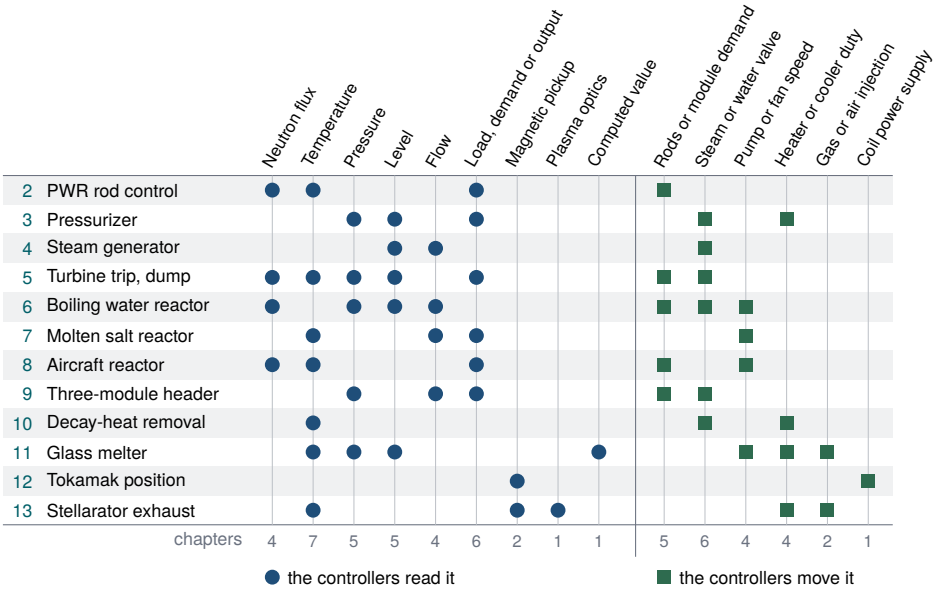

Figure 1.7 Sensors and final elements by chapter: a round mark where the chapter's controllers read the device and a square mark where they move it, as its variables table, control structure or technology section names them. A flow set through a flow loop counts as a valve, and a computed value is not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

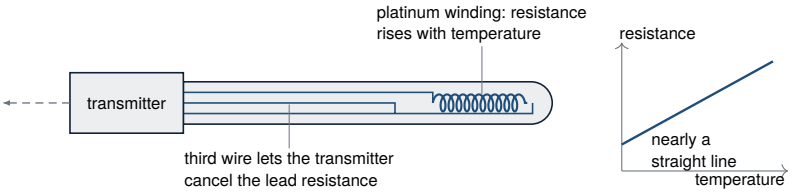

Figure 1.8 Principle of a resistance thermometer. Notice that the winding reads its own temperature, so the well around it sets how late the coolant temperature arrives [32], 3 s in Chapter 2 and 4 s in Chapter 5.

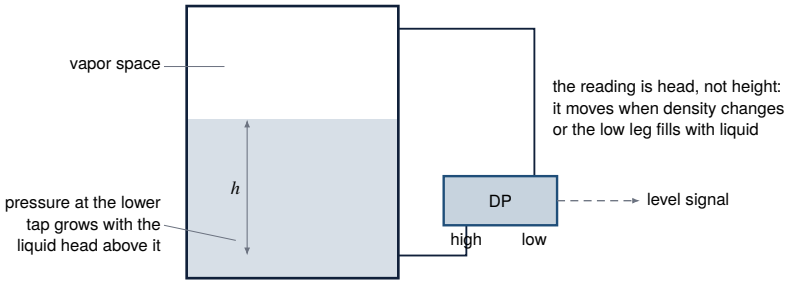

Figure 1.9 Principle of a level measured by differential pressure. Notice that the transmitter reads a head, so a lighter boiling mixture reads lower than cold water of the same height.

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.

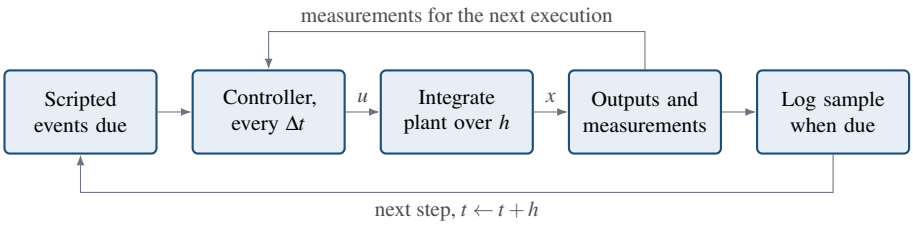

Figure 1.11 Work of the Engine in One Integration Step

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 $\frac{\mathrm{d}x}{\mathrm{d}t} = f(x,u,d,p)$ 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

CHAPTER 2

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

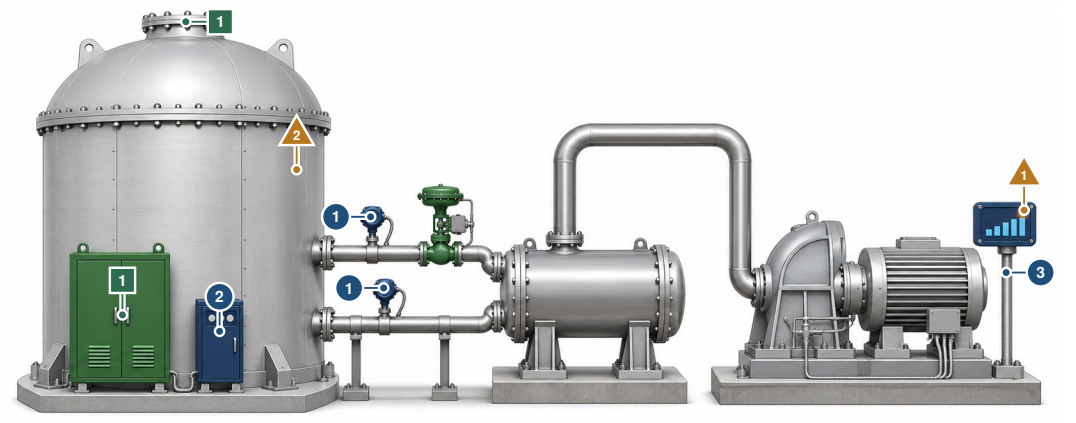

The controller moves 1 rod control cabinet and the rod drives on the vessel head

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

What disturbs it 1 turbine load change 2 a dropped rod

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

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

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

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

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

THREE WAYS TO READ THIS CHAPTER

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

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

Learning objectives

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

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

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

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

2.1 Introduction

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

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

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

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

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

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

2.2 Background

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

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

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

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

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

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

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

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

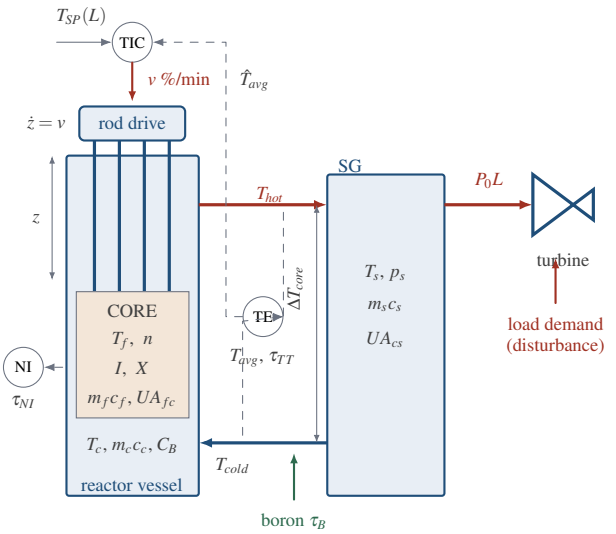

Figure 2.1 Primary circuit of the modeled four-loop plant with control bank D, one lumped steam generator and the turbine. Fission heats the fuel node T_f, which passes heat to the coolant node T_c through conductance UA_{fc}, and the coolant passes it to the secondary node T_s through UA_{cs}. Hot and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

2.3 Reactor Model

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

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

2.3.1 Point Kinetics Under the Prompt Jump Approximation

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

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

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

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

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

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

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

2.3.2 Reactivity Balance and Control Bank Worth

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

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

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

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

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

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

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

2.3.3 Fuel, Coolant and Steam Generator Energy Balances

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

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

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

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

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

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

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

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

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

2.3.4 Iodine, Xenon and Boron

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

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

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

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

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

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

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

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

2.3.5 Temperature Program, Rod Drive and Instruments

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

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

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

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

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

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

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

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

2.3.6 Parameters, Initialization and Numerics

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

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

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

2.3.7 Assumptions

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

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

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

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

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

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

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

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

2.4 Dynamic Analysis and Controller Design

2.4.1 Process Gain From Bank Position to Temperature

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

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

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

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

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

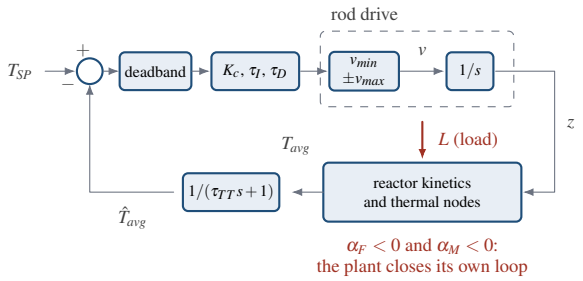

Figure 2.2 The rod control loop as implemented. The manipulated variable is rod speed, so the drive contains the integrator, and the controller ships proportional-only. The deadband acts on the error rather than on the output, which makes it a gap controller and keeps it free of windup, while the minimum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

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

2.4.2 Controller Implementation

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Gains in %/(min °C).

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

2.4.3 Frequency Response and the Case for Proportional-Only Control

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

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

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

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

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

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

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

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

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

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

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

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

2.5 Results

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

2.5.1 Inherent Feedback With the Bank Parked

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

Table 2.5 Scenarios and Completion Requirements

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

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

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

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

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

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

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

Pause. Try Exercise 2.2 before reading on.

2.5.2 Proportional Gain and the Measurement Lag

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

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

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

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

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

2.5.3 Deadband and Minimum Rod Speed

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

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

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

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

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

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

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

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

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

2.5.4 Feedforward and Cascade Against Two Disturbances

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

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

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

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

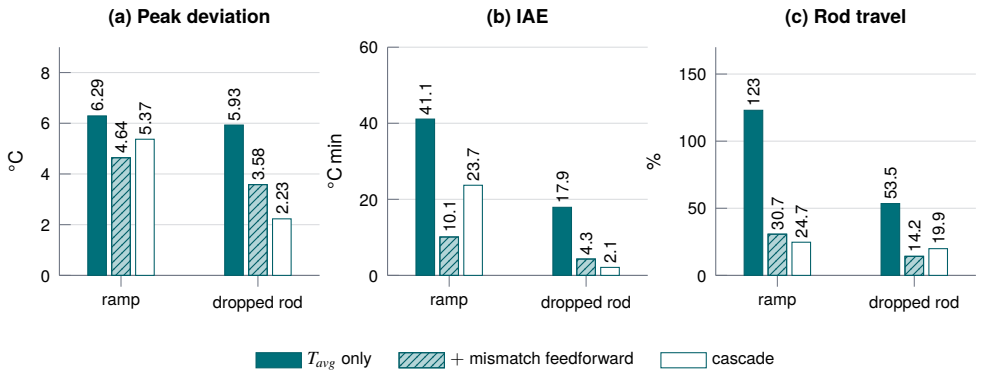

Figure 2.3 The three architectures of Table 2.8 against the 30 % turbine ramp at 20 %/min and the −150 pcm dropped rod: (a) peak deviation, (b) integrated absolute error and (c) rod travel. Solid: T_{avg} only. Hatched: with the power-mismatch feedforward. White: cascade. Notice that the feedforward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

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

2.5.5 Xenon Transient and the Rod Insertion Limit

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

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

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

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

2.5.6 Sensitivity to the Moderator Temperature Coefficient

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

2.5.7 Graded Challenge

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

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

2.6 Technology in Practice

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

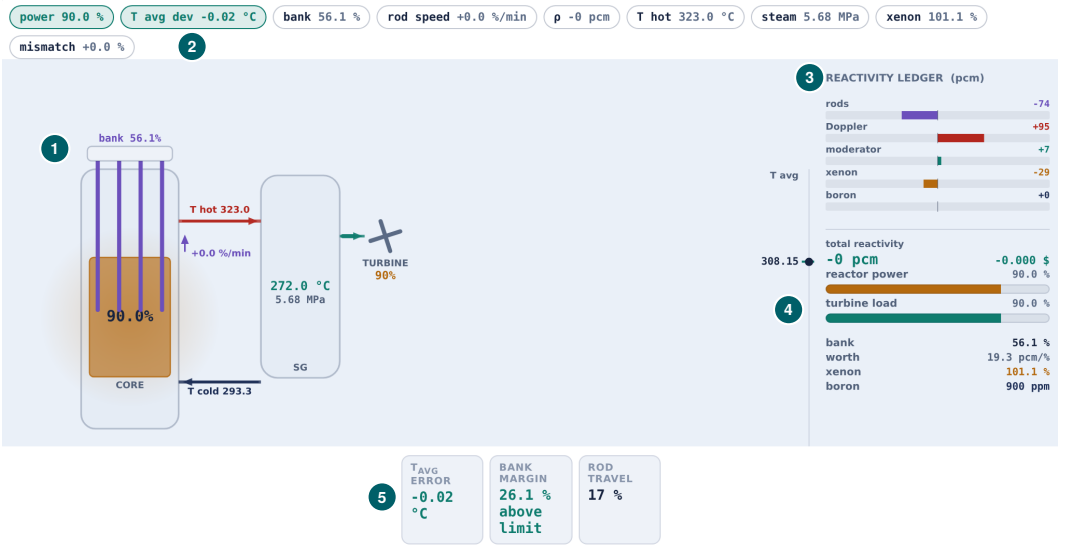

Figure 2.4 Scenario 4 at $t = 30$ min, twenty-eight minutes after the turbine load demand went from 100 to 90 % with $K_c = 3$ and no reset. The cut keeps the plant view with its badge row and three tiles of the control performance panel. Callouts mark (1) the control bank, (2) the deviation of T_{avg}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

What Measures It

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

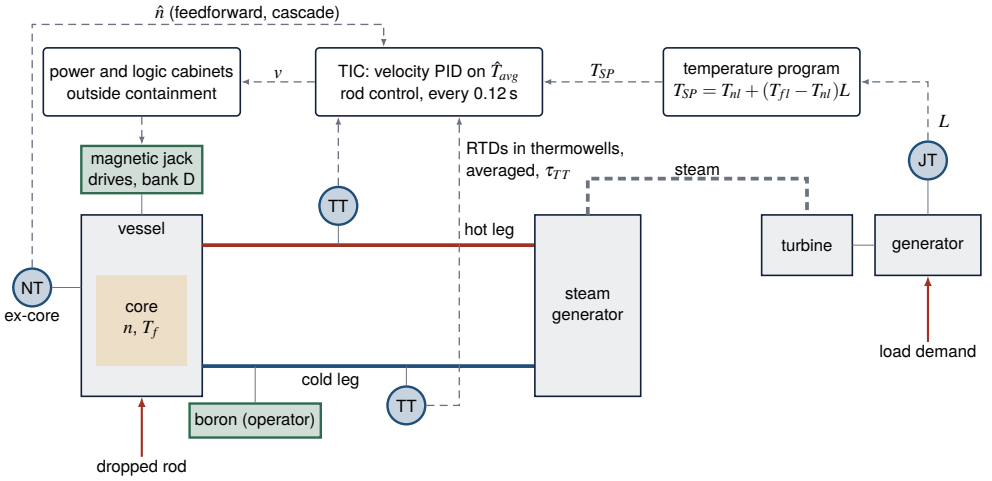

Figure 2.5 The rod control loop of Figure 2.1 with the hardware of a plant. Resistance thermometers (TT) in thermowells on the hot and cold legs give $\hat{T}_{avg}$, an ex-core detector (NT) gives nuclear power, and a load measurement (JT) schedules the setpoint. The velocity PID writes a rod speed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

TECHNOLOGY **Sensor: Resistance Thermometers in Thermowells**

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

What Moves It

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

TECHNOLOGY Final Element: The Magnetic Jack Rod Drive

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

Where the Controller Runs

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

IN PRACTICE Count Every Step

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

What Is Modernizing It

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

2.7 Limitations

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

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

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

2.8 Discussion

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

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

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

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

2.9 Conclusion

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

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

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

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

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

KEY POINTS

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.10 Exercises

2.1 ○ Concept check · 2 min

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

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

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

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

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

2.2 ○ Concept check · 1 min

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

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

2.3 Concept check · 2 min

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

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

2.4 Theory and calculation · 5 min

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

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

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

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

2.5 Theory and calculation · 8 min

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

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

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

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

2.6 Theory and calculation · 5 min

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

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

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

2.7 App challenge · 6 min

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

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

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

2.8 App challenge · 9 min

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

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

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

2.9 App challenge · 10 min

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

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

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

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

2.10 Application area · 4 min

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

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

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

2.11 Application area · 3 min

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

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

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

2.12 Review the agent · 5 min

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

CHAPTER 3

Pressurizer Simulator for Split-Range Control Across Asymmetric Actuators

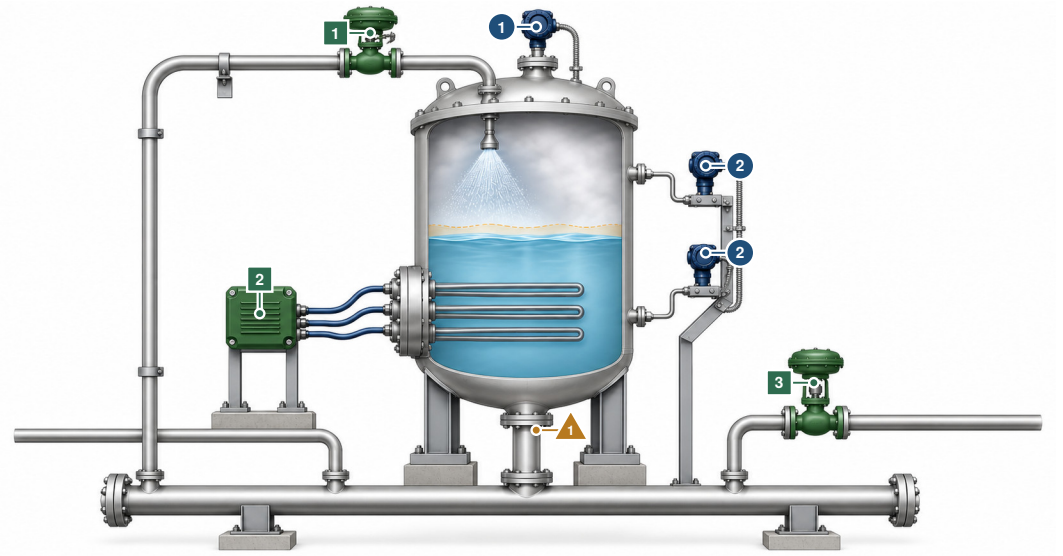

The controller moves 1 spray valve on the line from the cold leg
2 heater power, one modulated bank and four switched banks 3 charging valve into the coolant loop

The controller reads 1 pressure in the steam space 2 level, from taps in the steam and in the water

What disturbs it 1 a load change: coolant swelling up the surge line

The spray condenses steam at the top and the heaters in the water below make steam, two actuators that push the pressure opposite ways from one controller output. Illustration, not to scale.

IN PLAIN TERMS One Controller for Heaters and a Spray

The main idea. A nuclear reactor's cooling water is kept at high pressure by a tall tank with steam at the top. Heaters in the water make steam to raise the pressure, and a cold spray shrinks the steam to lower it.

The challenge. One controller runs both, and the spray works about nine times faster than the heaters. When the plant suddenly needs less steam, water swells into the tank faster than even a full spray can answer.

Why it matters. Knowing what a controller cannot fix shows where safety valves and simple safety rules must take over.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 3.2 and 3.3.8, the fact boxes of Sections 3.3 and 3.7, the try-it boxes of Section 3.7, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 3.7, Section 3.8, Section 3.9 and the open challenge

The control question Heaters make steam slowly and a spray condenses it about nine times faster, both on one controller output. How should the deadband, the valve trim and the loop structure be chosen, and which part of a load rejection can any controller win?

Learning objectives

- Compute the spray's condensing duty and the resting controller output from the model.
- Measure the two halves of a split range and explain their ratio.
- Choose a deadband and a valve trim from what fighting costs under noise.
- Derive the programmed level setpoint from the coolant expansion.
- Separate what tuning changes in a load rejection from what the plant fixes.

Prerequisites Sections 1.3, 1.4.5, 1.6.2 and 1.6.4, and Chapter 2. Section 1.6.3 for the predictive controller only.

The app https://apmonitor.com/apps/pressurizer/

See also Volume 4, Chapter 2 (a protective override), Volume 4, Chapter 3 (limiters on one valve through a selector) and Volume 4, Chapter 10 (three actuators of different speed and cost). Volume 11, Chapter 14 for a split range whose two ends cost money.

3.1 Introduction

Split-range control is usually drawn as one line. A controller output runs from zero to one hundred, one actuator is given the lower part and another the upper part, and a diagram of two triangles crossing in the middle is offered as the whole of the design. That picture makes the arrangement look like a wiring convenience. It is not. Two actuators on one output are two different plants sharing one controller, and the gain, the speed, the operating cost and the failure mode of each are properties of the equipment rather than of the split, so the single tuning has to be a compromise that the diagram gives no way to compute [1, 2]. The pressurizer of a pressurized water reactor is the sharpest available demonstration of that point. Its two actuators differ by a factor of 9.3 in authority, by a factor of five in response time, and by the direction in which each of them moves a second controlled variable, and every one of those numbers is measurable in the same experiment.

The class of problem is broader than the plant. A split range across asymmetric actuators with a deadband appears wherever a variable is pushed both ways by different equipment, such as heating and cooling on a jacketed reactor (Volume 8, Chapter 2), charge and discharge on a battery, acid and base on a neutralization tank, or fuel gas and a vent valve on a fuel header (Volume 11, Chapter 14). A second loop sharing the same vessel through an inventory coupling appears wherever a material balance and an intensive property are controlled on the same holdup. Both structures resist the standard multivariable treatment for the same reason, which is that the interaction lives in an equipment constraint rather than in a transfer function. The relative gain array identifies the pairing but says nothing about which actuator is cheap and which is scarce [3], and the industrial answers to the coupling are a programmed setpoint, a selector and an interlock rather than a decoupler [4]. Recent formalization of split range and of selector-based constraint switching has supplied the analysis those structures always lacked, and it treats them as control design rather than as plumbing [1, 2, 4].

Pressurizer models in the nuclear literature were built for accident analysis and for plant simulators rather than for controller design. The two-region treatment with mass and energy balances on a saturated liquid and a saturated vapor region dates from the 1960s, was extended with bubble rise and condensate drop terms to reproduce measured insurge and outsurge transients [5], and was carried into digital simulation with the spray, the heaters and the relief valves included [6]. Later work showed that thermal equilibrium between the regions is the assumption that fails first and that a non-equilibrium formulation is needed to follow the pressure through a fast surge [7], and comparisons of two-region and four-region models against plant data continue to be published [8]. However, none of that work is written for a student who has to choose a deadband, and a transfer function from heater power to pressure never discloses that the same output also commands a spray

valve nine times stronger. A block diagram of a split range shows the wiring and hides the equipment, which is exactly the information a course develops through worked examples rather than through the diagram itself [9, 10]. This chapter works with a browser pressurizer whose two actuators have to be measured one at a time before any tuning is possible, so that the asymmetry of a split range arrives as a measured number rather than as a caution. The simulator is a thirteen-state non-equilibrium model of the vessel, the reactor coolant system and the surge line, with a split-range mapping, an installed valve characteristic, a low selector, a heater interlock and an optional model predictive controller that plans the two actuators apart. Its scenarios measure the gain ratio, price the deadband, close the second loop and then meet a load rejection. The central result is that the peak of an insurge belongs to the plant rather than to the controller, which leaves the recovery, the level excursion and the equipment interlocks as the only discriminators between architectures.

Section 3.2 describes primary pressure control and its instruments. Section 3.3 gives every constant and every balance needed to rebuild the simulator, with its assumptions in Section 3.3.8. Section 3.4 measures the two halves of the split range and prices the deadband and the valve trim. Section 3.5 describes the three loops, the selector, the interlock and the predictive controller, and Section 3.6 closes the level loop on the same vessel. Section 3.7 follows a load rejection, an outsurge, two hidden failures and the graded challenge, and compares the split range with the predictive controller, and Section 3.8 sets the two loops in their plant hardware. Sections 3.9 to 3.11 state the limitations, discuss what generalizes and conclude, and the exercises follow.

3.2 Primary Pressure Control and Its Instruments

The pressurizer exists because water is nearly incompressible. A four-loop reactor coolant system holds roughly 306 $\mathrm{m^3}$ of liquid outside the pressurizer at about 310 °C, and a rigid circuit full of liquid has no way to accept the volume change that a temperature change forces on it. Attaching a partly steam-filled vessel to the hot leg solves both problems at once, because the bubble accepts volume and its saturation pressure sets the pressure of everything connected to it. Design pressure for the circuit modeled in this chapter is 15.51 MPa (2250 psia), at which saturation occurs at 344.95 °C, and the margin to departure from nucleate boiling in the core is demonstrated against that pressure rather than against a band.

Early plants held the pressure with heaters and a spray line under manual or simple on-off control, and the analysis that justified the arrangement came after the hardware. The first pressurizer models written for transient analysis treated the vessel as saturated liquid and saturated vapor in equilibrium, and adding bubble rise through the pool and condensate rain out of the steam space was what first reproduced the measured pressure during an insurge [5]. Digital simulation followed within a few years and put the spray, the proportional and backup heater banks and the relief valves into one model [6]. The equilibrium assumption survived until experiments showed that the pool and the bubble separate during a fast transient, and that forcing both onto the saturation line predicts the wrong pressure during an insurge and too rapid a recovery after an outsurge [7]. Non-equilibrium two-region formulations are now the standard control-oriented representation, and four-region models that resolve a subcooled layer are used where pool stratification matters [8, 11]. Spray effectiveness received its own treatment, because a droplet falling through a steam space does not reach saturation and the condensation it achieves depends on its size and fall height [12].

The control structure settled earlier than the modeling and has changed very little. Pressure is held by a single controller whose output is split, with the low half driving a modulating heater bank backed by contactor-switched banks and the high half driving spray valves supplied from the cold leg. Level is held by the chemical and volume control system on charging flow, against a setpoint programmed from reactor power rather than fixed, because the coolant inventory the pressurizer must absorb is a function of the average coolant temperature and therefore of load. Above the control band sit a power-operated relief valve at 16.10 MPa and spring-loaded code safety valves at 17.24 MPa, and reactor trips act at 16.55 and 13.10 MPa. They are the limiting and protection layers of Section 1.4.5.

The accident at Three Mile Island in 1979 began in this system, as the history segments of this chapter tell. The regulatory and design response that followed reshaped light water reactor safety practice well beyond the pressurizer itself [13], and it is the reason the distinction between what a control system can do and what a relief valve is for is worth teaching explicitly. Interest in the loop has returned with load following, because a plant asked to change power quickly puts the surge line to work several times a day rather than at startup [14, 15].

> ⧖ HISTORY **A Relief Valve That Did Not Close (1979)**
>
> On 28 March 1979 a pressurized water reactor at Three Mile Island was running at 97 percent power when its feedwater stopped. The coolant heated and expanded, the pressurizer level rose and the steam above it was compressed. Pressure reached 2255 psi, 100 psi above normal, and the power-operated relief valve opened as designed. It should have closed 13 seconds into the accident and instead stayed open for 2 hours and 22 minutes, while a control room light showed only that the signal to close it had been sent. As coolant escaped, the pressurizer level rose while the pressure fell. Operators trained to hold that level cut the high pressure injection from 1000 to under 100 gallons per minute, and much of the core was uncovered and severely damaged [16].

3.3 Process Model

Pressurizer PZR-101 is a 51 m^3 vessel of 3.575 m^2 cross-section connected to the hot leg of a four-loop reactor coolant system by a surge line. At full power it holds 30.6 m^3 of water (18 088 kg at 591 $\mathrm{kg/m}^3$) under a 20.4 m^3 steam bubble (2097 kg at 102.8 $\mathrm{kg/m}^3$), which is 60 % level. Heater bundle EH-101 sits in the bottom 17 % of the vessel and consists of a 300 kW modulating bank plus four 300 kW backup banks on contactors. Spray valves PCV-101 admit up to 28 kg/s of 292.75 °C cold-leg water into the steam space through a nozzle in the upper head, with a continuous bypass of 0.25 kg/s that keeps the line hot. Charging and letdown each pass 4.5 kg/s at the design condition. Pressure transmitter PT-101 and level transmitter LT-101 supply the two controlled variables. Every symbol used below, together with the instrument tags and the geometry the parameters refer to, is drawn in Figure 3.1.

> ✦ DID YOU KNOW? **A Pressurizer of the Same Size**
>
> The pressurizer of the Three Mile Island unit usually held 800 cubic feet of water under 700 cubic feet of steam [16], about 23 and 20 m^3. The vessel of this chapter holds 30.6 m^3 of water under a 20.4 m^3 bubble at full power.

Table 3.1 classifies the two controlled variables, the two manipulated variables, the disturbances and the thirteen integrated states. Working units are seconds, kilograms, cubic meters, megapascals, degrees Celsius and kilowatts throughout, and the pressure control band is 70 kPa either side of the setpoint.

3.3.1 Water and Steam Properties

Properties are least-squares polynomial fits to the IAPWS Industrial Formulation 1997 [17] evaluated on the saturation line at 12, 14, 16 and 18 MPa, with compressed-liquid points at 250 and 300 °C added for the liquid functions. The fitted quantities are $T_{sat}(P)$, $\rho_g^{sat}(P)$ and its inverse $P_{sat}(\rho_g)$, $h_{fg}(P)$, $u_g^{sat}(P)$, $h_g^{sat}(P)$, $\rho_f(T)$, $h_f(T)$ and $c_{p,f}(T)$. Residuals over the fitted range are 0.036 °C on T_{sat}, 0.34 % on saturated vapor density, 1.3 kJ/kg on latent heat, 0.010 MPa on $P_{sat}(\rho_g)$, 0.52 % on liquid density and 2.3 kJ/kg on liquid enthalpy. Arguments are clamped to 11–19 MPa and 200–366 °C, because the vapor-density fit turns over below about 11 MPa. At the design point the fits return $T_{sat} = 344.95\,^{\circ}\mathrm{C}$, $h_{fg} = 964.4\,\mathrm{kJ/kg}$ and $c_{p,f} = 7.697\,\mathrm{kJ/(kg\,K)}$.

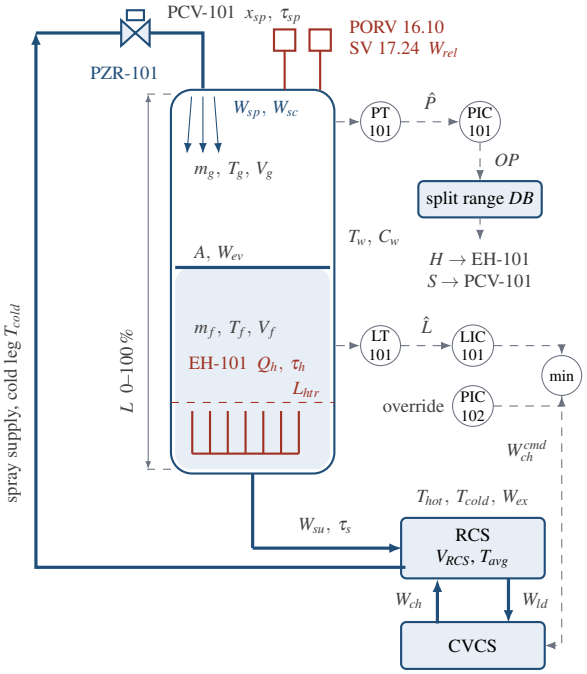

Figure 3.1 Pressurizer PZR-101 with the heater bundle, the spray line from the cold leg, the surge line to the reactor coolant system and the charging and letdown streams. Steam of mass m_g, temperature T_g and density ρ_g occupies V_g above water of mass m_f and temperature T_f occupying V_f, and the two exchange mass at W_{ev} across an interface of area A. Spray W_{sp} through PCV-101 condenses W_{sc} of steam directly, heaters EH-101 deliver Q_h into the pool and uncover below L_{htr}, and the surge line carries W_{su} set by charging, letdown, spray and the thermal expansion W_{ex} of the coolant at T_{avg}. PIC-101 splits one output OP into a heater demand H and a spray demand S across a deadband DB, and LIC-101 shares the charging valve with high-pressure override PIC-102 through a low selector.

3.3.2 Two Regions That Are Allowed to Disagree

The vessel carries four thermodynamic states rather than the usual two, because neither region is forced onto the saturation line. Mass conservation on the two regions is Equation (3.1).

$$\begin{aligned} \frac{\mathrm{d}m_g}{\mathrm{d}t} &= W_{ev} - W_{sc} - W_{rel} \\ \frac{\mathrm{d}m_f}{\mathrm{d}t} &= W_{su} + W_{sp} + W_{sc} - W_{ev} \end{aligned} \tag{3.1}$$

where m_g and m_f are the steam and liquid masses in kg, W_{ev} is the net interfacial evaporation, W_{sc} is the steam condensed by the spray, W_{rel} is the relief flow, W_{su} is the surge line flow positive into the vessel and W_{sp} is the spray flow, all in kg/s. The spray appears in both balances, once as liquid arriving and once as the steam it removes, which is the algebraic statement that a spray droplet is a mass transfer and a heat transfer at the same time. Pressure follows from the bubble alone, as shown in Equation (3.2).

$$P = P_{sat}(\rho_g) + zR\rho_g\left(T_g - T_{sat}(P_{sat}(\rho_g))\right) \tag{3.2}$$

where $\rho_g = m_g/V_g$ is the bubble density in kg/m^3, $R = 0.4615\,\mathrm{kJ/(kg\,K)}$ is the gas constant of steam and $z = 1.88$ is a real-gas multiplier on the ideal-gas $(\partial P/\partial T)_\rho$, because steam at 15 MPa departs strongly from ideal behavior. The second term is evaluated in kPa and scaled by 10^{-3} to match the MPa of the first. The second term in (3.2) is what an equilibrium model omits, and it is large.

Table 3.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
P (PT-101)	CV, measured	15.51 MPa	15.0–16.0 MPa SP
L (LT-101)	CV, measured	60 %	20–90 % SP
OP (PIC-101)	MV	40.75 %	0–100 %
W_{ch}	MV	4.5 kg/s	0–12 kg/s
W_{ld}	operator input	4.5 kg/s	0–16 kg/s
turbine load	DV, measured	100 %	0–110 %
T_{avg} bias	DV, unmeasured	0 K	−6 to 6 K
banks lost	DV, unmeasured	0	0–4
spray stuck open	DV, unmeasured	0 %	0–100 %
insulation	DV, unmeasured	1×	0.5–6×
m_g, T_g	states	2097, 344.5	kg, °C
m_f, T_f	states	18 088, 344.96	kg, °C
T_w	state	344.4 °C	40–380 °C
Q_h	state	141.8 kW	0–1500 kW
x_{sp}	state	0	0–1
T_{avg}	state	309.5 °C	—
P_{th}	state	100 %	0–120 %
W_{ch}, W_{ld} lags	states	4.5 kg/s	≥ 0
V_c	state	0 m^3	−8 to 8 m^3
W_{ex}	state	0 kg/s	±400 kg/s

Perturbing the bubble along the saturation line at the design density of 102.8 kg/m^3 gives a stiffness of 0.0936 MPa per kg/m^3, while compressing it adiabatically through the model's own energy balance gives 0.1949 MPa per kg/m^3, a factor of 2.08. An insurge therefore raises the pressure about twice as fast as a saturated-bubble model predicts, which is the mechanism the non-equilibrium pressurizer literature was written to capture [7, 8].

The interface transfers mass in whichever direction the pool departs from saturation, with two coefficients that differ by more than two orders of magnitude, as shown in Equation (3.3).

$$W_{ev} = -\frac{k\,(T_{sat} - T_f)}{h_{fg}}, \quad k = \begin{cases} k_c & T_f < T_{sat} \\ k_f & T_f > T_{sat} \end{cases} \tag{3.3}$$

where $k_c = 40\,\mathrm{kW/K}$ is the condensation conductance onto a subcooled pool and $k_f = 10000\,\mathrm{kW/K}$ is the bulk flashing conductance out of a superheated one. Condensation onto a quiescent pool is a surface process across 3.575 m^2, and with a pool heat capacity of 139 220 kJ/K its relaxation time $m_f c_{p,f}/k_c$ is 3480 s. Flashing is bulk nucleation through the whole inventory and relaxes in $m_f c_{p,f}/k_f = 13.9\,\mathrm{s}$. The 250:1 ratio is the reason a depressurization is self-limiting and a repressurization is not, and it is a property of the fluid rather than of the control system.

3.3.3 Spray, Heaters and the Vessel Wall

The spray condenses steam in proportion to the enthalpy deficit the droplets carry into the steam space, as shown in Equation (3.4).

$$W_{sc} = \eta_{sp}\, W_{sp}\, \frac{h_f(T_{sat}) - h_f(T_{cold})}{h_{fg}} \tag{3.4}$$

where $\eta_{sp} = 0.75$ is the fraction of that deficit actually used before the droplets reach the pool, which stands in for the droplet size distribution and fall history a spray model resolves explicitly [12]. At the design point $h_f(T_{sat}) = 1632.8\,\mathrm{kJ/kg}$ against $h_f(292.75\,^\circ\mathrm{C}) = 1299.4\,\mathrm{kJ/kg}$, a deficit of 333.5 kJ/kg

against a latent heat of 964.4 kJ/kg, so the full 28 kg/s of spray condenses 7.261 kg/s of steam. That is a condensing duty of 7.00 MW. The whole heater bank is 1.5 MW, and the first half of the asymmetry is already visible in those two numbers.

Spray flow follows the installed valve characteristic of Equation (3.5).

$$W_{sp} = W_{sp,\min} + \left(W_{sp,\max} - W_{sp,\min}\right) x_{sp}^{n} \qquad (3.5)$$
$$\tau_{sp} \frac{\mathrm{d}x_{sp}}{\mathrm{d}t} = \frac{S}{100} - x_{sp}$$

where x_{sp} is fractional valve travel, S is the spray demand in percent, n is the characteristic exponent with $n = 1$ linear and $n = 2.5$ approximating equal-percentage trim, and $\tau_{sp} = 4\,\mathrm{s}$ covers the valve stroke and the line transport. Travel is clipped to 0–1 after every step, and the derivative is zeroed at each bound so that a fixed-step integrator has no way to overshoot the stop and come back. Heater duty follows the demand through the bank logic of Equation (3.6).

$$Q_h^{cmd} = Q_p \min\left(\frac{H}{H_p}, 1\right) + Q_b \sum_{i=0}^{n_b-1} b_i \qquad (3.6)$$
$$\tau_h \frac{\mathrm{d}Q_h}{\mathrm{d}t} = Q_h^{cmd} - Q_h$$

where H is the heater demand in percent, $Q_p = 300\,\mathrm{kW}$ is the modulating bank rated full at $H_p = 20\,\%$, $Q_b = 300\,\mathrm{kW}$ is each of the $n_b = 4$ backup banks, b_i is the latched state of bank i and $\tau_h = 20\,\mathrm{s}$ is the sheath thermal lag. Bank i picks up at $H \geq 24 + 18i$ percent and drops out at $H < 24 + 18i - 6$, so the pick-up demands are 24, 42, 60 and 78 % and each bank carries 6 % of hysteresis against contactor chatter.

The vessel wall participates as a single lumped capacity, as shown in Equation (3.7).

$$C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} = hA_{wf}(T_f - T_w) + hA_{wg}(T_g - T_w) - \iota\, UA_{amb}(T_w - T_{amb}) \qquad (3.7)$$

where $C_w = 20000\,\mathrm{kJ/K}$ represents the inner few centimeters of the forging, hA_{wf} and hA_{wg} are the wall conductances to water and steam, ι is an insulation-condition multiplier and $UA_{amb} = 0.203\,\mathrm{kW/K}$ is the loss to containment. That loss is 59.8 kW at the design point.

3.3.4 Energy Balances on the Two Regions

The liquid energy balance collects the surge line, the spray, the condensate, the flashing, the heaters and the wall, as shown in Equation (3.8).

$$\begin{aligned} m_f c_{p,f} \frac{\mathrm{d}T_f}{\mathrm{d}t} = {} & \max(W_{su}, 0)\left[h_f(T_{hot}) - h_f\right] \\ & + W_{sp}\left[h_f(T_{cold}) - h_f\right] \\ & + W_{sc}(h_g - h_f) - W_{ev}\left[h_g^{sat} - h_f\right] \\ & + Q_h - hA_{wf}(T_f - T_w) \end{aligned} \qquad (3.8)$$

where $h_f = h_f(T_f)$ is the pool enthalpy and h_g is the enthalpy of the bubble including its superheat. An outsurge carries pool enthalpy away and therefore contributes nothing to (3.8), which is why only the positive part of W_{su} appears. The steam energy balance carries the compression work explicitly,

as shown in Equation (3.9).

$$\begin{aligned} Q_g &= W_{ev} h_g^{sat} - W_{sc} h_g - W_{rel} h_g - 1000\,P \frac{\mathrm{d}V_g}{\mathrm{d}t} \\ &\quad - hA_{wg}(T_g - T_w) - k_{ifg}(T_g - T_f) \\ c_{v,g} \frac{\mathrm{d}T_g}{\mathrm{d}t} &= \frac{Q_g - u_g \dot{m}_g}{m_g} - \frac{\mathrm{d}F}{\mathrm{d}\rho_g} \frac{\mathrm{d}\rho_g}{\mathrm{d}t} \end{aligned} \tag{3.9}$$

where $c_{v,g} = 3.30\,\mathrm{kJ/(kg\,K)}$, $k_{ifg} = 3.0\,\mathrm{kW/K}$ is a sensible conductance between steam and interface, and $u_g = F(\rho_g) + c_{v,g} T_g$ with $F(\rho) = u_g^{sat}(P_{sat}(\rho)) - c_{v,g} T_{sat}(P_{sat}(\rho))$ is the internal energy written so that $\dot{u}_g$ is explicit. The factor 1000 converts MPa m^3 to kJ. Both sides of (3.8) and (3.9) carry kW.

3.3.5 The Surge Line and the Reactor Coolant System

The reactor coolant system is rigid and liquid-full, so the mass it can hold is a function of its average temperature alone and everything it cannot hold arrives in the pressurizer. That closure is Equation (3.10).

$$\begin{aligned} W_{su} &= W_{ch} - W_{ld} - W_{sp} + W_{ex} \\ \tau_{loop} \frac{\mathrm{d}W_{ex}}{\mathrm{d}t} &= -V_{RCS} \frac{\mathrm{d}\rho_f}{\mathrm{d}T} \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} - W_{ex} \end{aligned} \tag{3.10}$$

where $V_{RCS} = 306\,\mathrm{m^3}$ and $\tau_{loop} = 8\,\mathrm{s}$ is the loop transport time. The lag matters. T_{avg} is one lumped number but the water it describes has to travel, so the expansion reaches the surge line spread over a loop transit rather than in the instant the turbine valve moves, and without it the whole insurge would land inside a single spray valve stroke. At 309.5 °C the coolant loses 2.601 kg/(m^3 K), so one kelvin of T_{avg} is 796 kg up the surge line. The primary loop itself is a single energy balance with a proportional rod controller on a programmed average temperature, the loop that Chapter 2 develops, as shown in Equation (3.11).

$$\begin{aligned} C \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} &= (P_{th} - \ell)\, q_{pct} - K_{sg}(T_{avg} - T_{prog}) \\ \tau_{rod} \frac{\mathrm{d}P_{th}}{\mathrm{d}t} &= \mathrm{clamp}[\ell + K_{rod}(T_{prog} - T_{avg}),\, 0,\, 115] - P_{th} \end{aligned} \tag{3.11}$$

where P_{th} is reactor thermal power in percent, ℓ is the turbine load in percent, $q_{pct} = 34\,110\,\mathrm{kW/\%}$, $C = 3.5 \times 10^6\,\mathrm{kJ/K}$, $K_{sg} = 90\,000\,\mathrm{kW/K}$ lumps steam generator self-regulation with the steam dump, $K_{rod} = 5\,\mathrm{\%/K}$, $\tau_{rod} = 25\,\mathrm{s}$ and $T_{prog} = 294 + 0.155\,\ell$ in °C is the programmed average temperature. Hot and cold leg temperatures are $T_{avg} \pm \frac{1}{2}\Delta T_{loop}\,\ell/100$ with $\Delta T_{loop} = 33.5\,\mathrm{K}$, which puts the cold leg at 292.75 °C at full load. The coolant is not perfectly rigid either, and its compressibility is carried as an equivalent volume held outside the vessel, as shown in Equation (3.12).

$$\tau_s \frac{\mathrm{d}V_c}{\mathrm{d}t} = \kappa V \left(P - P_{ref}\right) - V_c, \qquad V_f = \frac{m_f}{\rho_f(T_f)} - V_c \tag{3.12}$$

where $\kappa V = 0.40\,\mathrm{m^3/MPa}$ is 306 m^3 of 300 °C water at an isentropic compressibility of 1.3/GPa, $P_{ref} = 15.51\,\mathrm{MPa}$ and $\tau_s = 2.5\,\mathrm{s}$ is the surge line hydraulic response. Level is $L = 100\,V_f/V$, and $V_g = V - V_f$.

3.3.6 Relief, Protection and Instruments

The relief valves watch the true pressure rather than the transmitter, because they are mechanical. Two power-operated relief valves passing a combined 53 kg/s lift at 16.10 MPa and reseat at 15.86 MPa, and code safety valves passing 159 kg/s lift at 17.24 MPa and reseat at 16.90 MPa. Reactor trips

Table 3.2 Model Parameters. Every Value Is Adjustable in the Application. The Two Marked * Are Hidden From the User Interface.

Sym.	Meaning	Value	Unit	Sym.	Meaning	Value	Unit
V	vessel volume	51.0	m^3	Q_p	modulating bank	300	kW
A	interface area	3.575	m^2	Q_b	each backup bank	300	kW
L_{htr}	heater uncovery	17	%	n_b	backup banks	4	—
$V_{g,\min}$	smallest bubble*	0.5	m^3	H_p	demand for full Q_p	20	%
C_w	participating wall	20 000	kJ/K	τ_h	heater sheath lag	20	s
hA_{wf}	wall to water	100	kW/K	τ_{ch}, τ_{ld}	charging, letdown	8.0, 6.0	s
hA_{wg}	wall to steam	12	kW/K	V_{RCS}	coolant volume	306	m^3
UA_{amb}	loss to containment	0.203	kW/K	κV	coolant compliance	0.40	m^3/MPa
k_c	condensation	40	kW/K	τ_s	surge line response	2.5	s
k_f	flashing	10 000	kW/K	τ_{loop}	loop transport	8.0	s
k_{ifg}	steam to interface	3.0	kW/K	C	RCS heat capacity	3.5×10^6	kJ/K
T_{amb}	containment	50	°C	K_{sg}	SG self-regulation	90 000	kW/K
η_{sp}	spray efficiency	0.75	—	K_{rod}	rod control gain	5.0	%/K
$c_{v,g}$	steam c_v	3.30	kJ/(kg K)	τ_{rod}	rod and fuel	25	s
z	real-steam factor	1.88	—	ΔT_{loop}	hot–cold leg rise	33.5	K
$W_{sp,\max}$	spray, full travel	28.0	kg/s	P_{ref}	compliance ref.*	15.51	MPa
$W_{sp,\min}$	continuous bypass	0.25	kg/s	τ_{Pm}	pressure transmitter	1.0	s
τ_{sp}	spray valve and line	4.0	s	τ_{Lm}	level transmitter	4.0	s

are annunciated at 16.55 and 13.10 MPa and the run continues past them so that the transient stays readable. Both transmitters are first-order lags of 1.0 s on pressure and 4.0 s on level. Each is stepped by the discrete form $\hat{y}_k = \alpha \hat{y}_{k-1} + (1-\alpha) y_k$ with $\alpha = \exp(-\Delta t/\tau)$, which is the exact solution over a zero-order hold and is therefore stable at any period, and each accepts an optional colored-noise term scaled so that its stationary standard deviation equals the value set on the disturbance slider. A heater energized with the level below $L_{htr} = 17\,\%$ accumulates uncovered time, and 30 s of it destroys the bundle permanently. The interlock counts the heater command rather than the delivered duty, because a blocked heater is de-energized at once and what decays over the following two minutes is heat already stored in the sheath.

3.3.7 Parameters, Initialization and Numerics

Table 3.2 lists the constants that fix the behavior discussed in this chapter, and every scenario starts from a solved steady state rather than from a guess.

Initialization solves the steady state directly rather than iterating on the whole vessel. The requested level fixes V_f and V_g, the spray position fixes W_{sc} through (3.4), steady bubble mass requires $W_{ev} = W_{sc}$, and (3.3) then fixes the pool superheat exactly. Only the two-way coupling between the steam temperature and the wall needs a loop, and it converges in a handful of passes because both sit within a kelvin of the pool. The liquid balance (3.8) is finally solved for the heater duty rather than integrated, which is what makes the resting duty a computed quantity. At the design condition it returns 141.8 kW, of which 59.8 kW leaves through the insulation and the rest re-boils the 0.25 kg/s of continuous bypass spray. That duty is 9.45 % of heater demand and, through the split-range mapping of Section 3.4, a controller output of 40.75 %. The application verifies at startup that this state produces derivatives whose largest relative magnitude is below 10^{-9}.

The thirteen states are advanced by fourth-order Runge-Kutta on a fixed 0.05 s step, with the three controllers executing on a 0.5 s scan that matches a real pressurizer control system and the historian keeping ten records per simulated second. Every random draw comes from one generator seeded at 20260318, so a run with transmitter noise repeats exactly. Speed, seed and Reset behave as Section 1.6.1 describes.

3.3.8 Assumptions

Each simplification below is deliberate, and each fails somewhere.

One pool and one bubble. The whole vessel is two lumped regions with one interface. A real pressurizer stratifies, because the surge line discharges into the bottom of a pool that can be twenty kelvin colder at the top than at the nozzle, and computational studies of turbulent penetration during an in-surge resolve that structure [11]. The lumped treatment overstates how quickly the pool responds to a surge and produces no surge line thermal load at all.

The interface coefficients are judged, not correlated. Condensation at 40 kW/K and flashing at 10 000 kW/K are defensible orders of magnitude, and neither is a measurement. The 250:1 ratio between them is the physically important part and it is robust, while the individual values set how long the application takes to recover from a subcooling event.

Spray is one efficiency number. Equation (3.4) carries no droplet size distribution and no fall height, so the model cannot show spray becoming less effective as the bubble shrinks and the fall path shortens [12].

Properties are fits over a bounded range. Outside 11–19 MPa and 200–366 °C the arguments are clamped, so a depressurization that would take the loop two-phase cannot be followed.

The reactor coolant system is one node. A single average temperature, one lumped heat capacity and one self-regulation coefficient generate a realistic insurge on a load rejection and say nothing about steam generator level, secondary pressure or reactivity feedback beyond a proportional rod controller [14].

The relief valves are ideal. Fixed-capacity orifices with clean lift and reseat leave out the failure that matters most in practice, which is a relief valve that lifts and does not reseat [13].

Heater burnout is a timer. Thirty seconds uncovered and energized destroys the bundle. The real failure depends on sheath material, heat flux and how the water left.

Disturbances are idealized. A spray valve that sticks at exactly 45 % is not a failure a real valve has, but it isolates a continuous condensing load that the loop compensates perfectly. A real governor would round the corner off a load rejection, while a step gives a response that can be read directly from the trend.

> ⧖ HISTORY **Warnings Before the Accident (1977 to 1978)**
>
> The pattern had been seen before. In September 1977 a relief valve stuck open at a plant of the same design running at 9 percent power. Its pressurizer level rose while the pressure fell, operators interfered with the emergency injection, apparently relying on that rising level, and the block valve was closed about 20 minutes later. An engineer at another utility analyzed rising level with falling pressure more than a year before 1979, and relief valves had stuck open nine times at plants of that design [16]. At the accident unit itself a relief valve had stuck open a year earlier. The indicator light installed after that event showed that a closing signal had been sent, not whether the valve had closed [16].

3.4 The Split Range and Its Two Plants

3.4.1 The Mapping

One controller output is shared by two actuators across a deadband, as shown in Equation (3.13).

$$H = \frac{100\left[(50 - \frac{DB}{2}) - OP\right]}{50 - \frac{DB}{2}} \qquad S = \frac{100\left[OP - (50 + \frac{DB}{2})\right]}{50 - \frac{DB}{2}} \tag{3.13}$$

where OP is the controller output in percent, DB is the deadband width in percent of output and both H and S are clamped to 0–100 %, so between the two edges neither actuator is asked for anything.

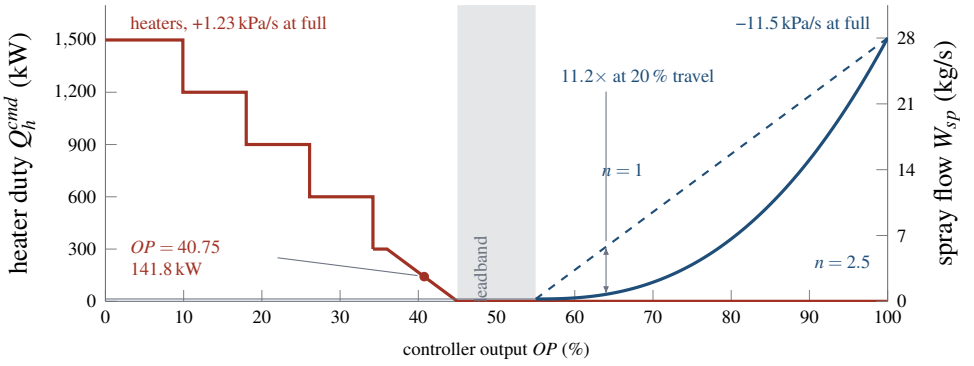

Figure 3.2 Split-range mapping from one controller output to two actuators of very different size. Heater duty is the staircase on the left axis, because the modulating bank saturates at 20 % demand ($OP = 36$) and the four backup banks pick up on contactors at 24, 42, 60 and 78 % demand. Spray flow is on the right axis, with the equal-percentage trim (solid curve) carrying 0.75 kg/s at 20 % travel against 5.80 kg/s for a linear valve (dashed line), and both trims reach the same 28 kg/s wide open. The design operating point sits at $OP = 40.75$, which is four points from the deadband edge rather than at the crossover.

At the default $DB = 10\,\%$ the heater band is $0 \le OP \le 45$ and the spray band is $55 \le OP \le 100$. Figure 3.2 draws the mapping through the actuators to the physical quantities they deliver.

3.4.2 Measuring the Two Halves

Scenario 1 places both actuators in manual at the design state and drives each one fully, and the measurement takes forty seconds. Driving the heater demand to 100 % in the first scan after Run raises the pressure at +1.2319 kPa/s averaged from $t = 10$ to 50 s. Driving the spray demand to 100 % from the same state lowers it at −11.5129 kPa/s. The ratio is **9.35**, from the two halves of one output. A demand set before Run acts half a second earlier, from $t = 0$, and gives a slightly smaller ratio over the same window.

Two mechanisms produce it, and they are separable. The duties differ, because 28 kg/s of water 52.2 K below saturation is the 7.00 MW condensing load of (3.4) against a heater bank of 1.5 MW. The paths differ as well. The spray removes steam directly from a bubble whose entire heat capacity is $m_g c_{v,g} = 6920\,\mathrm{kJ/K}$, while the heaters have to raise the pool above saturation before (3.3) produces any steam at all, and the pool and the participating wall carry 139 220 kJ/K and 20 000 kJ/K between them. Measured sixty seconds into a step to full heater demand, the delivered duty is 1431 kW and the flashing rate is 0.160 kg/s, which is 154 kW of steam production. Only 10.7 % of what the heaters put in is leaving as steam, and the rest is still going into sensible heat. A heater is a slow actuator because of the water in front of it and not because of its sheath.

The same asymmetry appears in the settling behavior, and it points the other way. Scenario 5 steps the pressure setpoint +100 kPa and then −200 kPa with the tuned loop in service. The upward move reaches a ±10 kPa band in 75.9 s with the heaters peaking at 1411 kW, overshoots the setpoint by 13.7 kPa and stays inside the band only from 207 s after the step. The downward move covers 154 kPa of its 200 kPa in the first forty seconds with the spray peaking at 17.5 kg/s, and then takes 126 s in total to reach the same band, because holding a lower pressure requires the pool to cool with the saturation temperature and eighteen tonnes of water do not cool quickly. Instantaneous authority and settling time are different questions, and a split range has to respect both.

3.4.3 Why the Answer Is Not a Smaller Gain

A single proportional gain has to serve a plant whose gain changes by 9.35 across the crossover. Reducing the gain until the spray half is comfortable gives up the heater half, which is already the

Table 3.3 Cost of the Deadband. Scenario 4 With 4 kPa of Transmitter Noise at the Shipped Seed, $K_c = 1500\,\%/\mathrm{MPa}$ and the Deadband Set Before the Run, Measured Over the 600 s Following the Setpoint Step at $t = 60\,\mathrm{s}$.

DB (%)	0	4	8	10	12	20	30
Fighting (s)	182.0	94.4	72.0	64.3	64.6	65.1	59.0
Heater energy (kWh)	47.01	43.03	41.64	41.34	41.11	40.67	39.83
Spray water (kg)	264	223	209	206	203	199	191

weak side. The installed valve characteristic solves the problem where it arises, at the actuator. With $n = 2.5$ in (3.5) the spray delivers 0.746 kg/s at 20 % travel against 5.800 kg/s for a linear valve, a factor of 11.2 less flow over the bypass at the position the controller normally occupies, and both valves deliver exactly 28.0 kg/s wide open where an insurge needs the capacity. The characteristic buys a lower gain in the region that matters without surrendering the capacity that a transient needs, which is the same argument that supports equal-percentage trim in any loop with a wide turndown [9]. The trim does not buy a faster setpoint response. With $n = 1$ the downward step of scenario 5 reaches its band sooner, in 85.6 s, because the linear valve has more gain near the seat. What the equal-percentage trim buys appears under noise. With 4 kPa of transmitter noise on the same 900 s run, the heaters and the spray are in service together for 234 s with $n = 1$ against 70 s with $n = 2.5$.

3.4.4 The Deadband Is Not Free

Scenario 4 runs the split range with the deadband at zero, a gain of 1500 %/MPa and 4 kPa of transmitter noise, and steps the setpoint +50 kPa at $t = 60\,\mathrm{s}$. The output has to pass through the crossover to move between actuators, and with noise on the measurement it does not pass through so much as loiter. Measured over 200 to 800 s, after the setpoint transient has cleared, the heaters and the spray are in service together for 139.5 s at $DB = 0$ and for 0.0 s at $DB = 10\,\%$. The noise is seeded, and over eight seeds the first figure ranges from 122 to 185 s while the second never exceeds 6.1 s. Over the whole 800 s run the plant consumes 57.61 kWh of heat and 340 kg of spray water at $DB = 0$ against 49.57 kWh and 256 kg at $DB = 10\,\%$, an increase of 16.2 % in heater energy and 33.0 % in spray water for no change in the pressure trace that an operator would notice.

Table 3.3 sweeps the width. The metrics improve steeply to about 8 % and then flatten, which makes the choice of ten points a design decision rather than a large number. Beyond 8 % every difference in the table is within the spread across noise seeds, which is 34 to 72 s of fighting over eight seeds between 8 and 30 %. The failure at the other end is not visible in this table, because a deadband wide enough to matter simply leaves the pressure sitting off setpoint until it drifts to an edge, and the run never asks it to.

The waste is the smaller part of the argument. Every spray cycle drops 292.75 °C water into a 344.95 °C pool and thermally cycles the nozzle and the surge line, and every heater cycle throws contactors rated for a finite number of operations. Neither is visible on a pressure trend, which is why the application counts the seconds directly.

3.5 Controller Implementation

3.5.1 The Shared Algorithm

The three loop controllers in the application are instances of the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement. In the notation of Equations (1.13) to (1.16), the running sum S_k of proportional and derivative increments is never clamped, v_k is the output the algorithm demands before the limit, and the integral I_k is the only term anti-windup touches, through the back-calculation tracking time T_t. The derivative filter on the pressure loop has $\tau_f = 4\,\mathrm{s}$. Manual operation and transfer are bumpless in both directions as Chapter 1 describes [18]. The scan period is $\Delta t_c = 0.5\,\mathrm{s}$, independent of the 0.05 s physics step.

3.5.2 The Three Loops

Loop action follows the convention of Section 1.3. PIC-101 is *direct* acting, so $a = -1$: a measurement above setpoint drives the output up into the spray half, and a rising output lowers the pressure. Its output limits are 0 and 100 % and its default tuning is $K_c = 800\,\%/\mathrm{MPa}$, $\tau_I = 240\,\mathrm{s}$ and $\tau_D = 12\,\mathrm{s}$. LIC-101 is *reverse* acting with $a = +1$, so a level below setpoint calls for more charging, and its output limits are 0 and 12 kg/s with $K_c = 1.0\,\mathrm{kg}/(\mathrm{s}\,\%)$ and $\tau_I = 400\,\mathrm{s}$. The high-pressure override PIC-102 is also reverse acting on the same valve, because charging water compresses the bubble, with $K_c = 90\,\mathrm{kg}/(\mathrm{s\,MPa})$, $\tau_I = 150\,\mathrm{s}$ and a setpoint of 15.72 MPa. PIC-101 and PIC-102 track at the shared default $T_t = \tau_I$, and LIC-101 uses $T_t = 30\,\mathrm{s}$. No output carries a rate limit, because the actuator lags of (3.5) and (3.6) already impose the physical rate. Pressure loop tuning is bounded to $0 \le K_c \le 2500$, $5 \le \tau_I \le 1200\,\mathrm{s}$, $0 \le \tau_D \le 120\,\mathrm{s}$ and $0 \le DB \le 40\,\%$.

The tracking time matters on this loop because the outsurge of scenario 7 takes the pressure a megapascal below setpoint, where the proportional and derivative sum S_k reaches −830 % against an output held at zero. Until 2026-10-08 the application tracked at $T_t = 30\,\mathrm{s}$ on PIC-101 and PIC-102. Back-calculation then drove I to +594 % so that v_k followed the clamp, and as the pressure recovered S_k shrank faster than I unwound, so v_k crossed the deadband with the pressure still 793 kPa low and the spray was commanded for 589 s while the pressure was more than 70 kPa below setpoint. That is integral windup produced by the anti-windup itself once the proportional term alone exceeds the output range eightfold. With $T_t = \tau_I$, I peaks at 71 % and the same condition lasts 12 s. The price is a slower unwind after a short saturation, visible as the 13.7 kPa overshoot of the upward step in scenario 5. The override suffered the same mechanism, and with $T_t = 30\,\mathrm{s}$ the selector handed it the charging valve at $t = 150.5\,\mathrm{s}$ in scenario 11 with the pressure 836 kPa below its setpoint.

Five architectures are selectable. MANUAL leaves all four stations to the operator. PRESSURE ONLY closes PIC-101 and leaves the level in manual. BOTH LOOPS adds LIC-101 against a fixed level setpoint with nothing between the two loops. BOTH PLUS OVERRIDES adds the three pieces of logic described in Section 3.6. MPC keeps that level side and replaces the split range with the predictive controller of Section 3.5.4.

3.5.3 The Selector and the Interlock

With the override architecture selected, LIC-101 and PIC-102 both run continuously and the charging valve takes whichever asks for less, as shown in Equation (3.14).

$$W_{ch}^{cmd} = \min[u_L,\, u_O]\,, \qquad I_{loser} \mathrel{+}= \frac{W_{ch}^{cmd} - u_{loser}}{T_{sel}}\Delta t_c \tag{3.14}$$

where u_L and u_O are the two unselected outputs and $T_{sel} = 20\,\mathrm{s}$. The loser is unwound toward the selected valve position by back-calculation and not by assigning its integral the selected output directly. The distinction is not cosmetic. Snapping the loser's integral forces it to reproduce a number the winner may itself be producing only because it is saturated, the two controllers then exchange roles on alternate scans, and the valve can sit shut while both loops call for it to open. That failure looks like a plant fault rather than a configuration error, and it is the classic way a selector scheme is got wrong [4]. Tracking the loser to the valve has its own consequence. The loser's output sits near the selected one, so its proportional term decides the hand-over, and a fast pressure rise passes the valve to PIC-102 before the pressure reaches its setpoint. With the override setpoint at its 16.20 MPa maximum, scenario 11 still hands it the valve for 23.5 s during the reload. This early proportional hand-over is retained.

The heater interlock is separate from both loops. When the measured level falls below $L_{htr} = 17\,\%$ in the override architecture, the heater demand is forced to zero regardless of what the pressure controller asks for. The loop that most wants to add heat when the pressure is low is exactly the loop that must not be allowed to when the level is low, and the two conditions arrive together during an outsurge.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.5.4 Model Predictive Control

The optional predictive controller plans the heater duty Q and the spray flow W as two inputs instead of one split output, within the general arrangement of Section 1.6.3, while charging stays with LIC-101, the programmed setpoint, the interlock and PIC-102. Every $T_s = 2\,\mathrm{s}$, or four scans, it solves a quadratic program over $N = 90$ steps (180 s) and holds the first move until the next decision [19, 20]. Its model is the balance of Section 3.3 with the declared values of Table 3.2, not those of the run, linearized at 15.51 MPa, 60 % level, full load, charging and letdown at 4.5 kg/s and the balancing 141.82 kW (largest derivative 7.4×10^{-13}) and sampled with a zero-order hold, since per-input step-test fits cannot share the pool and bubble both inputs act through. It keeps the twelve vessel and transmitter states and omits the reactor coolant system, the contactors and the trim, which the planned kilowatts and kilograms per second are mapped through afterwards. The design point is on the flashing side of (3.3), so the model relaxes a subcooled pool through k_f. Over 300 s steps it follows a 150 kW heater step to 0.02 kPa root-mean-square and a 2.78 kg/s spray step to 1.4 % of the response, and one minute after a 3 kg/s charging cut it predicts −47.7 kPa against −53.5 kPa, but after a 5 kg/s rise 79.5 kPa against 300.8 kPa. A subcooled-side model predicts 295.8 and −177.5 kPa, so no single linearization serves both.

An unmeasured surge d_W (kg/s) entering as W_{ex} does and an unmeasured heat d_Q (kW) entering with the heaters augment the state as random walks [21]. A steady-state Kalman filter on the model sampled at Δt_c, with random-walk variances of $(2\,\mathrm{kg/s})^2$ and $(20\,\mathrm{kW})^2$ per scan against $(0.002\,\mathrm{MPa})^2$ and $(0.1\,\%)^2$ on the transmitters, runs at every scan under every architecture and is told the requested duty and spray and the charging and letdown commands, so a lost bank or a stuck valve appears in d_Q or d_W. Over twelve move blocks of 2, 2, 2, 2, 4, 4, 6, 8, 12, 18, 30 and 90 s the plan minimizes Equation (3.15).

$$J = \sum_{k=1}^{N} \left(\frac{P_k - P_{\mathrm{SP}}}{P_b} \right)^2 + \sum_{k=0}^{N-1} \left(c_h \frac{Q_k}{Q_{\max}} + c_s \frac{W_k}{W_{sp,\max}} \right) + \lambda \sum_{j=1}^{12} \left[\left(\frac{\Delta Q_j}{Q_{\max}} \right)^2 + \left(\frac{\Delta W_j}{W_{sp,\max}} \right)^2 \right] + J_s \tag{3.15}$$

where P_k is predicted with d_W, d_Q and the flow commands held, $P_b = 70\,\mathrm{kPa}$, $Q_{\max} = 1500\,\mathrm{kW}$, $c_h = 0.05$ and $c_s = 50$ are costs per step at full duty and full spray, ΔQ_j and ΔW_j are block moves measured first from the applied inputs, $\lambda = 10$, and $J_s = 1000\,(s_P/P_b)^2 + 10\,(s_L/5\,\%)^2$ prices the slacks of the soft bands. A step at full spray thus costs a thousand times one at full heat. The duty is hard within 0 to 1500 kW (zero while the heaters are blocked) and the spray within 0.25 to 28 kg/s, from the declared ratings, while the pressure is soft within 13.20 to 16.00 MPa, 100 kPa inside the trip and the PORV, and the level within 20 to 92 %, with limits read from the run. There is no terminal weight and no load preview. A dual active-set solver [22] solves the 24 decisions. The costs and λ are the largest within 10 % of the least IAE, 471.4 MPa s, among 90 settings on scenario 12 that spend no more spray and heat than the overrides. A plan that is not optimal, finite and in range is rejected, and PIC-101, which tracks the applied inputs through the inverse of (3.13), takes over without a bump until the next valid plan. The controller reads only the transmitters, setpoint, flow commands, load, interlock state and declared parameters, so its model shares the plant's equations but not its run, and a parameter the learner edits is a mismatch for the estimator.

3.6 Two Loops on One Vessel

3.6.1 The Interaction Is Not Symmetric

Scenario 6 closes both loops against fixed setpoints and asks for ten points more level. Filling the pressurizer *raises* the pressure, and the pressure controller answers with spray. Over the 440 s fill the

level reaches 69.55 %, the bubble is squeezed from 20.40 to 15.53 m^3, the heater duty falls below 5 kW within 70 s, the spray runs up to 6.91 kg/s, and 496 kg of the original 2097 kg of steam is condensed away. Returning the setpoint reverses it badly, because charging shuts but the only way out of the vessel is a letdown orifice at 4.5 kg/s, so the drain takes half as long again as the fill (654 s against 436 s to come within half a point) and the pressure falls to 15.2613 MPa with the heaters at their full 1500 kW. Moving a level ten points and putting it back costs 1969 kg of spray water and 196 kWh of heat from the step at $t = 60$ s to $t = 1000$ s.

Scenario 3 makes the opposite measurement. A step to 60 % spray demand delivers 7.99 kg/s into the steam space, condenses 1.95 kg/s of steam for 1.92 MW, and takes the pressure down at −3.14 kPa/s averaged over the 100 s step. It moves the level by −0.214 % over the same one hundred seconds. The spray draws its water from the cold leg, so the rigid liquid-full coolant system pulls the same mass back down the surge line and only the condensate stays. Charging is a level actuator with a large pressure side-effect, and spray is a pressure actuator the level can barely see. A relative gain array computed on those two gains identifies the pairing correctly [3], and it says nothing about the fact that both loops also share the surge line disturbance, which is where most of the difficulty actually lives.

Pause. Try Exercise 3.1 before reading on.

3.6.2 The Level Program and What It Does Not Cover

The industrial answer to the coupling is not a decoupler. It is a level setpoint programmed from turbine load, so that the target moves with the disturbance and the level controller has less to correct. The application programs 23 % at zero load to 60 % at full load, a slope of 0.370 % of level per percent of load. Converting the coolant expansion to pressurizer volume over the same span gives Equation (3.16).

$$\frac{\mathrm{d}L}{\mathrm{d}\mathrm{load}} = \frac{V_{RCS}\left[\rho_f(T_{avg,lo}) - \rho_f(T_{avg,hi})\right]}{\rho_{f,pz}\,V} \tag{3.16}$$

The coolant density is 738.38 kg/m^3 at 294 °C and 701.66 kg/m^3 at 309.5 °C, and the pressurizer water $\rho_{f,pz}$ is at 591.12 kg/m^3. With these values Equation (3.16) gives 11 238 kg, 19.01 m^3 and a slope of 0.3728 % per percent, which matches the programmed 0.370 % to 0.8 %. The setpoint is the plant's own thermal expansion written down rather than an operating preference.

The program is nonetheless not the whole level swing, and measuring it directly shows why. Settling the model at a series of loads with the pressure loop closed and charging matched to letdown, so that only the surge line moves the inventory, gives a natural level slope of 0.4931 % per percent between 100 and 95 % load. The excess over (3.16) is exact and has a name. Falling level grows the bubble at constant pressure, and the steam to fill it has to be boiled out of the pool. Between 100 and 95 % load the pool loses 743 kg, of which 614 kg is coolant expansion up the surge line and 129 kg becomes steam in the growing bubble, and the same decomposition holds to the kilogram at 70 % load with 4360 kg, 3602 kg and 758 kg. The consequence for the control system is measurable. At 70 % load the program asks for 48.90 % level while the process would settle at 45.54 % of its own accord, so the level controller still charges about 840 kg. The remaining roughly 176 kg of the pool-level gap is steam that condenses. That residue is small against the 11 % error a fixed setpoint carries, and it is not zero.

3.7 Results

The thirteen scenarios are ordered as a laboratory sequence. Scenarios 1 to 3 measure each actuator by hand and by a step test, scenarios 4 and 5 set the deadband and the valve trim, and scenario 6 closes the level loop. Scenario 7 is the load rejection, scenario 8 an outsurge, scenarios 9 and 10 two failures that hold every setpoint, scenario 11 adds the overrides, scenario 12 the predictive controller, and scenario 13 is the graded challenge. Each one loads an architecture, a tuning, a set of

disturbance positions and a schedule of timed events, and its briefing, completion rule and Reset work as Section 1.6.4 describes. Reset keeps the learner's own settings, which is how the comparisons in scenarios 1, 4, 5 and 11 are made. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt.

3.7.1 The Load Rejection

Scenario 7 is the transient the application exists for. The turbine sheds thirty percent of its load in a step at $t = 30$ s, and two transients follow in opposite directions from the single disturbance.

The reactor cannot follow instantly, so T_{avg} climbs 0.674 K while the program is falling, 306 m^3 of coolant expand, and the surge line peaks at 34.72 kg/s four seconds after the step. The pressure reaches 15.7462 MPa at $t = 39.6$ s, a rise of +236 kPa in ten seconds, and the spray demand goes to 100 %, although the valve reaches only 96 % of its travel by $t = 44.5$ s, after the peak. Even full spray could not stop it. Full spray condenses 7.26 kg/s while the surge is compressing the bubble several times faster than that, and the remaining 354 kPa of margin to the relief valve is the plant's rather than the controller's.

> **TRY IT IN THE APP** **7 · The Load Rejection: An Insurge You Cannot Stop**
>
> Load scenario 7 and press Run. Watch the pressure after the turbine sheds thirty percent of its load at 30 s. The spray demand goes to 100 %, yet the pressure still peaks at 15.7462 MPa at 39.6 s. Keep watching: the reversed surge then takes the pressure down to 14.5028 MPa while every heater bank comes in.

Then rod control and the steam dump win, the dump being the subject of Chapter 5. T_{avg} settles at the 70 % program of 304.85 °C, 4.65 K below where it started, the coolant contracts, and the surge reverses to −100.7 kg/s. The level falls from 60 to 45.37 % and the pressure to 14.5028 MPa, a full megapascal below program. Every heater bank comes in and the pressure is back inside its 70 kPa band at $t = 208$ s. LIC-101 meanwhile charges at its 12 kg/s maximum from $t = 78.5$ to 297.5 s and is still at 8.1 kg/s at 700 s, and that 292.75 °C water compresses the bubble. The pressure overshoots by 66.8 kPa at $t = 270$ s, the output crosses into the spray half at $t = 221$ s, and from $t = 328$ to 700 s the heaters stay below 5 kW while the spray runs at 1.5 to 4.6 kg/s, condensing what the charging compresses. Ten minutes after the rejection the pressure is 10 kPa high at 15.5202 MPa. When the unit is reloaded the sequence runs backwards and the level peaks at 70.48 %, ten and a half points above a setpoint that never moved, and the second outsurge takes the pressure to 14.50 MPa at $t = 911$ s. Half an hour into the run the loop is recovering on heat alone, as shown in Figure 3.4.

3.7.2 The Peak Is Not the Part Tuning Wins

Table 3.4 puts the three closed-loop architectures on the same disturbance at the same tuning, with the architecture chosen before the run. The peak pressure moves by 20 kPa across all three, which is 8.3 % of the 236 kPa rise, while the spray water differs by 1107 kg and the time with the level more than five points off setpoint differs by a factor of 26. Figure 3.3 draws Table 3.4 as bars.

Tuning moves the peak more than architecture does, and it moves it far less than intuition suggests. Sweeping the pressure loop over the plausible box $K_c \in \{400, 800, 1600, 2500\}$, $DB \in \{0, 10, 20\}$, $n \in \{1, 2.5\}$ and $\tau_D \in \{0, 12, 40\}$ with the level loop closed against a fixed setpoint gives a lowest peak of 15.6427 MPa and a highest of 15.8180 MPa, a spread of 175 kPa on peaks that are 133 to 308 kPa above program. The lowest peak comes from $K_c = 2500$ with no deadband, linear trim and $\tau_D = 40$ s, which is the one combination a designer would refuse, because that configuration then uses 11 529 % of output travel to $t = 1500$ s against 1290 % for the shipped tuning, spends 731 s of the run with both actuators in service against 247 s, and under the 4 kPa of transmitter noise of scenario 4 fights for all 600 s of the measurement window and moves the output 15 295 % against 4498 % for the shipped tuning, which never fights. Buying 104 kPa of peak with a factor of 8.9 in valve travel is not a trade a plant makes.

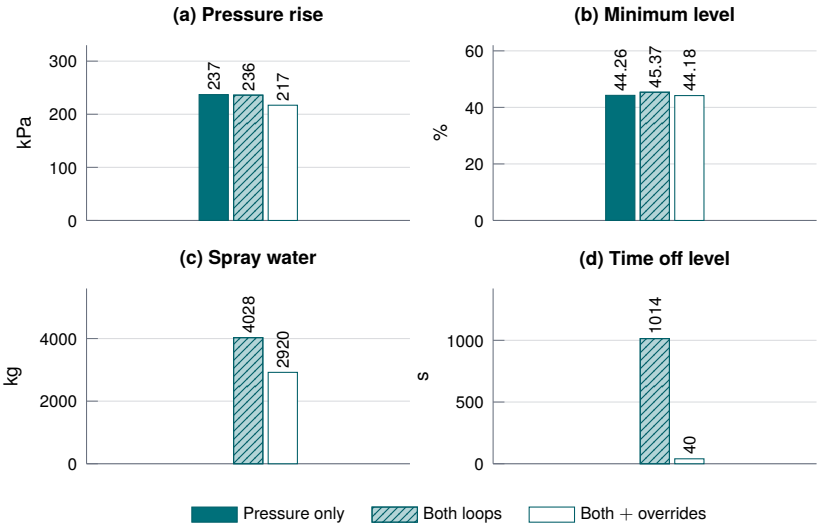

Figure 3.3 The −30 % load rejection of Table 3.4 under three architectures: (a) pressure rise, (b) minimum level, (c) spray water and (d) time with the level more than five points off setpoint. Solid: Pressure only. Hatched: Both loops. White: Both + overrides. Pressure only leaves the level in manual, so the table gives it no spray or level-time entry. Notice that the rise and the minimum level barely move, while the spray water and the time off level separate the architectures.

Table 3.4 The Same −30 % Load Rejection Under Three Architectures at the Shipped Tuning (K_c = 800 %/MPa, τ_I = 240 s, τ_D = 12 s, DB = 10 %, n = 2.5). Peaks Are Read Over the First 300 s and Totals at t = 1500 s.

	Pressure only	Both loops	Both + overrides
Peak P (MPa)	15.7466	15.7462	15.7269
Rise (kPa)	237	236	217
Minimum P (MPa)	14.4452	14.5028	14.4325
Minimum level (%)	44.26	45.37	44.18
Reload level peak (%)	—	70.48	63.04
P at 1500 s (MPa)	—	14.8501	15.2946
Spray water (kg)	—	4028	2920
Time off level (s)	—	1014	40

The same result appears in the graded challenge, where the disturbance is larger and the comparison is between two complete configurations rather than two gains. The challenge passes a run with no PORV lift at 16.10 MPa, the level always above 17 %, and the pressure within 100 kPa of 15.51 MPa over the last third of the run ($t \geq 1386$ s). Table 3.5 reports it. The shipped configuration and the suggested one peak within 26 kPa of each other, and neither lifts the relief valve. Everything after the peak separates them completely. The shipped configuration's worst error over the last third is twelve times the tolerance, although the level stays above 34.15 %. The suggested one is inside the tolerance by 73 kPa, with 75 kPa of PORV margin and the level above 33.2 %. No single change passes. Both loops with the overrides and K_c = 800 %/MPa, with nothing else changed, pass at 35 kPa. Tuning alone does not reach the target. The suggested split-range settings applied to the shipped two-loop architecture still end 320 kPa low, and 306 kPa low with the integral time halved to 120 s. The same settings with the override architecture end 1 kPa from program and pass even with the shipped level tuning, so the structure has to change before the gains do.

Table 3.5 Graded Challenge. A −55 % Load Rejection at $t = 60$ s, Read Over 2100 s, With the Last Third Judged From $t = 1386$ s as in the Application. The Shipped Configuration Has No Deadband, No Derivative, Linear Trim, a Gain Reduced to Stop Cycling and a Fixed Level Setpoint.

	Shipped	Suggested
Architecture	both loops	both + overrides
K_c (%/MPa), τ_I (s)	200, 500	800, 240
DB (%), n, τ_D (s)	0, 1, 0	10, 2.5, 12
Level K_c (kg/(s %)), τ_I (s)	2.5, 120	1.0, 400
Peak P (MPa)	16.0505	16.0253
PORV lifts	no	no
Minimum level (%)	34.15	33.20
Worst $\lvert P - P_{SP}\rvert$, last third (kPa)	1203	27
$\lvert P - P_{SP}\rvert$ at the end (kPa)	360	2
Spray water (kg)	2818	816
Heater energy (kWh)	519	524

✦ DID YOU KNOW? The Configuration That Passes Sprays Less

From the simulator: in the graded challenge the shipped configuration sprays 2818 kg of water and fails, while the suggested one sprays 816 kg and passes, for almost the same heater energy, 519 against 524 kWh (Table 3.5). The structure that holds the pressure also spends less water doing it.

The physical reason the peak resists tuning is in (3.2) and (3.4). The surge compresses the bubble at a rate set by the coolant expansion, and the only mechanism the controller has is condensation at a rate capped by the spray capacity and the cold leg enthalpy deficit. The one parameter that moves the peak at no cost is not a controller setting. Removing the coolant compressibility of (3.12) by setting κV to zero raises the rise from 236 kPa to 314 kPa, so the compliance of the reactor coolant itself takes 24.8 % off the peak. That is most of what the most aggressive tuning in the sweep buys, and it costs no valve travel at all.

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.7.3 Depressurization Is Self-Limiting

Scenario 8 opens a second letdown orifice with the level loop out of service, taking letdown from 4.5 to 14 kg/s against charging held at 4.5 kg/s. The level falls fifteen points from 60 to 45.04 % by the time the orifice shuts at $t = 430$ s, and the pressure falls only 305 kPa to 15.2050 MPa. The pool holds it up. As the bubble expands the saturation temperature falls, eighteen tonnes of water that were exactly at saturation become superheated, and (3.3) flashes them. Measured, the pool reaches 0.190 K above saturation at $t = 90$ s and flashes 1.95 kg/s, which is 1.90 MW and more than a quarter of what full spray can condense, and it is still flashing 1.84 kg/s when the orifice shuts. The heaters go to 1500 kW and stay there while the pressure falls anyway, which is the plainest available demonstration that heaters hold a steady state and do not answer transients.

3.7.4 Two Failures That Hold Every Setpoint

Scenario 9 removes two backup heater banks on a bus transfer and nothing happens. At 9.45 % demand the modulating bank alone is carrying the plant and no backup bank is in service, so the duty, the controller output, the pressure and every trend are exactly what they were. The failure becomes visible only when an outsurge asks for the heat, at which point the duty saturates at 900 kW instead of 1500 kW and the pressure sags 149 kPa to 15.3607 MPa. Scenario 10 sticks a spray valve 45 % open, which the equal-percentage trim makes 4.02 kg/s of cold water condensing 1.005 MW continuously. The loop holds the pressure to within 4 kPa of setpoint by running the heaters at 1500 kW with every

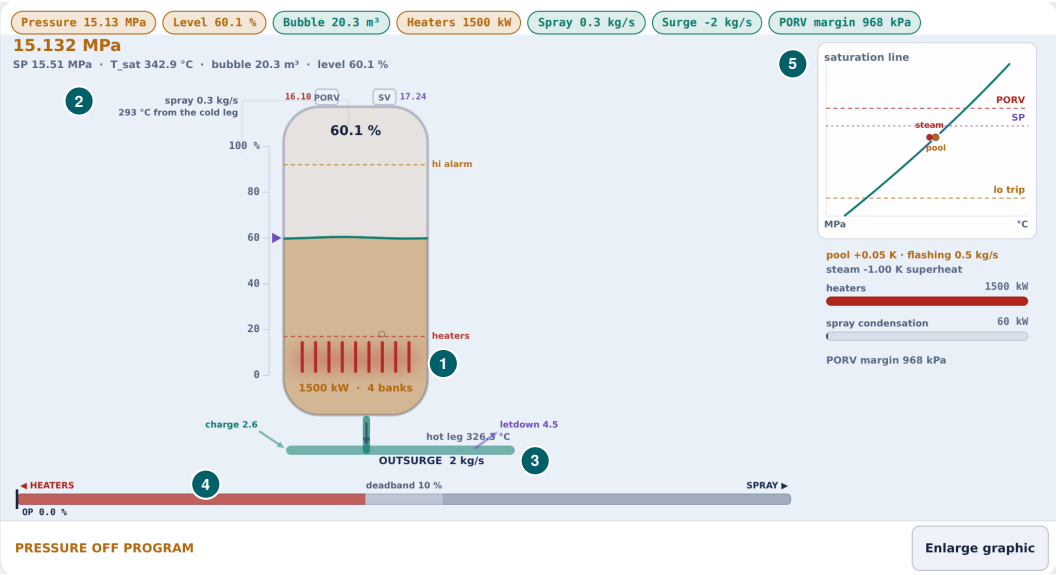

Figure 3.4 The load rejection of scenario 7 at $t = 1800$ s, during the heater-limited recovery from the outsurge that follows the reload. Pressure is at 15.132 MPa against a setpoint of 15.51 MPa and the banner reads PRESSURE OFF PROGRAM, the level is at 60.1 % against its fixed setpoint of 60.0 %, and the surge line is still running 2.1 kg/s outward. Every heater bank is in at 1500 kW while spray condensation supplies only 60 kW from the bypass, so the split-range bar under the vessel shows the output pinned at $OP = 0$ at the far end of the heater band with the whole deadband and the entire spray band unused. The saturation-line inset places the pool and the steam points below the setpoint line, and the readout under it reports 968 kPa of margin to the relief valve. The view is cut to the plant panel of the application. Callouts: (1) the heater bundle with every bank in, (2) the spray line from the cold leg, (3) the surge line running outward, (4) the split-range bar with the output pinned at zero, (5) the saturation-line inset.

bank in, around the clock. On the pressure trace everything is fine. What has happened is that the loop has spent all of its upward authority to stand still, and when a −15 % load change arrives at $t = 400$ s the pressure falls to 14.9779 MPa with nothing left to add. Two scorecard numbers would have caught it from the first minute, because the fighting counter reaches 369 s by $t = 400$ s and 869 s by $t = 900$ s, and the heater energy reaches 342 kWh over the run against the 35.5 kWh that 141.8 kW of standing duty would have consumed in the same 900 s. Neither is an alarm on a real plant and both are free to compute.

> **TRY IT IN THE APP** **10 · A Spray Valve That Will Not Shut**
>
> Load scenario 10 and press Run. After the spray valve sticks open at 30 s, the pressure stays within 4 kPa of setpoint. Watch the heater duty and the fighting counter instead: the heaters run at 1500 kW and the counter reaches 369 s by 400 s. Then the load change arrives and the pressure falls to 14.9779 MPa.

3.7.5 Comparison With the Predictive Controller

Table 3.6 runs scenario 7 (scenarios 7, 11 and 12 differ only in the architecture), scenario 8 and the challenge at its suggested settings with the same events, seed and instruments and the MPC as in Section 3.5.4.

Table 3.6 Split Range Against the Predictive Controller

Case	Structure	Peak P (MPa)	IAE (MPa s)	Off band (s)	Spray (kg)	Heat (kWh)	Fight (s)
Insurge[a]	Both loops	15.8653	589.8	891	4028	344.4	247
1500 s	Both + overrides	15.7394	441.1	949	2920	404.1	244
	MPC	15.8344	505.9	958	2835	371.0	133
Outsurge	Pressure only	15.5582	103.7	481	225	242.2	0
900 s	Both loops[b]	15.5681	26.7	90	1681	117.6	88
	MPC	15.5333	17.0	77	1670	111.8	77
Challenge	Both + overrides	16.0253	537.5	1077	816	523.8	24
2100 s	MPC	16.0609	535.5	1080	811	523.7	24

Totals and peaks are over the whole run. Off band is the time more than 70 kPa from setpoint, and Fight the time with more than 5 kW of heat and more than 0.375 kg/s of spray delivered together. [a]The peaks fall on the reload at 748 to 752 s, and the rejection's own peaks are 15.7462, 15.7269 and 15.7579 MPa. [b]Both + overrides is identical.

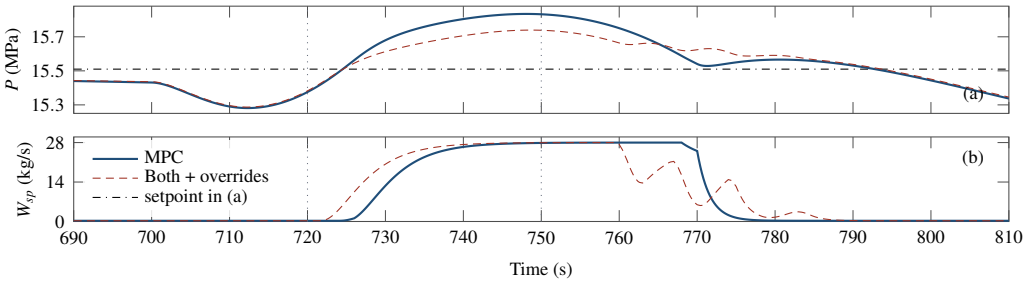

Figure 3.5 The reload of scenario 12 under the predictive controller and of scenario 11 under the split range with overrides, from the same initial state and the same load steps at 700 and 900 s. The dotted lines bound the decisive interval from 720 to 750 s. The predictive controller is the solid line, the split range with overrides the dashed line and the setpoint the dash-dotted line. (a) True pressure and setpoint. (b) Delivered spray flow.

On the outsurge, on the flashing side where the model was linearized, the MPC is ahead in every column. On the insurge it uses 2.9 % less spray and 8.2 % less heat than the overrides for a 14.7 % higher IAE. The loss is the reload's. The rejection peaks are 31 kPa apart near 39.5 s, the reload peaks 95 kPa apart at 748 s, as shown in Figure 3.5. The derivative of PIC-101 has the spray demand at 72 % at 722 s, with the pressure 74 kPa low, and full from 722.1 s, while the plans at 722 and 724 s keep the bypass, predicting 15.71 and 15.86 MPa at the horizon end, and the MPC's demand is full from 726.1 s. In diagnostic copies, not shipped, deciding every 0.5 s lowers the reload peak by 43 kPa, a subcooled-side model by 53 kPa, both by 96 kPa, and a zero spray cost by only 9 kPa. The spray cost sets the trade against pressure. With c_s at 0, 50 and 500, scenario 12 spends 3758, 2835 and 2457 kg of spray for an IAE of 265.3, 505.9 and 532.1 MPa s, while its reload peak stays within 18 kPa. Free spray buys a much lower IAE and leaves the reload peak nearly where it was, since the spray is already full on the reload under every price.

Neither structure ever commands heat and spray at once, with 0.0 s of the Fight condition on the commands in every case. Fight itself counts delivered overlap, and 116 of the MPC's 133 s and 225 of the split range's 244 s are heater duty decaying through the sheath lag after its command reached zero. The quadratic program adds a price on each actuator and pays it in pressure, and no setting in the sweep within the overrides' spray and heat reached their 441.1 MPa s.

On the challenge with its suggested settings the two are level on IAE, spray and heat and the MPC peaks 36 kPa higher. No plan was rejected in any scenario over 2100 s. The MPC solved 750, 450 and 1050 plans, and in Chrome a solve took a mean of 647 to 661 µs in three runs, at most 9.9 ms, about 0.03 % of one core, as wall clock on a 20-core desktop at a load average of 2 to 3.

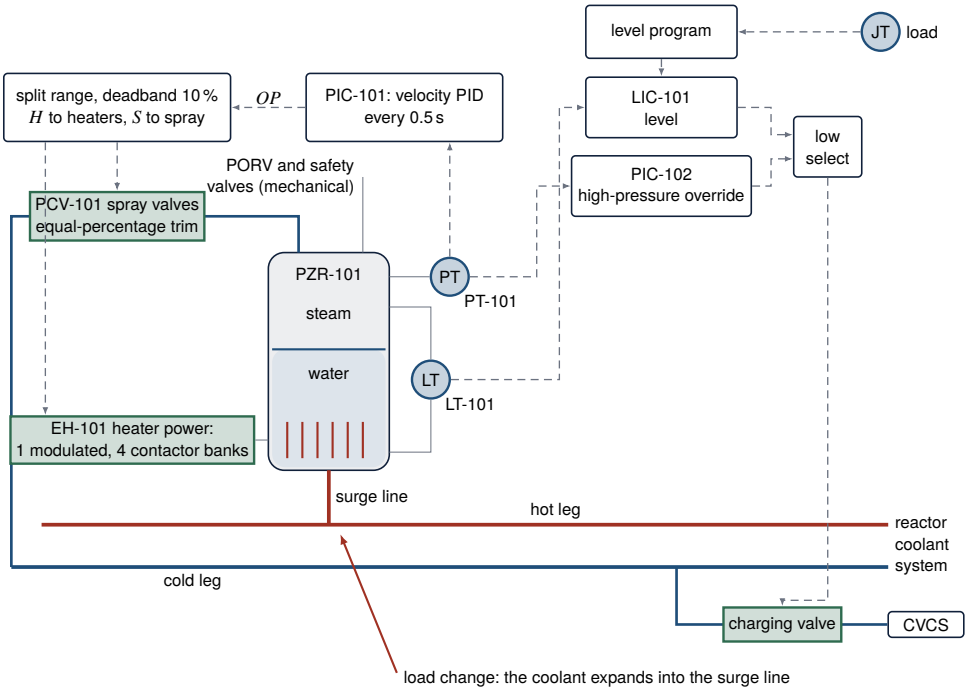

Figure 3.6 The pressure and level loops of Figure 3.1 with the hardware of a plant. A pressure transmitter (PT) on the steam space and a differential pressure level transmitter (LT) across the vessel feed the controllers, and a load measurement (JT) programs the level setpoint. PIC-101 runs every 0.5 s and its one output is split across the deadband between the heater power equipment and the spray valves, while LIC-101 and the override PIC-102 share the charging valve through the low selector. The relief valves act on their own. Thick solid lines carry hot-leg water (upper) and cold-leg water (lower and the spray line), and thin dashed arrows signals.

3.8 Technology in Practice

The tags of Figure 3.1 are, in a plant, a pressure transmitter and a level transmitter on the vessel, a heater bundle with its power equipment, spray valves on a line from the cold leg and the charging valve of the chemical and volume control system. Figure 3.6 draws the two loops with that hardware in place of the schematic's symbols.

What Measures It

The pressure of the whole primary circuit is read at the pressurizer, where the bubble sets it. Pressure transducers convert the force on a sensing element to a 4 to 20 mA signal or, in newer smart transmitters, to a digital one [23]. Level in a nuclear plant is most often measured as a differential pressure between taps at the top and the bottom of the vessel, against a filled reference leg [23]. The chapter carries both measurements as first-order lags, 1.0 s on PT-101 and 4.0 s on LT-101.

> **TECHNOLOGY** **Sensor: The Pressure Transmitter**
>
> A common pressure cell holds a taut diaphragm between two fixed plates, with a silicone fill fluid carrying the pressure in, so the deflection reads as a change of capacitance [24]. The transmitter sends a 4 to 20 mA signal, or a digital one from a smart transmitter, and the live zero at 4 mA reveals an open circuit [23]. Too much damping makes the transmitter report a slower process than the real one, and the instability that follows does not appear on the trend [24]. The chapter's PT-101 is a 1.0 s lag with optional noise.

What Moves It

Two final elements share the output of PIC-101. The heater bundle has one modulated bank and four banks switched by contactors, the two kinds of heater actuation found in plants, a regulated current and plain on and off [23]. The spray valves pass cold-leg water through a characterized trim, the equal-percentage form of Equation (3.5) with $n = 2.5$. Charging is the third element, shared by LIC-101 and PIC-102 through the low selector. The relief valves above them are mechanical and act on the true pressure, outside the control system, as the limiting and protection layers of Section 1.4.5.

> **TECHNOLOGY** **Final Element: Heaters and Spray on One Output**
>
> Valves that follow the output of one controller are called split-ranged. In the exclusive form both are shut at mid output and each opens toward its own end, so the fluid never passes through both at once [24]. An equal-percentage trim passes no more flow wide open than a linear one, but makes the installed flow closer to linear in stem position. A valve with stem friction moves in steps and ignores small reversals of its signal, and no tuning removes that [24]. Electric heaters are mostly switched on and off, though some are modulated by regulating their current [23]. PIC-101 drives both kinds of heater and the spray valve PCV-101.

Where the Controller Runs

Control loops of this kind run in field controllers with cycle times from milliseconds to minutes, under a supervisory layer that sends them setpoints [25]. The chapter's three controllers execute every 0.5 s above a 0.05 s physics step. In a function-block controller the split range is a block of its own: one PID block passes its output to a split-range block that drives two outputs, one for each final element [24]. Doing the split in the controller keeps each output a standard signal, and the deadband and the two bands of Equation (3.13) become settings of that block. The selector of Equation (3.14), the heater interlock and the level program of Equation (3.16) complete the chain.

> **IN PRACTICE** **Count the Thermal Cycles**
>
> A plant is designed for a set number of load cycles over its life. Under load following the pressurizer sees more thermal cycling, and the parts that suffer are the surge and spray lines, their nozzles and the vessel welds [26]. Plants that cycle therefore count the cycles on the surge line and monitor fatigue by measuring and recording wall temperatures against the design assumptions. Some change the pressurizer level and pressure control so that it makes fewer setpoint moves [26]. The fighting counter and the spray total on the chapter's scorecard are a first version of that count.

What Is Modernizing It

Three developments touch this loop. Smart transmitters diagnose themselves and report their own health [23], and on-line monitoring can show that a sensor has not drifted since its last calibration, so the calibration interval can be extended and the faults that maintenance itself introduces are avoided [23]. Continuous fatigue monitoring records wall temperatures and compares the observed load cycles with those the design assumed [26]. Plants that raise their power rating change the pressurizer pressure control to give finer control under reduced operating margins [23], which is where the deadband and trim choices of Section 3.4 return.

3.9 Limitations

The model is not a safety analysis and cannot support a claim about peak pressure, surge line loads or the licensing basis of any plant. Two lumped regions with judged interface coefficients cannot produce the stratified pool that a real in-surge creates, and the surge line thermal loads that follow are outside the model entirely [7, 11]. The spray is one efficiency number rather than a droplet field, so the loss of spray effectiveness as the bubble shrinks does not appear [12]. The reactor coolant system is a single node with one self-regulation coefficient standing in for both steam generator heat

transfer and the steam dump, which is enough to generate a realistic insurge and not enough to say anything about secondary pressure, steam generator level or reactivity feedback. The relief valves lift and reseat cleanly, so the failure that dominates the historical record is absent by design [13]. Charging, letdown and spray are first-order lags on commanded flow with no pump curves, no seal injection, no boron and no volume control tank, so the application can show what the continuous bypass spray costs in heater kilowatts but not why it exists. Reactor trips are annunciated rather than simulated. Even the repeatability is false, because a fixed integration step and a fixed scan give the same trajectory on every machine, and a plant control system is not that reproducible. What survived the simplification is what a controller has to be designed around, namely the 9.35 actuator asymmetry, the deadband, the installed characteristic, the inventory coupling, the selector and the interlock. Section 3.3.8 states where each assumption breaks down.

The predictive controller's model shares the plant's equations at the declared parameter values, so the comparison tests no structural model error, only the linearization side (79.5 kPa against 300.8 kPa one minute into an insurge) and the disturbance held over the horizon, and an edited plant parameter becomes a mismatch for the estimator. The held surge also puts the predicted pressure at the horizon end between 10.55 and 18.68 MPa on scenario 12 while the plant stays within 14.42 to 15.83 MPa.

OPEN CHALLENGE **Following Load Without Fatiguing the Surge Line**

Flexible operation aggravates cyclic loads, and a small change in operation can have an abrupt and large effect on fatigue. Thermal stratification in the surge line and actuations of the spray line are the examples given for the pressurizer [26]. These loads arise where hot and cold water meet in one pipe, which the chapter's two lumped regions cannot represent. Progress would be a pressure and level controller that knows the fatigue each spray cycle and each surge costs and trades it against pressure error, as the predictive controller of Section 3.5.4 already trades spray water against error.

3.10 Discussion

The result that generalizes beyond the pressurizer is the separation between what a control system can change and what it cannot. Table 3.4 shows a peak that moves 20 kPa across three architectures and a recovery that moves 445 kPa, so the peak belongs to the plant. Its value is set by the coolant expansion rate, by the bubble stiffness of 0.195 MPa per kg/m^3 under adiabatic compression, and by the 7.00 MW ceiling on spray condensation, and no arrangement of the same equipment beats those three numbers. The relief valves exist for that reason, and a control narrative promising to hold the pressure below the relief setpoint through every transient promises something the equipment does not support. The same reasoning applies wherever a fast disturbance outruns the strongest available actuator, which includes compressor surge (Volume 4, Chapter 2), batch reactor runaway and grid frequency after a large unit trip.

The three pieces of logic that do change the outcome are older and cheaper than multivariable control, and all three are auditable. Setpoint programming removes most of the interaction before any controller acts, because it makes the target a known function of the measured load. The low selector resolves a conflict over one valve with an explicit priority, and the interlock removes an actuator from the loop that would otherwise destroy it. Formal treatments of split range and of selector-based constraint switching now place these structures on the same footing as the optimization they approximate [1, 2, 4], which matters where a licensing basis is written around what the control system will do.

The model of Section 3.3 also supports methods beyond the implemented loops. The predictive controller of Section 3.5.4 plans the heaters and the spray with the relief setpoint as a soft constraint, and a third input on charging would replace the level loop and the selector with one three-input problem. State estimation on the same equations could infer the pool subcooling, which is not measured on a plant and which decides whether the next transient flashes or condenses. Coordinated

reactor and pressurizer control for load following is an active subject [14, 15], and the surge line closure of (3.10) couples the two problems.

3.11 Conclusion

This chapter has treated a pressurized water reactor pressurizer as an interactive laboratory for split-range control across actuators that are not the same size. The simulator integrates thirteen states, including separate steam and liquid temperatures so that the bubble superheats under compression at 0.195 MPa per kg/m^3 against 0.094 MPa per kg/m^3 on the saturation line, and so that the pool flashes at 1.95 kg/s on 0.190 K of superheat during an outsurge. Measured from the design state, full heaters move the pressure at +1.232 kPa/s and full spray at −11.513 kPa/s, a ratio of 9.35 that comes from a 7.00 MW condensing load against 1.5 MW of resistance heating and from a thermal path that holds 139 220 kJ/K of pool in front of the heaters and nothing in front of the spray. The deadband and the installed valve characteristic are the two design variables that make one gain serve both halves, because zero deadband costs 139.5 s of simultaneous operation in every ten minutes under 4 kPa of noise while ten points costs none, and an equal-percentage trim carries 11.2 times less flow above the bypass at 20 % travel with identical capacity wide open. On a −30 % load rejection the peak pressure moves only 20 kPa across three architectures at one tuning and 175 kPa across a wide tuning box, while the reactor coolant compressibility alone removes 24.8 % of it, so the peak is a plant property and the architectures separate on the recovery instead, at 15.2946 MPa against 14.8501 MPa after 1500 s and 2920 kg of spray water against 4028 kg. The anti-windup tracking time is a design variable on this loop, because a 30 s tracking time winds the pressure integral against a proportional term eight times the output range and sprays with the pressure still 793 kPa low, while tracking at the integral time keeps the spray shut until the pressure has recovered. Both pressure controllers retain this integral-time tracking. A predictive controller pricing the actuators spends 2.9 % less spray and 8.2 % less heat on the rejection and reload but peaks 95 kPa higher. Further work would add a stratified pool with a separate surge line node, a droplet-resolved spray model, a relief valve that can fail to reseat, and a steam generator secondary side so that the primary and secondary transients can be studied together. The lesson the app teaches by measurement is that two actuators on one output are two plants.

KEY POINTS

- Full spray moves the pressure 9.35 times faster than full heaters, and both hang on the two halves of one controller output.
- Zero deadband costs 139.5 s of simultaneous operation in ten minutes under 4 kPa of noise, while ten points cost none.
- On a −30 % load rejection the peak pressure moves only 20 kPa across three architectures, so the peak is a plant property.
- The architectures separate on the recovery, at 15.2946 MPa against 14.8501 MPa after 1500 s.
- Two actuators on one output are two plants.

Where else this applies. A split range across asymmetric actuators also sets heating against cooling on the jacketed reactor of Volume 8, Chapter 2 and fuel gas against a vent on the header of Volume 11, Chapter 14.

References

[1] A. Reyes-Lúa, C. Zotică, K. Forsman, and S. Skogestad, “Systematic design of split range controllers,” *IFAC-PapersOnLine*, vol. 52, no. 1, pp. 898–903, 2019.

[2] A. Reyes-Lúa and S. Skogestad, “Multi-input single-output control for extending the operating range: Generalized split range control using the baton strategy,” *Journal of Process Control*, vol. 91, pp. 1–11, 2020.

[3] E. Bristol, “On a new measure of interaction for multivariable process control,” *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[4] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[5] A. N. Nahavandi and S. Makkenchery, "An improved pressurizer model with bubble rise and condensate drop dynamics," *Nuclear Engineering and Design*, vol. 12, no. 2, pp. 135–147, 1970.

[6] R. C. Baron, "Digital model simulation of a nuclear pressurizer," *Nuclear Science and Engineering*, vol. 52, no. 3, pp. 283–291, 1973.

[7] S.-N. Kim and P. Griffith, "PWR pressurizer modeling," *Nuclear Engineering and Design*, vol. 102, no. 2, pp. 199–209, 1987.

[8] S. K. Moghanaki and M. Rahgoshay, "Simulation of two-region and four-region models for typical PWR pressurizer and benchmark obtained results using available results," *Annals of Nuclear Energy*, vol. 63, pp. 302–308, 2014.

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

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

[11] B. Tang and Y. Zhou, "Numerical investigation on turbulent penetration and thermal stratification for the in-surge case of the AP1000 pressurizer surge line," *Nuclear Engineering and Design*, vol. 378, p. 111176, 2021.

[12] J.-C. Jo, S.-K. Lee, and W.-K. Shin, "Heat transfer of a spray droplet in a PWR pressurizer," *KSME Journal*, vol. 5, no. 2, pp. 130–139, 1991.

[13] D. L. Phung, "Light water reactor safety before and after the three mile island accident," *Nuclear Science and Engineering*, vol. 90, no. 4, pp. 509–520, 1985.

[14] D. Bose, A. Hazra, S. Mukhopadhyay, and A. Gupta, "A co-ordinated control methodology for rapid load-following operation of a pressurized water reactor based small modular reactor," *Nuclear Engineering and Design*, vol. 367, p. 110748, 2020.

[15] H. K. Selim and N. M. El-Sahlamy, "Pressurizer system dynamic model for transient control in PWR," *Kerntechnik*, vol. 87, no. 6, pp. 672–682, 2022.

[16] 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.

[17] W. Wagner, J. R. Cooper, A. Dittmann, J. Kijima, H.-J. Kretzschmar, A. Kruse, R. Mareš, K. Oguchi, H. Sato, I. Stöcker, O. Šifner, Y. Takaishi, I. Tanishita, J. Trübenbach, and T. Willkommen, "The IAPWS industrial formulation 1997 for the thermodynamic properties of water and steam," *Journal of Engineering for Gas Turbines and Power*, vol. 122, no. 1, pp. 150–184, 2000.

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

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

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

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

[22] 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.

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

[24] 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.

[25] 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.

[26] 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.

3.12 Exercises

3.1 Concept check · 2 min

The spray valve opens to pass 8 kg/s of cold-leg water into the steam space, with charging equal to letdown and the average coolant temperature steady. Use the mass balances of Equation (3.1) and the surge closure of Equation (3.10). What does the spray do to the total mass of water and steam in the pressurizer?

(a) It rises by 8 kg/s, because all of the spray water enters the vessel

(b) It rises by the condensation rate W_{sc}, because only condensed steam is new water

(c) It stays the same, because the surge line returns the spray and condensation only moves mass

(d) It falls by the condensation rate W_{sc}, because condensed steam leaves the bubble and the vessel

3.2 Concept check · 1 min

In the outsurge that follows a load rejection the pressure is far below its setpoint and the level is falling toward 17 %. With BOTH PLUS OVERRIDES, which element keeps the pressure controller from adding heat, and why?

(a) The low selector, because it gives the heaters to whichever loop asks for less heat

(b) The high-pressure override, because it takes the charging valve whenever the pressure is off setpoint

(c) The deadband, because an output pinned at zero lies outside both actuator bands

(d) The heater interlock, because a heater energized while uncovered is destroyed after 30 s

3.3 Concept check · 2 min

Which statement about the peak pressure of the 30 % load rejection in scenario 7 does the chapter's evidence support?

(a) It is fixed mostly by the plant, since full spray condenses far less steam than the surge pushes in

(b) It is set mostly by the pressure gain, since the spray demand saturates only when the gain is low

(c) It is set mostly by the architecture, since the overrides cut charging before the pressure peaks

(d) It is set mostly by the deadband, since a narrow band opens the spray valves much sooner

3.4 Theory and calculation · 6 min

At the design point the saturated liquid enthalpy is 1632.8 kJ/kg, cold-leg water at 292.75 °C carries 1299.4 kJ/kg, the latent heat is 964.4 kJ/kg, the spray efficiency is $\eta_{sp} = 0.75$ and the wall loses 59.8 kW to containment. The heater bank totals 1500 kW.

(a) Find the steam that full spray (28 kg/s) condenses, its condensing duty and the ratio of that duty to the full heater bank. *Checkpoint:* 7.26 kg/s, 7.00 MW and 4.67.

(b) A spray valve sticks at 45 % travel with the $n = 2.5$ trim. Find the spray flow and the steady heater duty that holds the pressure, if all spray water leaves down the surge line at the saturated liquid enthalpy. *Checkpoint:* 4.02 kg/s and 1400 kW, leaving about 100 kW in reserve.

(c) The measured ratio of the two pressure rates in scenario 1 is 9.35. Find the factor that the duty ratio of (a) leaves unexplained, and name the quantity in the model that produces it. *Checkpoint:* About 2.0.

3.5 Theory and calculation · 7 min

The programmed level of the app runs from 23 % at zero load to 60 % at full load. The coolant volume outside the pressurizer is 306 m^3, its density is 738.38 kg/m^3 at 294 °C and 701.66 kg/m^3 at 309.5 °C, the pressurizer water is at 591.12 kg/m^3 and steam at 102.8 kg/m^3, and the vessel holds 51 m^3.

(a) Evaluate Equation (3.16): the mass expelled from zero to full load, its volume in the pressurizer and the level slope. Compare the slope with the program's. *Checkpoint:* 11 236 kg, 19.01 m^3 and 0.373 % per percent against 0.370 %.

(b) The measured natural slope is 0.4931 % per percent. From 100 to 95 % load, find how much the bubble grows and the steam mass that has to be boiled out of the pool to fill it. *Checkpoint:* 1.257 m^3 and 129 kg.

(c) At 70 % load the program asks for 48.90 % while the plant settles at 45.54 % on its own. With the pressure held, find the pool water the gap represents and the water the level controller must charge to close it. *Checkpoint:* 1013 kg of pool water, of which about 840 kg is charged.

3.6 Theory and calculation · 6 min

The resting heater demand is $H = 9.4544\,\%$. Use the split-range mapping of Equation (3.13).

(a) Find the resting controller output with the deadband at 10 % and at zero. *Checkpoint:* 40.75 and 45.27 %.

(b) At $DB = 10\,\%$, find the outputs at which the first and the fourth backup banks pick up (heater demands of 24 and 78 %). *Checkpoint:* 34.2 and 9.9 %.

(c) Scenario 4 uses $K_c = 1500\,\%/\mathrm{MPa}$ and 4 kPa (one standard deviation) of transmitter noise. Find the output noise from the proportional term, and the distance from the resting output to the nearest edge of the spray band in units of that noise, at both deadbands. *Checkpoint:* 6 points, then 0.79 and 2.38 standard deviations.

3.7 App challenge · 7 min

Load scenario 1, Two actuators, one output, and they are not the same size (MANUAL), at the default seed. Set Heater demand to 100 % before Run, run to 50 s and export the CSV file. Press the reload button, set Spray demand to 100 % before Run, run to 50 s and export again. Both files have the time column Time (s).

(a) From the true pressure, column `p_MPa`, find each average rate between the rows at 10 and 50 s, and their ratio. Accept 0.005 kPa/s on the heater rate and 0.02 kPa/s on the spray rate. *Checkpoint:* +1.243 kPa/s, −11.450 kPa/s and 9.21.

(b) In the heater file, read `heater_kW` and `evap_kgs` at 50 s and find the share of the duty that leaves as steam, with a latent heat of 964.4 kJ/kg. *Checkpoint:* 1388.5 kW, 0.155 kg/s and about 10.7 %.

Explain why the ratio differs from the chapter's 9.35.

3.8 App challenge · 8 min

Load scenario 4, Split range with no deadband. It runs Pressure only at $K_c = 1500\,\%/\mathrm{MPa}$ with 4 kPa of transmitter noise, and steps the setpoint +50 kPa at 60 s. Use the default seed and speed 10, and run each case to 800 s. The exported file then has the time column Time (min), and the rows at 200 and 800 s read 3.33333 and 13.3333.

(a) Run the scenario as shipped. From column `metric_fighting_s`, find the seconds of fighting between 200 and 800 s, and read the Heater energy and Spray water tiles at 800 s. *Checkpoint:* 139.5 s, 57.6 kWh and 340 kg.

(b) Press Reset, set Split-range deadband to 10 % and run again. Read the same three values. *Checkpoint:* 0.0 s, 49.6 kWh and 256 kg.

Accept 0.5 s, 0.1 kWh and 1 kg. Find the increase in heater energy and in spray water that the zero deadband costs.

3.9 App challenge · 10 min

Load the challenge, Catch a fifty-five percent load rejection, at the default seed and speed 10. Each run covers the full 2100 s.

(a) Change only Control architecture to Both + overrides and run. Read the challenge result: the worst pressure error over the last third, whether the PORV lifted, and the lowest level. *Checkpoint:* 159 kPa, no lift, 33.2 %: not passed.

(b) Set Pressure loop gain to 800 %/MPa, press Reset and run again. Read the same three values and the Spray water tile. *Checkpoint:* 35 kPa, no lift, 33.2 %: passed, with 1000 kg of spray.

Accept 1 kPa on the errors. Compare both runs with the shipped configuration in Table 3.5, and say what each change did.

3.10 Application area · 4 min

A plant engineer prices the deadband from Table 3.3, which compares zero and ten points over 600 s at $K_c = 1500\,\%/\mathrm{MPa}$ with 4 kPa of noise.

(a) If the differences in heater energy and spray water between zero and ten points persisted around the clock, find the extra heater energy and the extra spray water in a year of 8760 h. *Checkpoint:* About 298 MW h and 3050 t.

(b) Name the consequence the chapter ranks above the energy, and say why no pressure trend shows it.

3.11 Application area · 4 min

A relief valve that lifts and does not reseat began the accident at Three Mile Island. The two PORVs pass 53 kg/s of steam, and the bubble holds 2097 kg at the design point.

(a) Find how long the PORVs would take to pass the bubble's whole mass if nothing replaced it. *Checkpoint:* About 40 s.

(b) Use Equation (3.3) with $k_f = 10\,000\,\mathrm{kW/K}$ and $h_{fg} = 964.4\,\mathrm{kJ/kg}$ to find the pool superheat at which flashing would replace the relief flow. *Checkpoint:* About 5.1 K.

(c) Name two limitations of the model that keep the app from following such an event to its end.

3.12 Review the agent · 5 min

An AI assistant reviews scenario 7 and writes: "The load rejection peaks at 15.7462 MPa, only 354 kPa below the PORV. The chapter's tuning sweep finds a lower peak of 15.6427 MPa with $K_c = 2500\,\%/\mathrm{MPa}$, no deadband, linear trim and $\tau_D = 40\,\mathrm{s}$. Adopt these settings as the standard tuning. The deadband and the equal-percentage trim only matter under transmitter noise, and a plant transmitter is quiet enough to ignore it." Find the flaw and say how you would check it.

CHAPTER 4

Three-Element Drum Level Simulator for an Inverse Response That Moves With Load

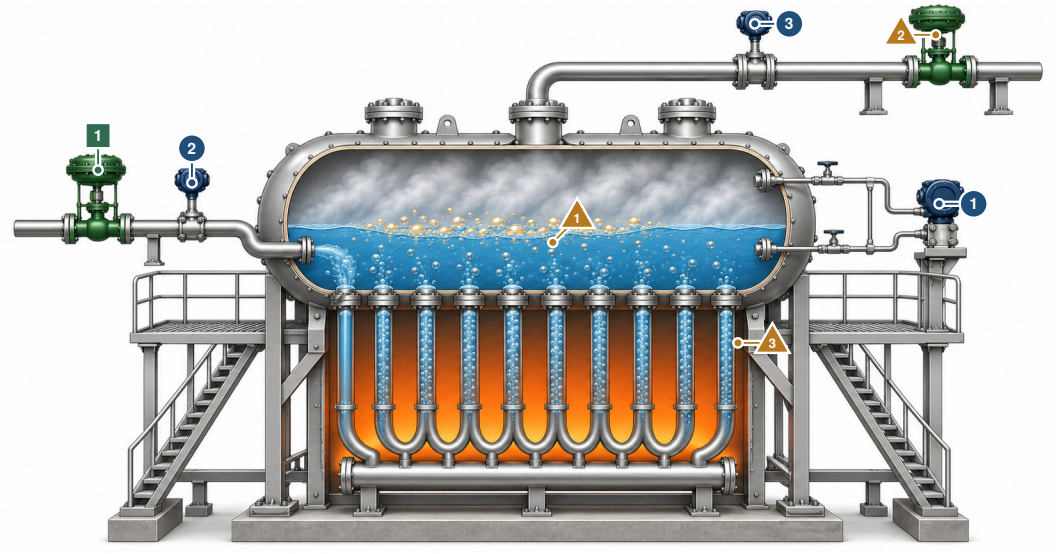

The controller moves 1 feedwater valve on the inlet line

The controller reads 1 drum level, from a pressure difference between a tap in the steam and a tap in the water
2 feedwater flow 3 steam flow to the turbine

What disturbs it 1 steam bubbles under the surface, which swell and collapse
2 turbine governor, opening with the steam demand 3 heat from the firing under the tubes

Notice the bubbles below the water line: they hold up the surface, so a load change or a burst of cold feedwater first moves the measured level the wrong way. Illustration, not to scale.

IN PLAIN TERMS Keeping Water in a Boiling Drum

The main idea. A boiler drum must keep its water level in a narrow band while steam leaves and new water comes in. The gauge sees the water and also the steam bubbles hidden under the surface.

The challenge. Cold water added to the drum first collapses bubbles, so the level drops before it rises. A controller that watches only the level is fooled, and it is fooled for longer when the boiler runs at low output.

Why it matters. Too little water harms the boiler tubes and too much sends water to the turbine, so either limit stops the plant.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 4.1 and 4.2, the fact boxes of Sections 4.3 and 4.6, the try-it boxes of Sections 4.4 and 4.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	25 min	The first page, Section 4.6, Section 4.8, Section 4.9 and the open challenges

The control question A drum level is the water inventory plus the steam held below its surface, and the two terms answer a feedwater or load change in opposite directions. What does the wrong-way response that this creates cost a level loop as the load falls, and what does each added measurement of the single-element, two-element and three-element ladder buy against it?

Learning objectives

- Compute the steam that a kilogram of subcooled feedwater condenses, and the void contribution to the level at a given load, from the balances and the property fits.
- Distinguish the depth of the inverse response, which scales with the step, from its duration, which moves with load.
- Explain why a steam-flow feedforward deepens the dip of a load rejection while it removes the overshoot, and why it is exact only when stated in flow.
- Show on a pump trip why an inner feedwater flow loop covers faults that no level tuning reaches.
- Evaluate the level error a differential-pressure cell ranged at the wrong pressure leaves behind a loop that holds its indication on setpoint.

Prerequisites Sections 1.4.4 (shrink, swell and the moving zero), 1.6.2 (velocity form) and 1.6.4 (how an exercise names a run). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/steam-generator/

See also Volume 3, Chapter 2 for matching loop structure to where a disturbance enters, Volume 3, Chapter 3 for the inner-loop speed a cascade needs and Volume 3, Chapter 9 for a right-half-plane zero that moves with load. Volume 1, Chapter 6 for a vessel whose level swells while it loses water, and Volume 1, Chapter 2 for a level loop whose dynamics move with the operating point.

4.1 Introduction

The water level in a steam drum is among the most heavily instrumented measurements in a power plant and among the most frequent reasons a unit comes off line. Neither direction is tolerant. A level low enough to uncover the steam-generating tubes overheats them within seconds, and a level high enough to carry water over with the steam delivers liquid to a machine turning at 3000 rpm. Protection is therefore hard-wired on both sides, the band between the two trips is a fraction of a meter, and the loop that has to stay inside it reads a measurement that moves the wrong way for the first minute of every disturbance. Level control in the steam generator of a pressurized water reactor is documented as a source of unplanned shutdowns, and the difficulty concentrates at low power rather than at full power [1, 2]. The same loop on fossil-fired and heat-recovery boilers carries the same reputation in the practitioner literature [3, 4].

> ⧖ HISTORY **Emergency Feedwater That Could Not Reach the Steam Generators (1979)**
>
> At 36 seconds after four in the morning on 28 March 1979, the pumps feeding the two steam generators of Unit 2 at Three Mile Island tripped while the reactor ran at 97 percent power, and the turbine was shut down automatically. Three emergency feedwater pumps started, but a valve was closed on each of the two emergency lines. The two control room lights that showed it were missed, and one was covered by a maintenance tag. About a minute and 45 seconds into the event the steam generators boiled dry. Eight minutes in, the closed valves were found and opened, and water rushed in. The commission found that this loss of emergency feedwater had no significant effect on the outcome, but it added to the confusion [5].

What the transmitter reports is not the water inventory. It is the height of a water surface, and beneath that surface is a population of steam bubbles that occupies real volume. The indicated level is therefore the sum of two states, a water volume and a sub-surface steam volume, and the second responds to the feedwater valve with the opposite sign to the first. Feedwater arrives subcooled, absorbs sensible heat on its way to saturation, and takes that heat out of the bubbles already in the drum. Each bubble that collapses gives back about forty times its own mass in volume, so admitting

water lowers the indication before it raises it. The opposite disturbance is more dangerous. Opening the turbine governor takes steam away faster than the firing rate can replace the energy, the drum pressure falls, saturated steam density falls with it, and every bubble under the water line expands. The level rises while the drum loses mass, which is the transient in which a controller reading only level closes the feedwater valve on an emptying boiler.

A linearized model of that behavior carries a numerator root in the right half plane, and the classical consequences follow. The root cannot be canceled, it cannot be moved by tuning, and it caps the achievable closed-loop bandwidth at a value no gain can raise [6], as Section 1.4.4 shows for a two-term drum model. The process control literature has treated inverse response as a distinct design problem since the early 1960s [7], and tuning rules exist for the specific combination of an integrator, dead time and a right-half-plane zero [8]. The drum adds a property those treatments usually hold fixed. Natural circulation is driven by the density head the heated risers develop, so the circulating flow falls with the firing rate and steam spends longer below the water surface before it separates. The zero therefore migrates toward the origin as the boiler is turned down, so the achievable bandwidth is lowest at the load a cycled plant spends its nights at and highest at the load it was commissioned on.

That class is wider than boilers. Any plant whose measurement is a sum of two states that answer the manipulated variable with opposite signs and different speeds has the same shape, and vessels with a phase change in them are full of examples, including flash drums, evaporator bodies, reboiler sumps and refrigeration receivers. The flash vessel of Volume 1, Chapter 6 swells on a rise in steam demand while it empties, for the same reason. The broader class is any nonminimum-phase plant whose zero location is a function of the operating point rather than a constant of the hardware. The fixed version of the problem as it appears in a textbook does not require the loop to be retuned, gain scheduled or transferred to a different architecture as the plant moves, and this one does.

Simulators of drum boilers are usually built to answer a modeling question rather than a control-architecture question [9–11]. However, a transfer function whose numerator root was chosen by the author of a textbook problem presents inverse response as a property of algebra, and a student who observes that a smaller move produces a smaller dip draws the wrong conclusion from it, because the depth of the dip belongs to the step while the interval the response spends traveling backward belongs to the plant. That presentation also cannot show the outcome the practitioner literature reports, which is that most drum level events are decided by which measurements were wired in and by whether the transmitter is reading correctly, long before anybody chooses a gain [3].

This chapter works with an interactive natural-circulation steam generator in which shrink and swell emerge from a mass and energy balance rather than from a numerator coefficient. It measures what a load-dependent right-half-plane zero costs and what each rung of the single-element, two-element and three-element ladder buys. The drum model is a teaching-scale reduction of published work [11, 12], and three-element control has been standard practice for decades [4, 13]. On this plant the three architectures share a single actuator, a single sample period and a single set of limits, so the comparison between them is a matter of wiring and of nothing else. The differential-pressure cell's density error is computed from the physics rather than injected as a bias. The scenario sequence separates what the plant decides from what the engineer decides, and it ends on two faults no gain reaches.

The principles the plant carries are inverse response and the distinction between the depth of a wrong-way excursion and its duration, feedforward and cascade in the form in which they are actually wired, the time-scale separation a cascade requires, actuator limits and anti-windup, and the rule that a loop holding a biased measurement holds the bias [14]. The physical fundamentals are three conservation statements on one vessel, a mass balance on the water, a mass balance on the steam held below the surface and an energy balance that carries the pressure, coupled through saturated properties and separated in time by a circulation that is fast against the drum and slow against the valve. The industrial warrant is direct, because drum level is among the most frequent causes of unplanned outages on both fossil and nuclear units, and the fault modes that cause those outages are structural rather than arithmetic [1, 3].

Section 4.2 traces the architecture ladder and the drum models behind it. Section 4.3 states the balances, the property fits, the constants and the integration scheme, so that the plant can be rebuilt from this chapter. Section 4.4 measures the inverse response in open loop and its dependence on load. Section 4.5 describes the three architectures, the PID algorithm they share, the limits of a feedforward delivered as a valve position and the predictive controller. Section 4.6 compares the architectures on load changes, a pump trip, a swell, a miscalibrated transmitter, a noisy level and the graded challenge, and sets the predictive controller against three-element control. Sections 4.7 and 4.9 discuss what the results mean for plant practice and where the model stops, Section 4.8 sets the loop in its plant hardware, and Section 4.10 sums up. The exercises follow in Section 4.11.

4.2 Background

Feedwater regulation is older than feedback control as a discipline. The first devices were mechanical, a float in the drum or a thermostatic expansion tube across the water line, acting directly on the feed valve. Both are proportional controllers whose gain is set by geometry and whose offset nobody measures, and both respond to the indicated level, which means both respond to shrink and swell in the wrong direction. Practitioner accounts of boiler control describe the sequence from those regulators through single-element, two-element and three-element schemes as the industry's response to exactly that defect [3, 4].

The quantitative case arrived with the first dynamic analyses of boilers. A 1958 study derived the drum water level response from mass and energy balances and identified the sub-surface steam volume as the cause of the wrong-way motion [9], and the effect of rapid steam take-off on natural circulation and drum water level was studied in its own right once turbines were expected to pick up load quickly [15]. A reduced nonlinear model of a drum boiler and turbine unit intended for control design appeared in 1972 [10] and reached its standard form as a fourth-order model whose drum, downcomer and riser balances reproduce the measured level response including the sub-surface steam volume [11]. Later work treated the void fraction distribution and the circulation loop as the mechanism that sets both the size and the duration of the excursion [12], and reviews written for the control audience place shrink and swell alongside pressure and temperature as the three couplings a unit controller has to respect [16].

The architecture settled before the modeling did. Single-element control uses drum level alone. Two-element control adds a steam flow measurement as a feedforward, on the argument that a kilogram leaving has to be replaced by a kilogram entering and that the steam flow signal moves the instant the governor does. Three-element control adds a feedwater flow measurement and an inner flow controller, so that the level controller writes a flow setpoint rather than a valve position. By the time the standard process control references were written the three-element scheme was described as the normal arrangement for any drum that swings load, with single-element control retained for low loads at which the feedwater flow measurement stops being trustworthy [4, 13]. The nuclear industry reached the same structure by a different route and has continued to publish on it, because a U-tube steam generator has stronger shrink and swell than a fossil drum and its worst case is at low power where the plant is least willing to accept an excursion [1, 2]. A boiling water reactor carries the same three-element structure on its vessel level, where the voids are also the moderator (Chapter 6).

Two developments have kept the loop current, and neither is a tuning problem. The first is instrumentation. A differential-pressure cell infers a height by dividing a pressure difference by a density difference, and the densities it uses are the ones it was ranged with rather than the ones in the drum, so its error moves with operating pressure and is shared by the trips that read the same transmitter [3, 17]. The second is duty. Thermal plants commissioned as base-load machines are now cycled to accommodate variable renewable generation, which moves them into the low-load, fast-ramp region where the level loop is hardest and where the architecture chosen at design time may no longer be the right one [18].

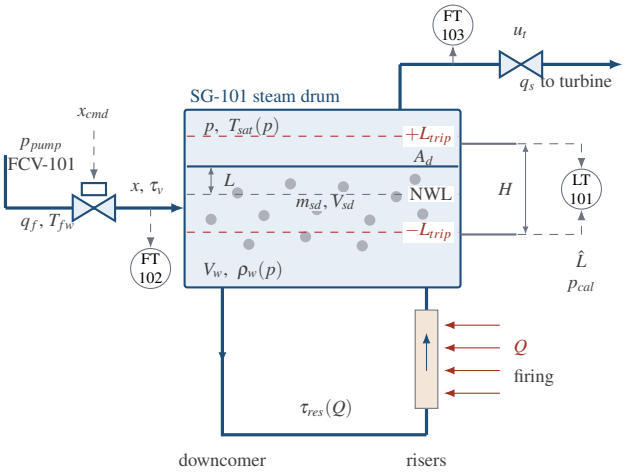

Figure 4.1 Natural-circulation steam generator SG-101 with its feedwater valve, turbine governor and drum level measurement. Water volume V_w occupies the drum below a surface of area A_d, and steam mass m_{sd} occupying volume V_{sd} is carried below that surface by a circulation whose residence time τ_{res} falls with the firing rate Q. Feedwater q_f at T_{fw} enters through FCV-101 at stem position x against pump pressure p_{pump}, and steam q_s leaves through the governor at opening u_t. LT-101 is a differential-pressure cell across taps H apart, ranged at pressure p_{cal} and reporting $\hat{L}$. Both level trips act on $\hat{L}$.

> **HISTORY** **From Three Elements to n Elements (1979)**
>
> A paper prepared in 1979 for a nuclear society meeting observed that plants were still being shut down by loss of steam generator water level control. It named the difficulties: a process with several interacting variables, a level that answers a steam flow change the wrong way, a delayed answer to feed flow, flows sensitive to valve position and uncertain measurements. Many were worse at low power, so the control schemes of the day ran automatically only above 15 percent power and left the range below to manual control. The paper proposed state feedback from a model of the steam generator, a natural extension of the classical three-element controller to n elements, one for each state of the model [19].

4.3 Steam Generator Model

Steam generator SG-101 is a natural-circulation drum boiler rated at 25 kg/s of saturated steam at 40 bar, which is 250.3 °C of saturation temperature and 53.0 MW of absorbed heat. Water leaves the drum down an unheated downcomer, boils in heated risers and returns to the drum as a two-phase mixture. The steam separates and leaves through the turbine governor. Feedwater at 160 °C replaces it through equal-percentage valve FCV-101, which sits 70 % open at the design condition against a feed pump holding 55 bar. Level transmitter LT-101 is a differential-pressure cell straddling two drum taps 1.2 m apart and reports a 600 mm span with normal water level at its center. The arrangement, the instrument tags and each symbol used below appear in Figure 4.1. Both level trips act at ±250 mm and the deviation alarms at ±150 mm, so the whole usable band is half a meter.

Table 4.1 sorts the signals into controlled, manipulated, disturbance and state. Five states are integrated. The working units are the ones used in boiler practice, mass in kg, flow in kg/s, volume in m^3, pressure in bar, enthalpy in kJ/kg, heat in kW and time in sec. Level is carried in meters, so the display factor of 1000 is the only unit conversion anywhere in the application. Every controller works in percent of transmitter span and percent of rated steam flow, because that is what a distributed control system faceplate shows and because it makes one gain number mean the same thing in all three architectures.

Table 4.1 What the Level Loop Interacts With. The Nominal Column Is the Full-Load Design Condition.

Symbol	Role	Nominal	Range
$\hat{L}$ (LT-101)	CV, measured	0 mm	±150 mm SP
q_f (FT-102)	inner CV, measured	100 %	0–160 % SP
x_{cmd} (FCV-101)	MV	70 %	0–100 %
steam demand	DV, measured	100 %	20–120 %
T_{fw}	DV, measured	160 °C	100–185 °C
p_{pump}	DV, measured	55 bar	40–70 bar
firing bias	DV, unmeasured	0 %	±25 %
LT bias	DV, unmeasured	0 mm	±40 mm
LT noise σ	DV, unmeasured	0 mm	0–12 mm
FT noise σ	DV, unmeasured	0 %	0–6 %
V_w	state	10 m^3	$\geq$ 0.2 m^3
m_{sd}	state	300 kg	$\geq$ 0
p	state	40 bar	5–90 bar
Q	state	53030 kW	0.05–1.4 Q_0
x (stem)	state	0.70	0–1

✦ DID YOU KNOW? Feedwater at Plant Scale

The two main feedwater pumps of Three Mile Island Unit 2 could supply its two steam generators with up to 15,500 gallons of water a minute [5], nearly a cubic meter every second. The boiler of this chapter is rated at 25 kg/s of steam.

4.3.1 Indicated Level, Water Volume and Sub-Surface Steam

The indicated level is algebraic rather than a state, and it is the relation every other result in this chapter depends on, as shown in Equation (4.1).

$$L = \frac{V_w - V_{w,ref}}{A_d} + k_{void}\,\frac{V_{sd} - V_{sd,ref}}{A_d}, \qquad V_{sd} = \frac{m_{sd}}{\rho_s(p)} \tag{4.1}$$

where L is the true level in m from normal water level, V_w is the drum water volume, A_d is the water-surface area, V_{sd} is the volume of steam held below the surface, k_{void} is the fraction of that volume that appears as indicated level, and the two reference volumes are the design values. Writing both terms as deviations makes L zero at the design condition by construction rather than by an offset. Setting k_{void} to zero deletes the second term, and with it the inverse response, which is the cheapest available test of what that term is responsible for. Equation (4.1) is the nonlinear counterpart of Equation (1.7). Its first term is the integrating water term of that equation, and its second term is the void term. The void relaxation time τ_{void} of Equation (4.3) below plays the part of τ_v there, since this chapter uses τ_v for the valve stroking time.

The water inventory is a pure integrator of the flow imbalance, corrected for the mass that has just left the liquid to become vapor, as shown in Equation (4.2).

$$\rho_w(p)\,\frac{\mathrm{d}V_w}{\mathrm{d}t} = q_f - q_s - \frac{\mathrm{d}m_{sd}}{\mathrm{d}t} \tag{4.2}$$

where q_f is feedwater flow and q_s is steam flow, both in kg/s. The bubble correction is worth about 2.5 % at 40 bar and it is what closes the mass balance. With $A_d = 20\,m^2$ and $\rho_w = 798.6\,kg/m^3$, a total loss of feedwater at full load drops the water term 1.57 mm/s.

The sub-surface steam is the state that carries shrink, and it relaxes toward the holdup the circulation would settle at, as shown in Equations (4.3) and (4.4).

$$\frac{\mathrm{d}m_{sd}}{\mathrm{d}t} = \frac{m_{sd}^{ss} - m_{sd}}{\tau_{void}} \tag{4.3}$$

$$m_{sd}^{ss} = \tau_{res}(Q)\left[\frac{Q}{h_{fg}(p)} - \frac{q_f\left(h_w(p) - h_{fw}\right)}{h_{fg}(p)}\right] \tag{4.4}$$

where τ_{void} is the void relaxation time, Q is the heat absorbed in kW, h_{fg} is the latent heat, h_w is the saturated water enthalpy and h_{fw} is the feedwater enthalpy, all in kJ/kg. The first term inside the bracket is steam made in the risers and the second is steam destroyed in the drum by the incoming subcooled water. The second term is shrink, written as a flow. Both m_{sd} and m_{sd}^{ss} are floored at zero.

Circulation is what makes the second term a function of load, as shown in Equation (4.5).

$$\tau_{res}(Q) = \frac{\tau_{res,0}}{\sqrt{0.25 + 0.75\,Q/Q_0}} \tag{4.5}$$

where $\tau_{res,0}$ is the riser residence time at full firing and Q_0 is the design heat input. Natural circulation is driven by the density head the risers develop, and that head falls with the firing rate, so the residence time runs from 12.0 sec at full firing to 13.6 sec at 70 %, 16.2 sec at 40 % and 18.1 sec at a quarter. The constant 0.25 floors the residence time at twice its full-load value, 24 sec, and it is a physical statement rather than a numerical convenience, because a drum boiler on standby still circulates and an unfloored (4.5) would predict a stagnant boiler at zero firing.

Pressure carries the energy balance and is where swell originates, as shown in Equation (4.6).

$$C_e\,\frac{\mathrm{d}p}{\mathrm{d}t} = Q - q_s\,h_{fg}(p) - q_f\left[h_w(p) - h_{fw}\right] \tag{4.6}$$

where C_e is the metal and water sensible heat per bar of saturation pressure, in kJ/bar. Saturated steam density is nearly proportional to pressure, so a pressure drop expands every bubble already under the water line through the $V_{sd} = m_{sd}/\rho_s(p)$ of Equation (4.1). Equations (4.1) and (4.6) taken together are swell. The feedwater term appears in both (4.6) and (4.4), which is why a feedwater step pushes the pressure down and the level down at the same time.

4.3.2 Firing Rate, Valves and Instruments

The fuel and air path is not part of the control problem. It is a load feedforward with a proportional pressure trim behind a first-order lag, as shown in Equations (4.7) and (4.8).

$$\tau_f\,\frac{\mathrm{d}Q}{\mathrm{d}t} = Q_{dem} - Q \tag{4.7}$$

$$Q_{dem} = Q_0\,u_t\left(1 + \tfrac{b_f}{100}\right) + K_m\,(p_0 - p) \tag{4.8}$$

where u_t is the governor opening, b_f is a firing-rate bias in percent, K_m is the boiler master pressure gain in kW/bar and Q_{dem} of (4.8) is clamped to $0.05\,Q_0 \le Q_{dem} \le 1.4\,Q_0$. The firing path is present because pressure is what carries a load change into the void volume. The governor follows the operator-set steam demand through a rate limit of $\dot{u}_t \le 0.5$ %/s by default.

Both flows are algebraic, as shown in Equations (4.9) and (4.10).

$$q_f = q_{f,\max}\,\frac{R^x - 1}{R - 1}\sqrt{\frac{\max(0,\ p_{pump} - p)}{\Delta p_0}} \tag{4.9}$$

$$q_s = K_t\,u_t\,p \tag{4.10}$$

where R is the equal-percentage rangeability, x is the stem position in the interval 0 to 1, $q_{f,\max}$ is the flow the valve passes wide open at design differential Δp_0, and K_t is the choked-nozzle coefficient in kg/(s bar). The characteristic in (4.9) is exactly zero at the seat and exactly one wide open. Its derivative is $R^x \ln R/(R-1)$, so at $R = 12$ the valve has 2.62 times the gain at 70 % open that it has at 31.2 % open, which is the opening that passes a quarter of rated flow. The square root in (4.9) is what makes a pump trip damaging, because a discharge pressure falling from 55 bar to 43 bar cuts the valve differential from 15 bar to 3 bar and the same stem position then passes $\sqrt{3/15} = 45\,\%$ of the flow it did. The stem itself strokes as a first-order lag of time constant τ_v toward the commanded position, with the derivative blocked at both stops so that the fixed-step integrator cannot pass through a limit and return.

Saturated properties come from quadratics fitted at 20, 40 and 60 bar, listed in Equation (4.11) with p in bar and T in °C.

$$\begin{aligned} \rho_s &= 0.4956 + 0.45663\,p + 8.306\times10^{-4}\,p^2 \\ \rho_w &= 923.09 - 3.7743\,p + 0.016556\,p^2 \\ h_w &= 677.3 + 12.8775\,p - 0.065625\,p^2 \\ h_{fg} &= 2098.1 - 11.32\,p + 0.04225\,p^2 \\ T_{sat} &= 161.745 + 2.849\,p - 0.0158625\,p^2 \\ h_{fw} &= 4.19\,T + 3.4\times10^{-4}\,T^2 \end{aligned} \tag{4.11}$$

The fits reproduce 30 and 50 bar to better than 0.5 %, and the pressure used for a property lookup is clamped to the interval 8 to 75 bar rather than extrapolated. At the design pressure they give $\rho_s = 20.09\,\mathrm{kg/m^3}$, $\rho_w = 798.6\,\mathrm{kg/m^3}$, $h_w = 1087.4\,\mathrm{kJ/kg}$ and $h_{fg} = 1712.9\,\mathrm{kJ/kg}$, against $h_{fw} = 679.1\,\mathrm{kJ/kg}$ for 160 °C feedwater.

The level transmitter is the last equation and the most consequential one. A differential-pressure cell infers a height by dividing a pressure difference by a density difference, and it uses the difference it was ranged with, as shown in Equation (4.12).

$$\hat{L} = \frac{\left(\frac{H}{2} + L\right)\left[\rho_w(p) - \rho_s(p)\right] + H\left[\rho_s(p) - \rho_s(p_{cal})\right]}{\rho_w(p_{cal}) - \rho_s(p_{cal})} - \frac{H}{2} \tag{4.12}$$

where H is the tap separation and p_{cal} is the pressure the cell was ranged at. Equation (4.12) returns L exactly when $p = p_{cal}$, which is what allows it to report only the real error. The controller then sees $\hat{L}$ plus an operator-settable bias plus band-limited noise, passed through a first-order transmitter lag τ_{LT}. Both flow measurements pass through their own lag τ_{FT}, and the steam flow measurement carries the flow transmitter noise because it reaches the valve through the feedforward.

4.3.3 The Design Point and the Numerics

Table 4.2 holds the constants. Everything else is algebraic in them and is recomputed once per controller execution rather than cached, because any parameter may be changed at any time. The derived design constants are $Q_0 = q_{s,0}(h_{fg,0} + \Delta h_{sub,0}) = 53030\,\mathrm{kW}$, $m_{sd,0} = \tau_{res,0}\,q_{s,0} = 300\,\mathrm{kg}$, $V_{sd,ref} = m_{sd,0}/\rho_s(p_0) = 14.93\,\mathrm{m^3}$, $q_{f,\max} = 58.6\,\mathrm{kg/s}$, $\Delta p_0 = 15\,\mathrm{bar}$ and $K_t = q_{s,0}/p_0 = 0.625\,\mathrm{kg/(s\,bar)}$, where $\Delta h_{sub,0} = h_w(p_0) - h_{fw} = 408.3\,\mathrm{kJ/kg}$ is the design subcooling.

Every scenario opens at a genuine equilibrium. The steam flow is computed from (4.10) at the scenario's demand, the feedwater is set equal to it, the firing is set to the value that makes (4.6) stationary, the sub-surface steam is set to (4.4), the stem is set to the opening that passes the required flow through (4.9), and the water volume is then chosen so that the *indicated* level sits on the operator's setpoint by inverting (4.12). A scenario therefore starts on setpoint even when its transmitter is miscalibrated. The application checks the result on itself, and the largest scaled state derivative at the design condition is below 10^{-12}.

Table 4.2 Constants of the Plant. Every One Is Operator Adjustable Except $V_{w,ref}$, Which Is Fixed in the Source.

Symbol	Meaning	Value	Unit	Visibility
$q_{s,0}$	rated steam flow	25	kg/s	adjustable
p_0	design drum pressure	40	bar	adjustable
C_e	pressure energy capacity	195000	kJ/bar	adjustable
$T_{fw,0}$	design feedwater temp.	160	°C	adjustable
A_d	drum surface area	20	m^2	adjustable
L_{span}	transmitter span	600	mm	adjustable
k_{void}	void-to-level coupling	0.526	—	adjustable
$\tau_{res,0}$	residence time, full fire	12	sec	adjustable
τ_{void}	void relaxation time	20	sec	adjustable
x_0	valve at design flow	70	%	adjustable
R	valve rangeability	12	—	adjustable
τ_v	valve stroking time	4	sec	adjustable
$p_{pump,0}$	design pump pressure	55	bar	adjustable
$\dot{u}_{t,\max}$	turbine ramp rate	0.5	%/s	adjustable
τ_f	firing response time	30	sec	adjustable
K_m	boiler master gain	2000	kW/bar	adjustable
τ_{LT}	level transmitter lag	3	sec	adjustable
τ_{FT}	flow transmitter lag	2	sec	adjustable
H	level tap separation	1.2	m	adjustable
p_{cal}	transmitter range pressure	40	bar	adjustable
L_{alarm}	deviation alarm	150	mm	adjustable
L_{trip}	low-low and high-high trip	250	mm	adjustable
$V_{w,ref}$	water volume at NWL	10	m^3	hidden

The engine of Section 1.6.1 integrates the plant by fourth-order Runge-Kutta on a fixed 0.05 sec step, which is eighty steps per valve stroke and is therefore sized against the fastest mode in the model. Both controllers execute on a 0.5 sec period, which is a distributed control system scan, with the first execution 0.5 sec after the one that initializes the controllers at $t = 0$, and the historian logs ten samples per simulated second. Noise is drawn from a generator seeded at 20260829, so scenario 13 replays identically on every machine. The speeds run from $0.5\times$ to $10\times$.

4.3.4 Assumptions

The model is detailed where a controller has to face the consequence and thin everywhere else. Each item below names the point at which the simplification stops being safe.

Five states and lumped saturated properties. This is a teaching-scale reduction of the published drum-boiler balances [11]. It reproduces shrink, swell, their dependence on load and the right-half-plane zero they create, and it says nothing quantitative about any particular boiler.

The whole circulation loop is one number. Equation (4.5) replaces the riser and downcomer system with a single residence time, so there is no distributed void profile up the riser, no slip between the phases and no downcomer subcooling [12]. There is therefore no circulation instability, which matters because a real natural-circulation boiler at very low load can circulate unstably and low load is where this application is otherwise most interesting.

Two volume effects are neglected. The thermal expansion of the drum water during a pressure transient is not tracked as a separate contribution, which understates swell by roughly one part in forty of the water-side term. The steam space above the water line is not a state, and at about 10 m^3 of steam space a pressure ramp of 0.05 bar/s is a 0.26 kg/s error against flows of 25 kg/s, so it costs about 1 % during the fastest transient and nothing at steady state.

The fuel side is not modeled as a control system. Equations (4.7) and (4.8) carry no air-fuel cross limiting, no firing rate limit and no pressure controller, all of which a real boiler master carries [16]. The firing bias is the only way to make the fuel side misbehave.

The turbine is a choked nozzle with a ramp limit. There is no steam line volume, no drum-to-throttle pressure drop and no superheater, so steam pressure control cannot be studied here and neither can its interaction with level control.

The instrument chain is idealized apart from the density error. Equation (4.12) is exact, but reference-leg boiling, tap plugging and the standing waves of a real drum surface are absent, and each of those is a documented reason a drum level indication is wrong [3, 17].

The valve has no stiction, backlash or positioner. Feedwater valve stiction is a common cause of real level cycling and is exactly the fault an inner flow loop half conceals, which makes its absence worth stating.

A trip does not stop the plant. Both level trips latch and annunciate without stopping the run, which no real protection system would do. The transient after a trip is the part worth reading, so it is kept on the trend. The only protective action is the feedwater interlock of Section 4.5.

The trajectory repeats better than a plant does. Fixing the step, the scan period and the seed produces one history on every machine, which a real distributed control system does not.

4.4 Inverse Response and Its Dependence on Load

The experiment that establishes the plant's central property takes no controller at all. In MANUAL at full load, with feedwater and steam balanced at 25.0 kg/s, FCV-101 is opened from 70 % to 78 % at 60 sec (the scripted step of scenario 1) and the feedwater rises to 31.8 kg/s. More water is entering than leaving and the drum is gaining mass, yet the level falls to −3.95 mm 19 sec after the step and does not return to where it started until 38 sec. The level is reported decomposed into its two terms, and at the bottom of the dip the water contribution is +7.1 mm against a void contribution of −11.1 mm. The mass balance is behaving correctly and the void term is larger.

> **TRY IT IN THE APP** **1 · Drive the Valve by Hand (Manual)**
>
> Load scenario 1 and press Run. At 60 sec the script opens the feedwater valve from 70 to 78 %, so more water enters than leaves. Watch the drum level: it falls first, to −3.95 mm 19 sec after the step, and does not return to where it started until 38 sec, because the entering water collapses steam bubbles.

The arithmetic is entirely in the enthalpy of Equation (4.4). Saturated water at 40 bar carries 1087.4 kJ/kg and 160 °C feedwater carries 679.1 kJ/kg, so every kilogram entering the drum absorbs 408.3 kJ before it can boil. Latent heat is 1712.9 kJ/kg, so each kilogram of feedwater condenses 0.238 kg of steam on its way to saturation. That steam occupied space at 20.09 kg/m^3 and the water replacing it occupies space at 798.6 kg/m^3, so a bubble that collapses returns 39.8 times its own mass in volume. The void loss exceeds the mass gain for the first half minute and is smaller than it afterward.

Pause. Try Exercise 4.1 before reading on.

Two properties of the response behave differently, and Table 4.3 separates them. Depth is a gain. Increasing the feedwater by 10, 25, 50 and 100 % at full load troughs at −1.47, −3.70, −7.46 and −15.21 mm, a factor of ten across a factor of ten in the step. The crossing time is a dynamic property. All four come back through zero between 38 and 39 sec. Confusing the two leads to the attempt to fix an inverse response by moving less, which buys a smaller excursion and no speed at all. The rows are model runs rather than interface procedures, because each starts from the exact equilibrium opening, the step positions fall between the half-point increments of the manual valve station and the scripted step of scenario 1 is removed. Table 4.3 reports the drum level rather than the transmitter output, because the pressure falls during the step and the differential-pressure cell of Equation (4.12)

Table 4.3 Open-Loop Feedwater Step in MANUAL, With the Valve Moved From Its Equilibrium Opening to the Position That Passes the Stated Increase in Feedwater. Trough Depth Scales With the Step and the Crossing Time Does Not. Times Are Measured From the Step.

	Trough (mm)	Time of trough (sec)	Crossing (sec)
Step size, at 100 % load			
+10 % feedwater	−1.47	19	38
+25 % feedwater	−3.70	19	38
+50 % feedwater	−7.46	20	38
+100 % feedwater	−15.21	20	39
Load, at a +25 % feedwater step			
100 % load	−3.70	19	38
70 % load	−4.05	23	47
40 % load	−3.73	26	58
25 % load	−3.04	29	65

therefore returns a trough about 7 % shallower at full load and 4 % shallower at a quarter, and a crossing about 2 sec earlier.

🕮 **Pause.** Try Exercise 4.2 before reading on.

Changing the load does the opposite. The trough stays within 0.7 mm of −3.7 mm across the whole range while the crossing time runs from 38 sec at full load to 65 sec at a quarter. Equation (4.5) is the mechanism. Slower circulation holds steam under the water line longer per kilogram being made, so the sub-surface steam volume falls only from 14.93 m^3 at full load to 8.05 m^3 at 40 % while the steam production falls to two fifths, and the void contribution to the indicated level runs from 0 mm at full load to −181 mm at 40 % and −244 mm at a quarter. The zero has moved toward the origin, and the closed-loop bandwidth a feedback controller can carry has come down with it [6], through the bound that follows Equation (1.7).

No numerical zero location appears anywhere in this chapter, and the omission is deliberate. The coefficients of (4.4) and (4.5) are functions of the operating point, so a single value would be a linearization about one load presented as a property of the boiler. The crossing time in Table 4.3 is what a step test on the plant returns, and it moves for the same reason the zero does. The output voltage of a boost converter carries a right-half-plane zero that also moves with load, and Volume 3, Chapter 9 measures what an inner current loop returns against it.

The second term of (4.1) can be tested by deletion. Zeroing k_{void} and rerunning the closed-loop load rejection of Section 4.6 gives no undershoot at all, +15.9 mm of overshoot and 6 % of valve travel, against −29.8 mm, +37.7 mm and 36 % with the term in place. The drum becomes an ordinary integrating level and every difficulty in the application disappears with it.

Swell is the same mechanism with the sign reversed. Steam density in (4.1) is a function of pressure, so anything that moves the pressure moves the indicated level with no change in inventory at all. A load pickup is the dangerous direction, because the firing rate lags the demand by $\tau_f = 30$ sec in (4.7) while the governor opens at up to 5 %/s, so the pressure falls, every bubble expands and the level climbs while the drum empties. That effect was quantified for rapid steam take-off on a real natural-circulation boiler decades ago [15].

4.5 Controller Design

Every architecture in the application drives the same feedwater valve through the same first-order stem, so the manipulated variable in the trends and in the comma-separated export is always FCV-101 and what changes is only what decides it. The three forms are given in Equation (4.13) and drawn in

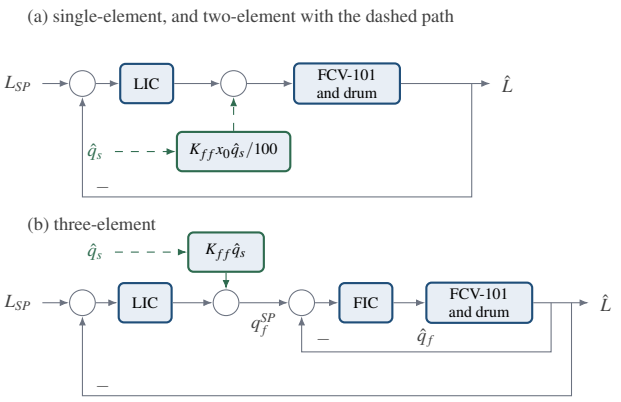

Figure 4.2 The three drum level architectures as implemented. All of them drive the same valve. Single-element control is the solid path in (a). Two-element control adds the dashed steam-flow feedforward, which has to be delivered as a valve position because there is no flow loop to receive it. Three-element control in (b) makes the level controller write a feedwater flow setpoint q_f^{SP}, states the feedforward in flow, and closes an inner loop on the measured feedwater flow $\hat{q}_f$. The output of the block on the right is the commanded valve position x_{cmd} of Figure 4.1.

Figure 4.2.

$$\begin{aligned} \text{single: } x &= \tfrac{1}{100}\,\mathrm{LIC}\big(L_{SP},\hat{L}\big) \\ \text{two: } x &= \tfrac{1}{100}\left[K_{ff}x_0\frac{\hat{q}_s}{100}+\mathrm{LIC}\big(L_{SP},\hat{L}\big)\right] \\ \text{three: } q_f^{SP} &= K_{ff}\hat{q}_s+\mathrm{LIC}\big(L_{SP},\hat{L}\big), \\ x &= \tfrac{1}{100}\,\mathrm{FIC}\big(q_f^{SP},\hat{q}_f\big) \end{aligned} \tag{4.13}$$

where $\hat{q}_s$ and $\hat{q}_f$ are the lagged steam and feedwater flow measurements in percent of rated steam flow, K_{ff} is the feedforward gain and x_0 is the design valve opening in percent. The feedforward is delivered as the output bias of the level controller rather than as a separate summing junction, so the integral term of the LIC carries exactly the difference between what the feedforward asked for and what the plant needed, and the application reports those two numbers side by side. A fourth structure, the model predictive controller of Section 4.5.3, writes the same flow setpoint q_f^{SP} to the same inner loop.

4.5.1 Control Algorithm

LIC-101 and FIC-102 run one algorithm between them, the velocity-form PID law of Section 1.6.2 [20], written here in the loop's own symbols. The filter and the three increments are Equations (4.14) to (4.17), and the recursion that assembles them is Equation (4.18). The running sum s_k is the S_k of Equation (1.14), and the demanded output c_k is the v_k of Equation (1.15).

$$\hat{L}_k^f = \alpha\,\hat{L}_{k-1}^f + (1-\alpha)\,\hat{L}_k, \qquad \alpha = \mathrm{e}^{-\Delta t_c/\tau_{df}} \tag{4.14}$$

$$\Delta u_k^P = K_c\,(e_k - e_{k-1}) \tag{4.15}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t_c \tag{4.16}$$

$$\Delta u_k^D = -\,\frac{K_c\tau_D\left(\hat{L}_k^f - 2\hat{L}_{k-1}^f + \hat{L}_{k-2}^f\right)}{\Delta t_c} \tag{4.17}$$

$$
\begin{aligned}
s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\
c_k &= u_{bias,k} + s_k + I_{k-1} \\
u_k &= \min(u_{max},\ \max(u_{min},\ c_k)) \\
I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\,(u_k - c_k)
\end{aligned}
\tag{4.18}
$$

where $e_k = L_{SP,k} - \hat{L}_k$ is the error in percent of transmitter span, $u_{bias,k}$ is the feedforward term of (4.13), $\hat{L}^f$ is the measurement filtered with a default τ_{df} of 2 sec for the derivative alone, Δt_c is the controller period, s_k accumulates the proportional and derivative increments without limit, c_k is the output the algorithm demands and T_t is the anti-windup tracking time. Equation (4.15) is written for a fixed gain. The app forms the increment as the change in $K_c e_k$, so a gain change acts on the whole error at the next execution. Both loops are reverse acting in the ISA sense, action $+1$ in the convention of Section 1.3, because more valve raises both the level and the feedwater flow, so a measurement above setpoint is met with less valve. The inverse response is a dynamic effect and does not change the sign of either process gain, which is the reason it is mistaken for one. The filter of (4.14) is advanced on every execution in MANUAL as well, so no derivative transient survives a transfer into AUTO. Windup protection is the back-calculation in the last line of (4.18), with T_t equal to each loop's own integral time (300 to 700 sec for the level controller in the shipped scenarios, and 8 or 120 sec for the flow controller). Only the integral is corrected, and the running sum s_k is never truncated. The LIC output is clamped to 0–100 % in single-element and two-element control, where it means a valve position, and to 0–160 % in three-element control, where it means a flow setpoint and a drum that has just swelled needs the headroom. The FIC output is clamped to 0–100 %. Neither controller carries an output rate limit, because the physical rate limit is the 4 sec valve stroke. Both execute every 0.5 sec, independently of the 0.05 sec integration step.

MANUAL inverts the same recursion for the integral, $I_k = u_{man} - u_{bias,k} - s_k$, as in Section 1.6.2. It places each controller output on the value it would have to deliver, the operator's valve position for the FIC and for the single-element and two-element LIC, and the measured feedwater flow for the three-element LIC. Changing architecture at run time is treated the same way, because the outer controller's output means a valve position in two of the three schemes and a flow setpoint in the third. For one execution after the switch both controllers track the valve the plant already has, which is the only way to avoid commanding a valve position of 25 % when the number in the integral was a flow setpoint of 25 %.

One protective action is implemented. After a high-high level trip, an interlock drives the feedwater valve shut, because water carried into a turbine is unrecoverable, and both controllers track the plant as in MANUAL while it holds, so neither integral winds up against it. It releases when the measurement is back inside the 150 mm deviation alarm rather than latching, and the release condition was chosen by measurement. Latched permanently, the interlock emptied the drum after every high trip, so a run that reached the high trip always finished at the low one and the transient that caused the original event became unreadable underneath a second, larger one.

4.5.2 Limits of a Feedforward Delivered as a Valve Position

The feedforward in three-element control is exact and needs no calibration, because a kilogram of steam leaving requires a kilogram of feedwater entering and that statement contains no valve, no pump and no operating point. $K_{ff} = 1$ is therefore the right answer rather than a tuned one. Two-element control has no flow loop to hand a flow setpoint to, so the same signal has to arrive as a valve position, and (4.13) scales the design opening x_0 by the steam flow fraction. That is a straight line through two points while the equal-percentage characteristic of (4.9) is a curve through the same two points, so the calculation is exact at one load and wrong everywhere else, as listed in Table 4.4.

The largest error in Table 4.4 is 14.5 points of valve travel at 40 % flow, and the sign reverses above the design point. Nothing in two-element control can remove it, because the scheme contains no measurement of what the valve delivered. Three-element control removes it by stating the feedforward

Table 4.4 Two-Element Feedforward Against the Valve It Is Driving. The Error Is Carried by the Level Controller's Integral Term.

Feedwater flow (% of rated)	Feedforward asks for (%)	Valve needs (%)	Error (points)
25	17.5	31.2	13.7
40	28.0	42.5	14.5
60	42.0	53.9	11.9
85	59.5	64.7	5.2
100	70.0	70.0	0
115	80.5	74.7	−5.8

in the units of the thing it is about, and the inner flow loop then absorbs the valve entirely, so the level controller faces a feedwater flow that follows its setpoint rather than an equal-percentage stem whose gain changes by a factor of 2.62 across the range.

4.5.3 Model Predictive Control

The model predictive controller (MPC) takes the place of LIC-101 in the three-element wiring of Equation (4.13) and writes q_f^{SP} to an unchanged FIC-102. It follows the receding-horizon scheme of Section 1.6.3 [21, 22], re-solving every $T_s = 2\,\mathrm{sec}$, or fourth execution. The voids of Equation (4.3) follow the measured steam flow, so a model can forecast a swell before the transmitter reports it.

The controller's model is a four-state reduction of Equations (4.1) to (4.8) with no pressure state, written from the parameters of Table 4.2 as they stood at the last reset, as shown in Equation (4.19).

$$\begin{aligned} \frac{\mathrm{d}\ell_w}{\mathrm{d}t} &= c_w\left(q_f - \hat{q}_s - \hat{d}\right), \qquad \tau_{void}\frac{\mathrm{d}b}{\mathrm{d}t} = b^{*} - b \\ \tau_{fl}\frac{\mathrm{d}q_f}{\mathrm{d}t} &= v - q_f, \qquad \tau_{th}\frac{\mathrm{d}z_f}{\mathrm{d}t} = q_f - z_f \\ b^{*} &= k_v\left[\tau_{res}(z)\,(z - \varepsilon q_f) - \tau_{res,0}\,q_{s,0}\right] \end{aligned} \tag{4.19}$$

Here ℓ_w is the level contributed by the water mass, counted from its estimate $\hat{L}_w$, b is the level contributed by the sub-surface steam net of the liquid it displaces, q_f is the delivered feedwater flow, z_f is that flow as the firing has replaced its sensible heat, v is the flow setpoint, $\hat{q}_s$ is the measured steam flow and $\hat{d}$ is an estimated unexplained outflow, all flows in kg/s. The heat input is counted as the steam flow it supports, $z = Q/h_{fg} = z_s + \varepsilon z_f$ with $\tau_{gov}\dot{z}_s = \hat{q}_s - z_s$, where $\varepsilon = \Delta h_{sub,0}/h_{fg} = 0.238$ and $\tau_{res}(z)$ is Equation (4.5) with $Q/Q_0 = z/[q_{s,0}(1+\varepsilon)]$. At 40 bar, $c_w = 1/(\rho_w A_d) = 6.261 \times 10^{-5}\,\mathrm{m/kg}$ and $k_v = (k_{void}/\rho_s - 1/\rho_w)/A_d = 1.2465 \times 10^{-3}\,\mathrm{m/kg}$, so b^{*} is k_v times the departure of Equation (4.4) from the design holdup, −172.3 mm at 40 % load against the plant's −180.9 mm, because the model holds the properties at design pressure. The predicted indication is $y = \hat{L}_w + \ell_w + b$.

The lags were identified by step tests on the simulator. The closed inner loop, stepped by ±10 % at the shipped flow tuning, fits a dead time of 1.15 to 1.25 sec and a lag of 2.05 to 2.43 sec at full load with residuals of 2.0 to 2.2 %, folded into $\tau_{fl} = 3.4\,\mathrm{sec}$ and scaled elsewhere as $\tau_{fl} = 3.4\,\mathrm{sec}\cdot R^{x_0 - x_q}(0.6/K_{c,F})(\tau_{I,F}/8\,\mathrm{sec})$, with x_q the opening that passes q_f. That gives 6.7 sec at 40 % load, where the fits sum to 6.9 to 7.4 sec, and 510 sec for the detuned loop of scenario 9 at full load. The firing follows a steam-flow change with $\tau_{gov} = 12.6\,\mathrm{sec}$ (fits of 11.99 to 13.25 sec) and replaces the sensible heat of extra feedwater through the pressure trim with $\tau_{th} = 104\,\mathrm{sec}$ (fits of 101.9 to 106.0 sec), with residuals of 11 to 16 % from the second-order shape of the trim. Run open loop beside the plant for 300 sec, the model's indication departs from the transmitter by 7.4 to 9.2 % of the excursion in root-mean-square after a valve step and by 15 to 20 % after a governor ramp. At each decision the model is linearized about the estimated heat input and the measured feed flow and discretized with an exact zero-order hold. The measured steam flow, held over the horizon, is the only

preview and enters through z_s as a known term, and it is the signal the three-element feedforward reads.

A steady-state Kalman filter estimates L_w, b, $\hat{d}$ and the transmitter output at every 0.5 sec execution under every structure, from the transmitter reading with an assumed noise of 2 mm and process noise per execution of 0.2 mm, 3 mm, 0.02 kg/s and 0.1 mm respectively, which gives the gain $[0.158,\ 0.914,\ -7.58\,\mathrm{kg/(s\,m)},\ 0.425]$. A persistent level error is absorbed into the random walk $\hat{d}$, the disturbance model that removes steady offset [23].

The plan chooses setpoints $v_1,\dots,v_{12}$ held over blocks of $1,1,1,1,2,2,3,4,6,9,15$ and 15 steps of a horizon of $N = 60$ steps, or 120 sec, longer than the 65 sec crossing time of Table 4.3 at a quarter load, to minimize Equation (4.20).

$$J = \sum_{k=1}^{N} e_k^2 + P e_N^2 + \lambda \sum_{j=1}^{12} \left(v_j - v_{j-1}\right)^2 + \rho_b s^2 + \sigma_b s \tag{4.20}$$

Here $e_k = y_k - L_{SP}$ in mm, v_0 is the setpoint applied now, $P = 10$, $\lambda = 300\,\mathrm{mm^2\,s^2/kg^2}$, and $\rho_b = 100$ per $\mathrm{mm^2}$ and $\sigma_b = 100$ per mm price the single slack $s \geq 0$ of the soft band $|e_k| \leq 150\,\mathrm{mm} + s$ at the deviation alarm. The only hard constraint is $0 \leq v_j \leq 1.6\,q_{s,0}$, the range of the three-element LIC. The valve keeps the clamp of FIC-102, neither PI loop carries a rate limit, and the trips are not constraints. Decisions are posed as deviations from v_0, so the small ridge that keeps the program strictly convex cannot pull the feed toward zero, and the program goes to the Goldfarb and Idnani dual active-set algorithm [24]. Only λ was tuned, on scenario 14, by the rule the predictive controllers of these apps share: among values from 0 to 3000 that reach no trip and use no more stem travel than three-element control, 152.6 %, keep the largest whose IAE at 600 sec lies within 10 % of the best. The best is 13801 mm s at $\lambda = 1$, and the rule keeps $\lambda = 300$ at 14844 mm s.

Any plan whose solver status is other than optimal, or whose first move is non-finite or outside the setpoint range, is discarded. LIC-101 tracks the applied setpoint at every execution, so it then takes over bumplessly and hands back once a valid plan returns. The first plan is solved at $t = 0$, and no plan was discarded in any run reported here. The controller reads only the transmitter, the measured flows, the setpoint and its copy of the parameters, which keeps the five-state plant, the instrument chain and the four-state model as three distinct objects.

4.6 Results

The app runs on the framework of Section 1.6.4. Manual mode positions the feedwater valve directly, automatic mode closes the architecture selected in Section 4.5, and every run exports as a comma-separated file. Fifteen scenarios are arranged as a laboratory sequence and are listed in Table 4.5. Each one presets the mode, the architecture, the tuning, the disturbances, any parameter it overrides and a list of timed events, and each figure quoted below came off the running application at the settings named next to it, changed before the run is started unless a Reset is named. Reset returns the plant to $t = 0$ and keeps every setting the learner changed, and reloading the scenario restores its own. Integrated absolute error keeps accruing after a transient ends, so every value reported here is the application's own running integral read 710 sec after the event that starts the transient, which makes the numbers comparable across controllers and meaningless in isolation. Valve travel accumulates in the same way and is read at the same instant. Two cases name their own window instead. The pump trip is still recovering at 710 sec, so its figures are read over the 900 sec following the trip, and the noise scenario applies no event at all and is read over its full fifteen minutes. The comparison with the predictive controller in Section 4.6.7 reads every case at the clock on which its lesson is credited instead. Peak deviation needs no reading instant. Every deviation quoted is the drum level rather than the transmitter output, except where the transmitter is the subject.

Each scenario also carries a completion rule, given in the last column of Table 4.5. A case is credited when the learner performs the listed actions and the simulation reaches the listed state, and

Table 4.5 Scenarios and Completion Requirements

#	Scenario (kind)	What it exercises	Required to complete
1	Drive the valve by hand *experiment*	inverse response with no controller	In MANUAL, the scripted valve step from 70 to 78 % at 60 sec followed to 120 sec. After Reset, a smaller step to between 70.5 and 76.5 % made before 60 sec, and after the scripted step feedwater matched to steam within ±0.3 kg/s for 30 sec.
2	Inverse response at 40 % load *experiment*	the zero moving with load	In MANUAL at 40 % load, the scripted valve step from 42.5 to 57.5 % at 60 sec followed to 150 sec, then the steam demand moved at least 10 points from 40 % and held there for 60 sec.
3	Single-element, 100 to 85 % *experiment*	one measurement that moves the wrong way	Single-element control at $K_c = 1.6$ held from the demand step at 90 sec to 600 sec, then after Reset a run at $K_c \geq 2.5$ held over the same window.
4	Same change at 40 % load *demonstration*	the same disturbance, harder plant	Single-element runs at 40 % load with $K_c = 1.6$, 1.0 and at least 2.5, each held from the step to 25 % at 90 sec to 600 sec.
5	Two-element feedforward *experiment*	a feedforward that fixes half the problem	Two-element control at $K_{ff} = 1$ with $K_c = 1.6$ and then $K_c \geq 3.5$, each held from the demand step at 90 sec to 600 sec.
6	Two-element at 40 % load *experiment*	a feedforward carrying a valve model	Two-element runs at 40 % load with valve rangeability 12 and then 2, each held from 90 sec to 600 sec.
7	Three-element control *experiment*	cascade plus an exact feedforward	Three-element control at $K_c = 3.5$ and then $K_c \geq 6$, each held from the demand step at 150 sec to 600 sec.
8	Three-element at 40 % load *experiment*	gain that no longer suits the load	Three-element runs at 40 % load with $K_c = 3.5$ and then $K_c \leq 2.5$, each held from 90 sec to 600 sec.
9	Inner flow loop detuned *demonstration*	the time-scale separation a cascade needs	Three-element runs with the flow loop at $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, at 0.2 and 30 sec, and at 0.6 and 8 sec, each held from 90 sec to 600 sec.
10	Feedwater pump trip *experiment*	a loop closed around the actuator	The pump trip at 90 sec ridden to 600 sec on two-element, three-element and single-element control in turn, each architecture in place from 90 sec.
11	Fast pickup, 25 to 120 % *demonstration*	swell, and the low-level excursion it leaves	Single-element control at $K_c = 1.0$ run until a level trip is reached, then runs held from the pickup at 60 sec to 600 sec at $K_c \geq 2.5$, on three-element control at $K_c = 3.5$, and on three-element control at $K_c = 3.5$ with the turbine ramp cut to at most 1 %/s.
12	DP cell ranged at 28 bar *demonstration*	an instrument fault no tuning reaches	The firing upset at 90 sec ridden to 600 sec with the cell ranged at 28 bar, then with it ranged at 40 bar.
13	Noisy level with derivative action *experiment*	actuator travel bought for nothing	Three idle runs held from the start to 900 sec, with $\tau_D = 20$ sec unfiltered, with $\tau_D = 0$, and with $\tau_D = 20$ sec behind a 5 sec derivative filter.
14	Three-element against MPC through a swell *experiment*	a model that forecasts the voids	At 40 % load with the turbine picking up to 100 % at 60 sec and 2 %/s, a run on the MPC held from 60 sec to 600 sec, then after Reset a run on three-element control selected before the pickup and held over the same window.
15	Graded challenge, three upsets *challenge*	choosing an architecture under time	Neither level trip reached in 900 sec through the pickup at 60 sec, the heater trip at 220 sec and the rejection at 380 sec, with the drum level inside ±30 mm of setpoint from 600 to 900 sec.[a]

[a] As shipped (single-element, $K_c = 1.0$, flow loop at 0.06 and 120 sec, governor ramp 2 %/s) the drum reaches the high-high trip at 559.5 sec and the worst level error over the last five minutes is 269.9 mm, 239.9 mm outside the band. The suggested settings (three-element, $K_c = 2.0$, $\tau_I = 500$ sec, $K_{ff} = 1$, flow loop at 0.6 and 8 sec) reach neither trip and hold the level within 14.8 mm, a margin of 15.2 mm, with the same result whether they are applied before the run or carried through Reset.

the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

4.6.1 Single-Element Control and the Limit on Gain

The architecture lessons require their settings to be held through the post-event comparison and are credited at $t = 600$ sec. Their Learn cards report that clock reading. Longer comparisons below state their horizon explicitly. In particular, Table 4.6 uses equal time after each event, accounting for scenario 7's later step.

Scenario 3 puts the level controller in service with drum level as its only measurement and backs the turbine from 100 % to 85 %, so the drum begins to gain water at 3.75 kg/s. The level falls to −29.8 mm at 168 sec, which is swell with the sign reversed. Less steam leaving raises the pressure, higher pressure compresses the bubbles, and the void contribution has to fall by 38 mm between 100 % and 85 % load. The controller reads a low level and opens the feedwater valve, which condenses more steam and holds the level down further, and by the time the mass balance exceeds the void term the drum has been overfilled. The level peaks at +37.7 mm at 286 sec, which is 67 mm peak to peak for a fifteen-point load change on the easiest part of the range.

Gain does not rescue it. Sweeping K_c through 1.0, 1.6, 2.0, 2.5 and 3.0 read at $t = 800$ sec gives integrated errors of 16074, 11486, 9938, 10414 and 18307 mm s for 17, 36, 60, 124 and 299 % of valve travel. The error stops improving at 2.0 and the loop oscillates at 3.0, on a disturbance that needs a five-point valve move. The number that barely moves is the instructive one. The trough runs only from −29.5 to −32.4 mm across the whole sweep, because the trough is the void collapse and feedback cannot reach it. What the gain buys is a faster recovery afterward and what it costs is overshoot, from +34.8 mm at $K_c = 1.0$ to +61.3 mm at $K_c = 3.0$. That is the right-half-plane zero setting a ceiling and gain converting the wrong-way response into an oscillation rather than into performance [13].

4.6.2 The Same Load Change at Reduced Load

Scenario 4 is the identical controller and the identical 3.75 kg/s at 40 % load rather than at full load. Peak deviation goes from −29.8 to −66.3 mm and integrated error from 11486 to 25341 mm s, a factor of 2.2. Two effects stack. The void swing per point of load is larger, 63 mm instead of 38, and the wrong-way part of the response now lasts 58 sec instead of 38 (Table 4.3), so the loop has to be slower at exactly the moment the disturbance has grown. Raising the gain makes it worse. At $K_c = 2.5$ the trough deepens to −83.8 mm and the overshoot reaches +123.3 mm with 162 % of valve travel, which is a full oscillation rather than a control response.

The obvious version of that comparison is wrong, and the application reports the reversal rather than hiding it. Measured at the same *relative* step instead, a fifteen percent load reduction from 40 % down to 34 %, the low-load case is slightly *better* than the full-load one, at −23.2 and +18.7 mm for 9107 mm s and 20 % of travel, against −29.8 and +37.7 mm for 11486 mm s and 36 %. That run is made on scenario 4 by setting the demand to 34 % just after its scripted step at 90 sec, which the 0.5 %/s governor ramp makes identical to a step to 34 %. The equal-percentage valve's gain falls faster with load than the bandwidth ceiling comes down, so a proportional disturbance at the bottom of the range is not harder than the same proportional disturbance at the top. The statement about the absolute change is the operationally relevant one. A house load, a soot blower, an auxiliary steam demand and a mill trip are fixed quantities of steam, every one of them is a larger fraction of a turned-down boiler than of a full one, and the same 3.75 kg/s arrives 2.2 times worse at 40 % load.

4.6.3 Comparison of the Three Architectures

Table 4.6 carries the three architectures on the identical load rejection at both loads. Two-element control makes the dip at full load *deeper*, from −29.8 to −50.2 mm, and removes the overshoot completely, and that combination is the most useful result in the application. The feedforward makes feedwater equal steam, which is exactly right for the water inventory of (4.2) and does nothing at all about the voids of (4.3). The bubble volume genuinely has to shrink by 38 mm between the two loads, so holding the mass constant means the level genuinely has to fall by 38 mm. What the feedforward

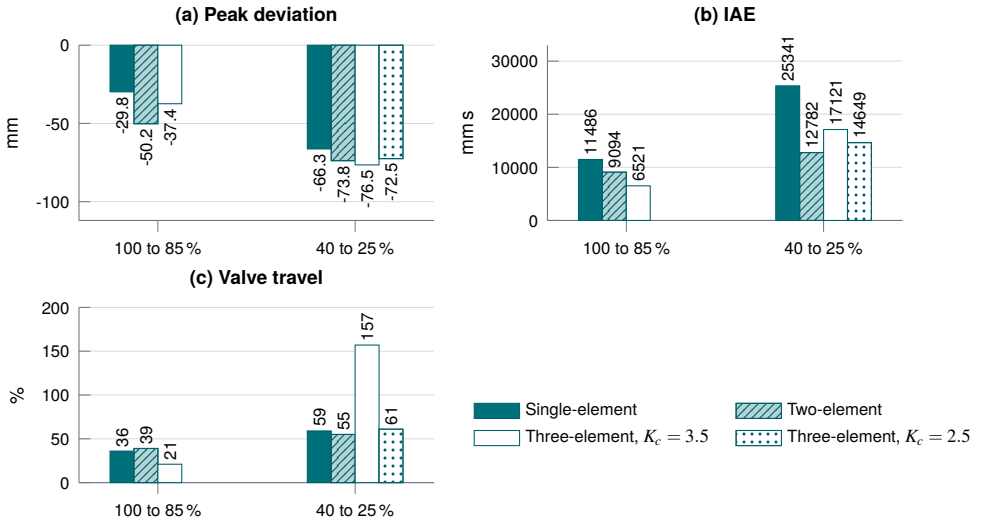

Figure 4.3 The fifteen-point load rejection of Table 4.6, from 100 to 85 % load and from 40 to 25 % load, under the three architectures: (a) peak deviation, (b) integrated absolute error read 710 sec after the load change and (c) valve travel. Solid: single-element. Hatched: two-element. White: three-element at $K_c = 3.5$. Dotted: three-element at $K_c = 2.5$, run at the lower load only. Notice that the ladder ranks cleanly at full load, while at the lower load two-element control posts the smallest error and the full-load gain of three-element control costs 157 % of travel.

Table 4.6 The Architecture Ladder on the Identical Fifteen-Point Load Rejection. Every Integral and Travel Is Read 710 sec After the Load Change. The Low-Load Runs Are the Same 3.75 kg/s of Steam.

	K_c	Peak (mm)	IAE (mm s)	Travel (%)
100 % to 85 % load				
Single (sc. 3)	1.6	−29.8	11486	36
Two (sc. 5)	1.6	−50.2	9094	39
Three (sc. 7)	3.5	−37.4	6521	21
40 % to 25 % load				
Single (sc. 4)	1.6	−66.3	25341	59
Two (sc. 6)	1.6	−73.8	12782	55
Three (sc. 8)	3.5	−76.5	17121	157
Three (sc. 8)	2.5	−72.5	14649	61

removed is the part of the excursion the controller was adding on its own. Figure 4.3 draws Table 4.6 as bars.

🕮 **Pause.** Try Exercise 4.3 before reading on.

The two-element feedforward gain is more interesting than one for one suggests. For an independent comparison, the learner reloads scenario 5, sets the gain before Run and reads every result at $t = 800$ sec. The initialized integral balances the selected feedforward. Gains 0, 0.25, 0.5, 0.75 and 1.0 give 11486, 8328, 5601, 5776 and 9094 mm s, with the same values when carried through Reset. Zero feedforward reproduces single-element under this startup procedure. A gain changed during a run can also disturb the existing integral balance, so it is a separate experiment. A half-strength feedforward splits the difference between the mass error and the void error at full load. At 40 % load the same 800-second comparison gives 18536 mm s at half strength against 12782 mm s at unity. The

preferred gain depends on load, which is a limitation of delivering feedforward as valve position. Three-element control instead states the feedforward in flow.

Three-element control is also the only architecture that tolerates gain. At $K_c = 6$ it still works, at 5809 mm s for 58 % of travel, where two-element control at $K_c = 3.5$ has already lost stability at 31635 mm s and 629 % of travel. What it does not do is remove the physics. The remaining 37.4 mm in Table 4.6 is the void collapse itself, and no architecture can prevent a void volume that has to shrink from shrinking.

The lower block of Table 4.6 does not rank the same way, and this chapter reports that rather than selecting the load at which the expected ordering holds. On the load rejection at 40 % load it is two-element control that posts the lowest integrated error at 12782 mm s, against 17121 mm s for three-element control at the full-load gain of 3.5 and 14649 mm s with the gain backed to 2.5. Backing the gain off is worth doing anyway, because it cuts the travel from 157 to 61 % for a 14 % improvement in error, and it is the practical case for gain scheduling this loop on load. What the load rejection does not separate, the pump trip does.

4.6.4 Inner Flow Loop Tuning and the Pump Trip

Scenario 9 breaks the cascade without changing the wiring. The following comparisons are read at $t = 800$ sec, 710 seconds after its load step. Detuning the flow controller to $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, which is slower than the level loop above it, raises the integrated error to 20683 mm s and leaves the level 45.2 mm high at 569 sec. Valve travel falls to 11 %, the lowest figure anywhere in the application, and the low figure is the misleading part. A flow loop that takes two minutes to reach setpoint is a slow valve positioner rather than a flow loop, so the level controller is facing the equal-percentage valve directly again. Nearly all of the recovery arrives with the first step away from the sluggish setting, at 7305 mm s for 0.2 and 30 sec and 6521 mm s for 0.6 and 8 sec. An inner loop does not have to be fast in absolute terms. It has to be several times faster than the loop above it. A fermenter oxygen cascade makes the same inner-loop speed a measured quantity (Volume 3, Chapter 3).

Scenario 10 decides the argument, and it involves no load change. A feedwater pump trips and the discharge pressure falls from 55 bar to 43 bar, so the differential across FCV-101 collapses from 15 bar to about 3 bar and (4.9) passes 45 % of the flow it did. On two-element control the feedwater falls from 25.0 to 10.4 kg/s within ten seconds with the valve still sitting at 69 % open, and nothing in the scheme measures it. The level has to fall far enough for the level controller to respond, and by then the level is falling at more than a millimeter a second. The bottom is −149 mm at 340 sec with the valve on its stop, the integrated error over the 900 sec following the trip is 86511 mm s, and the level is still 62 mm low thirteen minutes later. Single-element control is worse again at −164 mm and 99200 mm s. Three-element control does not encounter it as a level problem at all. FT-102 reports the collapse on the next scan, the inner controller drives the valve from 70 to 95 % inside ten seconds, and the level reaches −7.6 mm before recovering, for 1459 mm s over the same window and about the same valve travel. That is a factor of 59 on the integrated error for a fault that no level tuning addresses.

TRY IT IN THE APP **10 · Feedwater Pump Trip, on Two-Element Control**

Load scenario 10 and press Run. When the pump trips, the feedwater falls from 25.0 to 10.4 kg/s within ten seconds while the valve stays near 69 % open, and the level bottoms at −149 mm at 340 sec. Reload, select three-element control before the trip at 90 sec and run again: the level now reaches only −7.6 mm.

The property the inner loop supplies is coverage rather than speed. Everything that can go wrong between the controller output and the water entering the drum, a pump, a strainer, a stuck stem, worn trim or a suction pressure swing, is inside the inner loop and outside every other scheme.

4.6.5 Transmitter Range Error, Swell and Derivative Action

Scenario 12 ranges the differential-pressure cell at 28 bar and runs the drum at 40 bar. Equation (4.12) then reads 19.7 mm below the true level from the first sample. Denser water makes more pressure per millimeter, so a cell that assumes the water is denser than it is reports fewer millimeters than are there. The controller holds the indication on setpoint exactly, so the water sits 19.7 mm above normal water level and that much high-side trip margin has been spent by a calibration sheet rather than by an engineer. The firing-rate upset moves it in the direction that matters. Drum pressure climbs from 40.0 to 42.4 bar, the error grows to −23.6 mm, and at 600 sec the true level is +26.1 mm while the transmitter reports +2.6 mm. A correctly ranged cell subjected to the same 2.4 bar rise produces −3.5 mm of error, so the residual sensitivity is about 1.5 mm/bar even when nothing is wrong, which is small until a unit moves its pressure across a ten bar range in normal operation. Two consequences follow and neither is a control problem. Feedback cannot see the error, because a loop holding a biased measurement holds the bias and reports nothing, and the trips read the same transmitter [17].

Scenario 11 produces a trip. The boiler idles at 25 % load and the turbine picks up to 120 % at 5 %/s, which is faster than a large machine is permitted to move. Steam leaves the instant the governor opens, the firing rate needs half a minute to catch up, the pressure falls, and every bubble under the surface expands. The level rises to +225 mm at 136 sec, which the transmitter reports as +228.5 mm and is 21.5 mm short of the high-high trip, while feedwater falls short of steam by up to 31.1 kg/s. The controller reads a high level and closes the valve. The voids then stop growing, the mass balance is left with the deficit, and the level falls through zero, through the low alarm and into the low-low trip at 421 sec, bottoming at −254 mm. The excursion that threatens the boiler is therefore the second one and not the swell that caused it.

Two of the three available remedies work and the obvious one does not. Raising the gain to 2.5 avoids the low trip at −159.6 mm and costs 310 % of valve travel by 770 sec, and three-element control at $K_c = 3.5$ holds the falling half to −53.9 mm but reaches the high-high trip at 124 sec on the way to +265 mm, so each remedy trades one trip against the other. Slowing the governor ramp does not help at all. Reducing it from 5 to 2, 1, 0.5 and 0.25 %/s on the shipped tuning brings the indicated peak down from +228.5 mm through 221.6, 198.5 and 138.3 to +86.0 mm, and every one of those runs still reaches the low-low trip, at 432, 450, 494 and 627 sec respectively, because the mass deficit that produces the second excursion is set by the size of the load change and by how long the controller holds the valve shut, and not by how quickly the load arrives. A ramp rate limit protects the high side of the drum and leaves the low side to the controller. The two together divide the work. Three-element control at $K_c = 3.5$ with the ramp cut to 1 %/s reaches neither trip, with the indication peaking at +226 mm and the level bottoming at −29.6 mm.

Scenario 13 applies no disturbance at all and is one of the most expensive runs in the application. A drum surface at full boil is not flat, so the transmitter carries about 6 mm of noise, and derivative action is switched on at $\tau_D = 20$ sec with the filter removed. Fifteen minutes at steady load cost 2091 % of valve travel, twenty full strokes of a stem that is holding a constant level, and finish at 1121 mm s of integrated error against 1018 mm s with the derivative switched off at 208 % of travel. Three additional seeds give 2024 % to 2121 % of travel with derivative against 203 % to 210 % with PI, and derivative adds 227 mm s to 304 mm s of integrated error. Derivative action is wrong on this

Table 4.7 Three-Element Control and the Predictive Controller on the Same Cases, Read at the Lesson Clock

Case	Structure	Peak (mm)	IAE (mm s)	Stem (%)	Command (%)	Limit (%)	HH (mm)	LL (mm)	Last 5 min (mm)	Trip
7, load 100 to 85 %	three	−37.4	5451	20.9	21.3	0	238	212	14.3	none
	MPC	−22.2	3101	14.5	31.0	0	246	228	8.6	none
8, load 40 to 25 %	three	−76.5	14595	136.4	138.4	0	198	174	51.8	none
	MPC	−44.9	6455	42.2	50.5	0	242	205	7.6	none
9, detuned flow loop	three	+45.2	13501	8.1	8.2	0	205	218	45.2	none
	MPC	−29.5	6257	11.8	15.0	0	232	220	17.6	none
10, pump trip	three	−7.6	1083	31.1	36.8	0.6	246	242	1.6	none
	MPC	−5.7	554	32.6	34.4	10.7	246	244	0.5	none
11, 25 to 120 % at 5 %/s	three	+264.7	44448	184.7	191.1	15.7	−19	195	53.9	HH
	MPC	+217.0	25600	108.5	130.5	14.0	30	232	16.8	none
14, 40 to 100 % at 2 %/s	three	+167.6	27202	152.6	154.7	8.3	80	199	49.9	none
	MPC	+122.1	14844	102.7	116.7	5.7	127	238	11.4	none
15, suggested settings	three	−154.2	24868	69.7	74.4	0	214	97	14.8	none
	MPC	−83.4	14540	67.3	161.8	0	229	168	14.7	none
15, as shipped	single	+285.7	66180	211.1	211.7	13.8	−45	72	269.9	HH
	MPC	+106.1	41337	28.3	30.4	44.2	141	160	106.1	none

Peak and IAE use the true level, IAE being the scorecard value. Stem is the travel of the valve stem x and Command that of the valve command x_{cmd}, the scorecard's manipulated-variable travel. Limit is the share of samples with the valve command at 0 or 100 % or the flow setpoint at 0 or 160 %, whereas the scorecard's saturation counts the valve command alone and reads 0 on every MPC row but the pump trip. HH and LL are the smallest margins of the transmitter reading to the trips, and Last 5 min is the largest true-level error over the final 300 sec. Each MPC run made one plan per 2 sec with none rejected.

loop for a second reason beyond the noise. It acts on the direction the measurement is moving, and on a process with an inverse response the measurement moves the wrong way for the first minute, so an accurate derivative drives the valve away from where the mass balance needs it.

4.6.6 Graded Challenge

The challenge, scenario 15, applies a load increase to 115 % at 60 sec, a feedwater heater trip at 220 sec and a fifty-point load rejection to 65 % at 380 sec, starting on single-element control at $K_c = 1.0$ with the flow loop detuned to $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, and with the governor ramp at 2 %/s. The target is neither trip reached and the level inside ±30 mm for the last five minutes. As shipped the boiler trips, swinging to +285.7 mm and reaching the high-high trip at 559.5 sec after having been 181.5 mm low, for 211 % of travel. Two-element control at $K_c = 1.6$ trips the other way, reaching the low-low trip at 486.6 sec and −254.4 mm, retaining the shipped $\tau_I = 700$ sec after changing the architecture and gain. Three-element control at $K_c = 2.0$, $\tau_I = 500$ sec, $K_{ff} = 1$ and an inner loop at 0.6 and 8 sec rides all three events, with 24868 mm s of error at $t = 900$ sec and a worst excursion of −154.2 mm, 70 % of travel and the level inside 14.8 mm for the last five minutes. The middle event carries the mechanism. Feedwater at 110 °C instead of 160 °C raises the subcooling in (4.4) from 408 to 622 kJ/kg, so every kilogram entering the drum condenses 0.363 kg of steam instead of 0.238 kg. Nothing about the flows changed. The shrink is permanently 52 % stronger, and the level is the only measurement in the scheme that reports it.

4.6.7 Comparison With the Predictive Controller

Table 4.7 sets the MPC against three-element control on six scenarios and the challenge, each read at the clock on which its lesson is credited, 600 sec or 900 sec on the challenge, so its integrals do not match those of Table 4.6. Three-element control keeps each case's own outer tuning, switched to three-element before the run in scenarios 10 and 11, and the suggested settings on the challenge. The MPC keeps the settings of Section 4.5.3 throughout, both act through the case's own flow loop and read the same measured steam flow, and both see the same events, instruments and noise seed. The challenge is also run as shipped against the MPC selected as the only change.

Scenario 14 is the case added with the controller. The boiler idles at 40 % load, and from 60 sec the turbine picks up to 100 % at 2 %/s. The void contribution rises from −180.9 mm to −0.5 mm at 600 sec, so with the indication on setpoint the inventory has to fall by as much. Three-element control raises its flow setpoint with the steam flow, to 81.2 % at 89.5 sec, and drives it to zero only once the

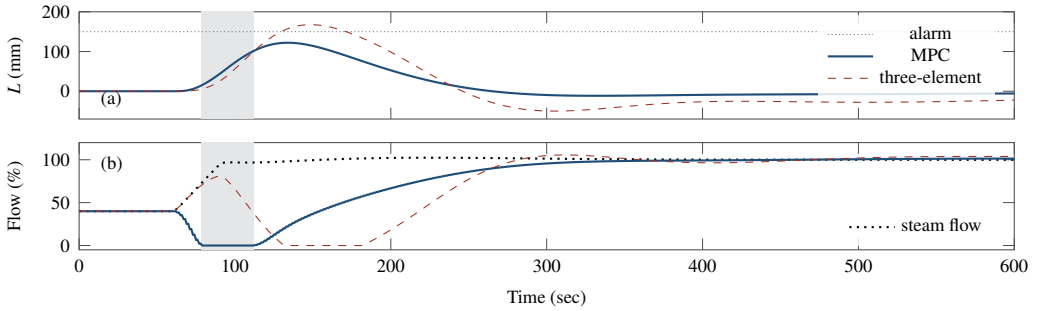

Figure 4.4 Scenario 14 under the predictive controller and under three-element control, with the same seed, over the ten minutes to the lesson clock. The turbine picks up from 40 to 100 % load from 60 sec at 2 %/s, and the shaded span from 78 to 112 sec is the interval in which the MPC holds the feedwater flow setpoint at zero. (a) True drum level under the MPC (thick solid) and under three-element control (dashed), and the 150 mm deviation alarm (thin dotted). (b) Feedwater flow setpoint under each structure, in the same line styles, and the measured steam flow of the MPC run (thick dotted), in percent of rated flow, the latter within 3.1 points of that of the three-element run.

swell reaches the transmitter, from 131.5 to 181.6 sec, after which the drum must be refilled. The MPC forecasts the void growth from the same steam flow and holds its setpoint at zero from 78 to 112 sec, while the steam flow is still rising, so the inventory falls before the voids arrive, as shown in Figure 4.4. The level peaks at +122.1 mm at 133.9 sec against +167.6 mm at 148.4 sec, beyond the 150 mm alarm, the transmitter comes within 127 mm of the high-high trip against 80 mm, and the trough that follows is −11.4 mm against −49.9 mm. The feedforward encodes only the steady mass balance, while the model also carries the void growth that occupies the drum first.

On scenarios 7 to 11 the MPC also lowers both the peak and the IAE, by factors of 1.7 to 2.3 in IAE. At 40 % load it does so with 42.2 against 136.4 % of stem travel, where the full-load gain of 3.5 suits three-element control least. On the pump trip the valve command sits at its upper stop for 64 sec, the 10.7 % of the Limit column, while the setpoint averages 6.7 points above what the tripped pump delivers. On the fast pickup of scenario 11, three-element control reaches the high-high trip at +264.7 mm, while the MPC peaks at +217.0 mm with its band slack reaching 62.8 mm, a plan that priced the band rather than obeyed it. On the detuned loop of scenario 9 it more than halves the IAE with more stem travel, 11.8 against 8.1 %.

The stem column understates what the MPC asks of the valve. Each 2 sec setpoint step draws an immediate command move of $K_{c,F}$ times the step from Equation (4.15), partly returned as the flow follows, and the 4 sec stem filters most of it out. On scenario 14 the command travels 116.7 % against 102.7 % on the stem, still below three-element control's 154.7 %. On the challenge with the suggested flow loop it travels 161.8 % against 74.4 %, more than twice as far, although the two setpoints travel 217.9 and 204.4 % and the two stems 67.3 and 69.7 %, and it also travels further on scenarios 7 and 9.

With the suggested settings both structures pass the challenge, the MPC at 14540 mm s against 24868 mm s. Selected as the only change to the challenge as shipped, the MPC reaches no trip but fails. From 564.1 sec its setpoint sits at zero while the detuned flow loop, $K_{c,F} = 0.06$ and $\tau_{I,F} = 120$ sec, closes the valve command from 65.1 to 54.0 %, 1.98 points a minute, and the level ends 106.1 mm above setpoint against the 30 mm allowed. No setpoint in range could have closed the valve faster, so the remedy is the inner loop, as for three-element control, and with the suggested flow tuning the same MPC passes. In Chrome on scenario 14 the mean solve took 238 to 279 µs over five runs and the largest 2.6 to 3.7 ms, measured on a 20-core host at a load average near 2, or 0.014 % of a core at real time.

✦ DID YOU KNOW? No Trip, Still a Fail

From the simulator: on the challenge as shipped, selecting the predictive controller as the only change cuts the stem travel to 28.3 % against 211.1 % for the shipped single-element loop and reaches no trip. The largest level error over the last five minutes is still 106.1 mm against the 30 mm allowed (Table 4.7).

4.6.8 Numerics of the Running Application

Three clocks run, as Section 1.6.1 describes, the 0.05 sec Runge-Kutta step for the physics, the 0.5 sec period for both controllers and ten historian samples per simulated second. A scheduled scenario event is applied at the start of the step that begins at its time, before the controller execution due on that step. Within an execution the order is fixed, the governor ramp first, then the instrument chain with its density error, bias, noise and lag, then the feedforward, then the Kalman filter of Section 4.5.3, then either the controllers or the manual tracker, then the integration, then the latched trips and the travel counters, then the outputs and the log. No controller reads a state directly. The export carries 87 columns at ten samples per simulated second, including the true level, the indicated level and the level the controller acted on as three separate columns, the water and void contributions apart from each other, the feedforward and trim split, and the predictive controller's status, slack, solve time and estimates, which is enough to reproduce every number in Section 4.6 away from the browser. The application also runs nineteen checks on itself, covering the exact equilibrium of the design condition, the sign of each flow against its valve, the collapse of the sub-surface steam under extra feedwater, the expansion of the voids under falling pressure, the growth of the residence time at reduced firing, the exactness of (4.12) at its calibration pressure, the valve characteristic and the unity feedforward. One of them runs the plant open loop for 520 sec and asserts that the level goes down first and up afterward. Seven cover the predictive controller, among them its limits on every scenario, its source boundary, a least-squares check on one unconstrained move and a bumpless forced rejection.

Figure 4.5 catches scenario 7 in the middle of its recovery and carries two results worth stating. The first is the integrated absolute error, 3256 mm s at $t = 300$ sec against the 6521 mm s the same run reports 710 sec after the load change, which is what accumulation after a transient looks like and is the reason every integral in this chapter names its reading instant. The second is the differential-pressure cell error, +0.6 mm on a correctly ranged transmitter, because the drum pressure has fallen 0.42 bar below the pressure the cell was ranged at and Equation (4.12) is exact at one pressure only.

4.7 Discussion

Drum level has an unusual property among control loops. It is a frequent cause of unit trips, and the cause is almost never the tuning [1, 3]. The application is built so that the numbers behind that statement can be produced rather than asserted.

Nothing in scenario 10 is mistuned. The controller does exactly what its measurements support and takes the level to −149 mm because two-element control contains no measurement of feedwater flow. Nothing in scenario 12 is mistuned either. The controller holds the indication perfectly while the water sits 19.7 mm above where it is reported to be, and every self-diagnostic on the transmitter passes. An investigation that opens with the tuning constants finds nothing wrong with them in either case, which is a large part of why these events repeat. The general shape is the most dangerous class of automation fault, a system working correctly on information that is wrong, with the evidence available only outside the loop. In a boiler house that evidence is the gauge glass, which is an independent measurement with a different failure mode.

Trip settings are a second place where a control engineer's numbers should be informing a safety decision. The ±250 mm trips protect tubes and turbine blades and their values come from the boiler maker and the turbine maker, but whether they are reachable is a control question with a measurable answer. Scenario 11 reaches the low-low trip from a swell that stops 21.5 mm short of the high-high trip, and the challenge reaches the high-high trip from a tuning that looks unremarkable on a faceplate.

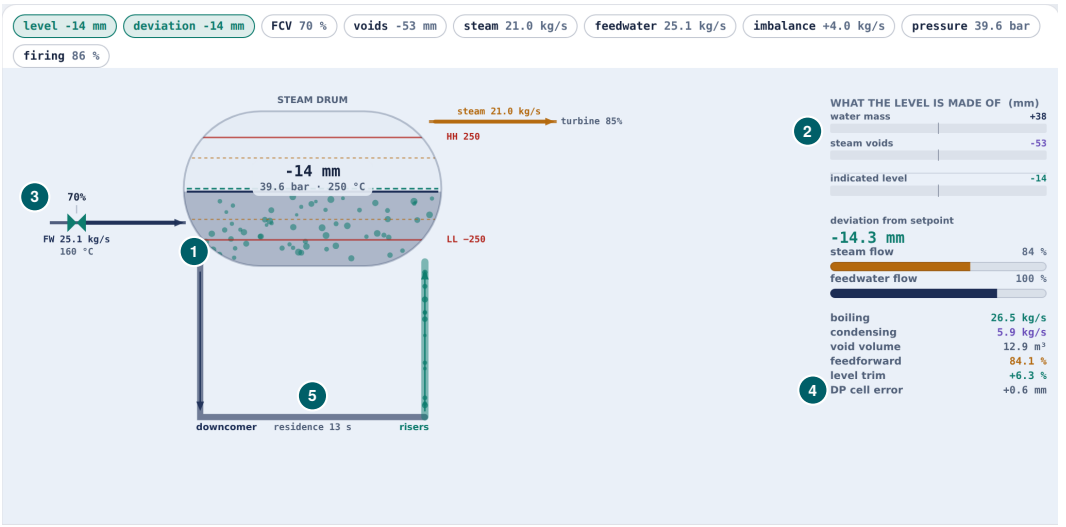

Figure 4.5 Three-element control riding the fifteen-point load rejection of scenario 7, caught at $t = 300$ sec with the disturbance 150 sec old, in the plant view of the app. The steam flow has ramped to 84.1 % of rated and the feedwater is at 100.2 % following an inner setpoint of 99.5 %, so the drum is gaining 4.03 kg/s while the level reads −14.3 mm. The level is decomposed into a water contribution of +38.3 mm against a void contribution of −52.6 mm, recovering from −59.7 mm at 265 sec. The drum pressure has fallen back to 39.58 bar from 40.59 bar at 202 sec. The crop keeps the plant view, whose readouts round these values, and cuts the trends, the control panel and the scorecard. Callouts: (1) steam bubbles held below the water surface, (2) the level split into its water and void terms, (3) feedwater valve FCV-101 with the feedwater flow and temperature, (4) the error of the differential-pressure cell, (5) the downcomer and the risers with their residence time.

The useful product is a table of worst excursion against load, ramp rate and architecture, which states a trip setting and an operating envelope in the same units.

Valve travel belongs on the same page as the error. Scenario 13 spends 2091 % of travel and finishes 103 mm s worse on integrated error than the same loop with the derivative switched off, and scenario 9 posts the best travel figure in the application, 11 %, alongside one of the worst errors. Read on its own, neither number says what happened.

The architecture is the decision with the longest consequence. Moving between the rungs of Figure 4.2 is not a tuning session. It is wiring, a flow element in the feedwater line, a second transmitter and a second controller in the distributed control system. A plant commissioned on two-element control because it was going to be a base-load machine, and now cycled because the grid asks it to be, is running the wrong architecture and no amount of retuning fixes it [18]. Scenario 4 against scenario 8 is that argument with numbers on it, and the judgment that has to be made at design time is a forecast about how the plant will be operated in fifteen years.

> **OPEN CHALLENGE** **Level Loops Tuned for Full Load, Run at Part Load**
>
> Plants asked to follow the load find level loops tuned for another duty. Feedwater heater level controls are typically tuned for full load and may not hold the level at reduced loads, which wears heaters, drain lines and valves, causes water hammer and, in extreme cases, trips the turbine [25]. The effective water level of a steam generator also changes with load, and one fleet's study of deeper load following lists steam generator level control among the items to assess [25]. Progress would be level control whose tuning or structure follows the load, the gain schedule that Section 4.6 argues for.

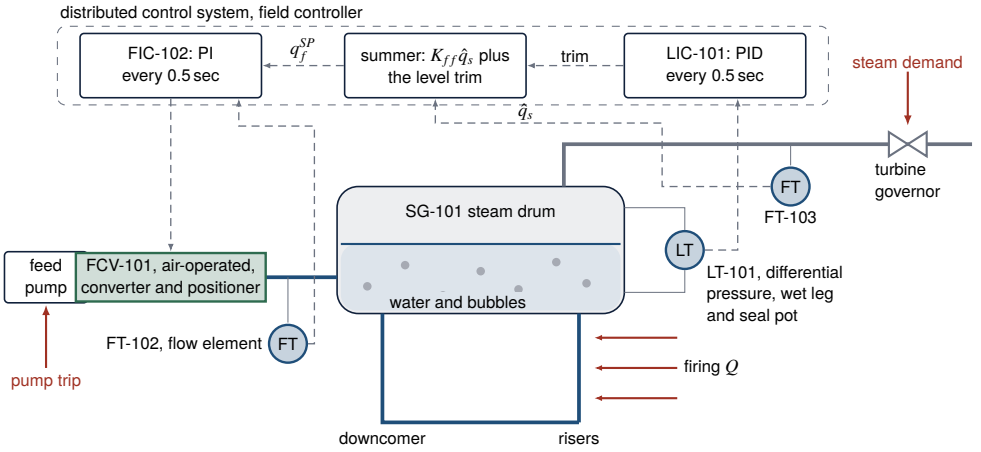

Figure 4.6 The three-element loop of Figure 4.1 with the hardware of a plant. A differential-pressure transmitter (LT) on two drum taps, its upper line a condensate-filled wet leg, reads the level, and flow transmitters (FT) on flow elements read the feedwater and the steam. In the field controller LIC-101 writes a trim that a summer adds to the steam flow feedforward, and FIC-102 holds the resulting feedwater setpoint with the air-operated valve FCV-101, both every 0.5 sec. The turbine governor follows the steam demand and is not moved by the loop. Thick solid lines carry water (blue: the feed line, the downcomer and the risers) and steam (gray: the line to the turbine), thin dashed arrows signals, and red arrows the disturbances.

4.8 Technology in Practice

The tags of Figure 4.1 are, in a plant, a differential-pressure level transmitter on the drum, two flowmeters, an air-operated feedwater valve and the controllers of a distributed control system. Figure 4.6 draws the three-element loop with that hardware in place of the schematic's symbols.

What Measures It

The three elements are three measurements. Drum level comes from a differential-pressure transmitter across two taps on the drum, the most common level method in nuclear plants [26]. Feedwater and steam flow are most often measured the same way, as the pressure drop across an orifice, a Venturi tube or another flow element, a square-root signal that the transmitter or the signal processing makes linear [26]. The chapter carries LT-101 as Equation (4.12) behind a 3 sec lag and both flow transmitters behind 2 sec lags.

> **TECHNOLOGY** **Sensor: Drum Level From a Differential Pressure and a Wet Leg**
>
> Drum level is read as the difference in pressure between a tap in the water and a tap in the steam. Steam condenses in the upper line and fills it, and the constant head of this wet leg is canceled by an intentional zero shift at calibration. Seal pots are standard on steam drums, because a wet leg that has lost water refills slowly and the level reads wrong until it does [27]. The transmitter sends 4 to 20 mA, and the live zero at 4 mA reveals an open circuit [26]. LT-101 adds the density error of Equation (4.12) and a 3 sec lag.

What Moves It

The final element is FCV-101, an equal-percentage valve whose stem follows its command as a 4 sec lag. In pressurized water reactors the feedwater regulating valves are air operated, and they are the valves to which smart positioners have most often been fitted [28]. In safety-related designs, solenoid valves between the positioner and the actuator can dump the air so that the spring returns the valve to its fail-safe position [28]. The inner flow loop shields the level controller from loads in the feedwater

supply and keeps the valve's nonlinear response its own problem [27], which is what the pump trip of scenario 10 measures.

> **TECHNOLOGY** **Final Element: The Air-Operated Feedwater Valve**
>
> In an air-operated valve, air on a diaphragm moves the stem against a spring, and the spring moves it back when the air is released. A converter turns the controller's 4 to 20 mA into 3 to 15 psig, and a positioner adjusts the actuator air until the valve reaches the demanded position. It reduces the hysteresis that packing friction causes and makes up for changes in spring rate and in the pressure on the stem [28]. A sticking valve still moves in steps and ignores small reversals of its signal, and tuning cannot fully remove that [27]. FCV-101 has no stiction and strokes as a 4 sec lag.

Where the Controller Runs

Loops of this kind run in field controllers with cycle times from milliseconds to minutes, under a supervisory layer that sends them setpoints [29]. LIC-101 and FIC-102 execute every 0.5 sec, the scan of a distributed control system, above a 0.05 sec physics step. In the classic three-element arrangement a summing block adds the steam flow signal and the level controller's output, the trim, and passes the sum as a remote setpoint to the feedwater flow controller [27]. The chapter delivers the same sum as the output bias of the level controller, so its integral carries only the trim, and the predictive controller of Section 4.5.3 writes to the same setpoint.

> **IN PRACTICE** **Two Independent Ways to See the Water**
>
> Boiler rules ask for a second, independent view of the drum level. United States rules for marine boilers, which modify paragraph PG-60 of Section I of the boiler and pressure vessel code, require two independent means of indicating the water level, connected directly to the drum, one of them a gauge lit from the emergency supply. The gauge glasses must stay in service whenever the boiler is steaming [30]. The glass is read with care: its water cools and becomes denser than the drum water, so the two columns balance at different heights [27].

What Is Modernizing It

Three developments reach this loop. Digital valve positioners calibrate themselves, diagnose air leakage, friction and dead band, and allow partial stroke tests on line. On the main feedwater valves of one unit they allowed the steam generator levels to be tuned to more stable control [28]. Plants that raise their power rating fit more accurate feedwater flow measurement, because the energy balance that confirms the reactor's thermal power rests on the feedwater flow, temperature and pressure [26]. Model-based control of this level has been proposed since at least 1979 [19], and predictive control has been applied to the steam generator level of a pressurized water reactor [31], the structure that Section 4.6.7 sets against three-element control.

4.9 Limitations

The application demonstrates inverse response arising from a mass and energy balance, a right-half-plane zero that moves with operating point, what each rung of the architecture ladder buys and does not buy, cascade as a loop closed around an actuator, and the drift of a differential-pressure level measurement with operating pressure. It is not a boiler analysis tool.

The circulation is one number, so the low-load region where the application is most interesting is also where the model is least complete, because a real natural-circulation boiler at very low load can circulate unstably and that failure mode is absent entirely. The fuel side is a feedforward and a proportional trim rather than a boiler master, and there is no steam line volume and no superheater, so neither firing rate limits nor the interaction between steam pressure control and level control can be studied. Blowdown, attemperation, the economizer and the drum internals are absent, and separator carryover is represented only by the high-high trip.

Three control-side limitations matter for how the results should be read. Four structures exist, with no automatic transfer to single-element control at low load, although real plants commonly make that transfer below the load at which the feedwater flow measurement stops being trustworthy. The valve has no stiction, no backlash and no positioner of its own, which matters because feedwater valve stiction is a common real cause of level cycling and is exactly the fault an inner flow loop half conceals. There is one feedwater valve and one pump characteristic, while real feedwater systems carry a startup valve, variable-speed drives, minimum-flow recirculation and pumps in parallel, all of which change what the inner loop is controlling. Perfect repeatability is itself a property of the implementation rather than of a plant, and the fixed step, the fixed scan and the fixed seed are the reason for it. Where each individual assumption breaks is set out in Section 4.3.4.

The predictive controller's model has no pressure state, so the immediate expansion of the bubbles when the pressure falls reaches it only through the filter's swell estimate. It assumes the design feedwater temperature, so the heater trip is a model error absorbed by $\hat{d}$, and it ignores the density error of Equation (4.12) and any parameter changed since the last reset. Its single 510 sec lag for the detuned inner loop predicts 2 % of a setpoint step after 10 sec where the loop delivers 15 % through its proportional part. Acting only through the flow setpoint, it fails the challenge as shipped, and its 2 sec steps cost more command travel than three-element control on three cases.

> **OPEN CHALLENGE** **One Automatic Level Controller From Startup to Full Load**
>
> In pressurized water reactors, unplanned trips caused by steam generator level control have been significant contributors to lost availability, and conventional feedwater control cannot cover the whole range from zero to full load. At low power the wrong-way response is strongest and the flow measurements are often unreliable [31]. A 2001 laboratory study of a robust model-based controller reported only preliminary results and called for more work [31]. Progress would be one controller that holds the level from startup to full load with no manual operation.

4.10 Conclusion

This chapter studied a natural-circulation steam generator as a laboratory for a nonminimum-phase plant whose zero depends on the operating point. The simulator integrates five states, the drum water volume, the sub-surface steam mass, the drum pressure, the absorbed heat and the feedwater valve stem, with the indicated level written algebraically as a water term plus a void term so that shrink and swell emerge from the balances rather than from a numerator coefficient. An open-loop feedwater step measures the zero directly. Trough depth scales with the step, from −1.47 mm to −15.21 mm across a factor of ten, while all four crossings back through the starting level fall between 38 and 39 sec, and changing the load leaves the trough near −3.7 mm while moving the crossing from 38 sec at full load to 65 sec at a quarter. On a fifteen-point load rejection the three architectures give −29.8 mm, −50.2 mm and −37.4 mm of peak deviation for 11486, 9094 and 6521 mm s of integrated error, and the same 3.75 kg/s of steam is 2.2 times worse at 40 % load while the same *relative* step is slightly easier there. A feedwater pump trip separates the architectures by a factor of 59 for the same valve travel, and a cell ranged at 28 bar while the drum runs at 40 bar holds an exact indication over water sitting 19.7 mm too high. Slowing the governor ramp from 5 to 0.25 %/s cuts the swell peak from 228.5 to 86.0 mm and does not prevent the low-low trip that follows it. A predictive controller that forecasts the voids from the measured steam flow holds the pickup from 40 to 100 % load to +122.1 mm against +167.6 mm under three-element control, and like three-element control it needs a properly tuned inner flow loop to pass the challenge. A fuller model would add valve stiction, a steam pressure loop, an automatic low-load transfer between architectures and a distributed riser model that admits circulation instability. A nonlinear predictive controller on the five-state balances could hold the two trips as explicit constraints, and an estimator over Equations (4.2) to (4.6) could reconstruct the water inventory and the transmitter density error from the level, the two flows and the pressure. The steam generator measures the price of a right-half-plane zero that moves with the load.

KEY POINTS

- A drum level is water plus the steam bubbles below the surface, so cold water lowers the reading before it raises it.
- The depth of the wrong-way dip scales with the step, while its duration grows from 38 sec at full load to 65 sec at a quarter.
- An inner feedwater flow loop covers a pump trip that no level tuning reaches, a factor of 59 on integrated error.
- A loop holding a biased measurement holds the bias, here 19.7 mm of water from a cell ranged at 28 bar.

Where else this applies. Three-element control also holds the vessel level of Chapter 6, and a zero that moves with load limits the boost converter of Volume 3, Chapter 9.

References

[1] M. V. Kothare, B. Mettler, M. Morari, P. Bendotti, and C.-M. Falinower, "Level control in the steam generator of a nuclear power plant," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 1, pp. 55–69, 2000.

[2] M. G. Na, K. C. Kwon, C. S. Ham, and Z. Bien, "A study on water level control of PWR steam generator at low power and the self-tuning of its fuzzy PI controller," *Fuzzy Sets and Systems*, vol. 74, no. 1, pp. 43–51, 1995.

[3] D. Lindsley, *Power-Plant Control and Instrumentation: The Control of Boilers and HRSG Systems*. London, UK: Institution of Electrical Engineers, 2000.

[4] S. G. Dukelow, *The Control of Boilers*, 2nd ed. Research Triangle Park, NC: Instrument Society of America, 1991.

[5] 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.

[6] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. Chichester, UK: John Wiley & Sons, 2005.

[7] K. Iinoya and R. J. Altpeter, "Inverse response in process control," *Industrial & Engineering Chemistry*, vol. 54, no. 7, pp. 39–43, 1962.

[8] W. L. Luyben, "Identification and tuning of integrating processes with deadtime and inverse response," *Industrial & Engineering Chemistry Research*, vol. 42, no. 13, pp. 3030–3035, 2003.

[9] K. L. Chien, E. I. Ergin, C. Ling, and A. Lee, "Dynamic analysis of a boiler," *Transactions of the ASME*, vol. 80, no. 8, pp. 1809–1818, 1958.

[10] K. J. Åström and K. Eklund, "A simplified non-linear model of a drum boiler-turbine unit," *International Journal of Control*, vol. 16, no. 1, pp. 145–169, 1972.

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

[12] H. Kim and S. Choi, "A model on water level dynamics in natural circulation drum-type boilers," *International Communications in Heat and Mass Transfer*, vol. 32, no. 6, pp. 786–796, 2005.

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

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

[15] K. L. Poon, "Effect of rapid steam take-off on natural circulation and water level in boilers," *Journal of Engineering for Power*, vol. 97, no. 4, pp. 645–654, 1975.

[16] C. Maffezzoni, "Boiler-turbine dynamics in power-plant control," *Control Engineering Practice*, vol. 5, no. 3, pp. 301–312, 1997.

[17] R. J. M. Welsh, "PWR steam generator water level measurement systems," *Measurement and Control*, vol. 26, no. 5, pp. 151–155, 1993.

[18] M. A. Gonzalez-Salazar, T. Kirsten, and L. Prchlik, "Review of the operational flexibility and emissions of gas- and coal-fired power plants in a future with growing renewables," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 1497–1513, 2018.

[19] J. J. Feeley, "Natural circulation steam generator model for optimal steam generator water level control," in *American Nuclear Society Meeting, San Francisco, CA, November 1979*. Idaho National Engineering Laboratory, Jun. 1979, report CONF-791103-30, OSTI 5962048.

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

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

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

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

[24] 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.

[25] 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.

[26] ——, *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.

[27] 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.

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

[29] 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.

[30] U.S. Coast Guard, "Water-level indicators, water columns, gauge-glass connections, gauge cocks, and pressure gauges (modifies PG-60)," Code of Federal Regulations, Title 46, Section 52.01-110, revised as of October 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[31] A. M. H. Basher and J. March-Leuba, "Development of a robust model-based water level controller for U-tube steam generator," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2001/166, Sep. 2001.

4.11 Exercises

4.1 ⌕ Concept check · 1 min

In MANUAL at full load the feedwater valve is opened from 70 to 78 %, and the feedwater rises above the steam flow. The drum level first falls. Which explanation is right?

(a) The valve strokes for 4 sec, and the feedwater flow dips while the stem is moving

(b) The level transmitter lag of 3 sec makes the reading run behind the true level

(c) The subcooled feedwater condenses steam below the surface, and the volume lost exceeds the water gained

(d) The extra feedwater raises the drum pressure, which compresses the bubbles below the surface

4.2 ⌕ Concept check · 2 min

Open-loop feedwater steps of 10, 25, 50 and 100 % are made at full load from the same equilibrium. What does Table 4.3 show about the inverse response?

(a) The trough and the crossing time both grow with the step, so a smaller move gives a faster recovery

(b) The trough is the same for every step, and the crossing time grows with the step

(c) The trough grows with the step, and the crossing time falls because a larger flow refills the drum sooner

(d) The trough grows in proportion to the step, and the crossing time stays near the same value

4.3 Concept check · 2 min

On the load rejection from 100 to 85 %, two-element control at $K_c = 1.6$ dips to −50.2 mm, deeper than the −29.8 mm of single-element control, and has no overshoot. Why is the dip deeper?

(a) The feedforward cuts the feedwater to the new steam flow, so no extra water offsets the void collapse

(b) The feedforward carries the noise of the steam flow transmitter into the valve

(c) The valve-position feedforward is most in error at full load, where it asks for too little valve

(d) The feedforward doubles the effective gain of the level controller, which then overshoots downward

4.4 Theory and calculation · 6 min

Use the property fits of Equation (4.11) at the design pressure of 40 bar.

(a) Compute the subcooling $h_w - h_{fw}$ of 160 °C feedwater and the mass of steam that one kilogram of it condenses on its way to saturation. *Checkpoint:* 408.3 kJ/kg and 0.238 kg.

(b) Compute the ratio ρ_w/ρ_s, and from it the ratio of the steam volume lost to the water volume gained for each kilogram of feedwater. *Checkpoint:* 39.8 and 9.5.

(c) A feedwater heater trip drops the feedwater to 110 °C. Repeat part (a) and state how much stronger the shrink is per kilogram. *Checkpoint:* 622.4 kJ/kg, 0.363 kg, 52 % stronger.

4.5 Theory and calculation · 6 min

At any steady load the drum of SG-101 sits at its design pressure of 40 bar with feedwater equal to steam. Take $\tau_{res,0} = 12$ sec, $q_{s,0} = 25$ kg/s, $k_{void} = 0.526$, $A_d = 20\,\mathrm{m}^2$ and $\rho_s = 20.09\,\mathrm{kg/m}^3$, and let the firing fraction Q/Q_0 equal the load fraction.

(a) Use Equations (4.4) and (4.5) to show that the steady sub-surface steam is $m_{sd} = \tau_{res}\, q_s$, and compute τ_{res} at 85, 40 and 25 % load. *Checkpoint:* 12.7, 16.2 and 18.1 sec.

(b) Compute the void contribution to the level, the second term of Equation (4.1), at 85, 40 and 25 % load. *Checkpoint:* −38, −181 and −244 mm.

(c) Compare the change in void contribution for the load rejection from 100 to 85 % with that from 40 to 25 %. Both remove 3.75 kg/s of steam. *Checkpoint:* 38 mm against 63 mm.

4.6 Theory and calculation · 7 min

Level transmitter LT-101 has taps $H = 1.2$ m apart and reports level by Equation (4.12). In scenario 12 it is ranged at $p_{cal} = 28$ bar while the drum runs at 40 bar.

(a) Compute $\rho_w - \rho_s$ at 28 and at 40 bar from Equation (4.11). *Checkpoint:* 816.5 and 778.5 kg/m^3.

(b) The controller holds the indication $\hat{L}$ at zero. Invert Equation (4.12) for the true level L. *Checkpoint:* +19.7 mm.

(c) A correctly ranged cell, $p_{cal} = 40$ bar, sees the drum pressure rise to 42.4 bar with the true level at zero. Compute its error and the sensitivity per bar. *Checkpoint:* −3.5 mm, about 1.5 mm/bar.

4.7 App challenge · 7 min

Load scenario 1, Drive the valve by hand (MANUAL), and change no setting. The feedwater valve steps from 70 to 78 % at 60 sec. Run to 160 sec at the default speed and seed, and export the CSV file.

(a) In column `level_m` (true level, in meters), find the lowest level after the step and the time after the step at which it occurs. *Checkpoint:* −3.95 mm, within 0.05 mm, 19.4 sec after the step, within 0.5 sec.

(b) In the row at 79.4 sec, read the water and void contributions, columns `level_water_m` and `level_void_m`. *Checkpoint:* +7.16 mm and −11.11 mm, each within 0.05 mm.

(c) Find the time after the step at which `level_m` passes back through zero, and compare it with the crossing at a quarter load in Table 4.3. *Checkpoint:* 38.1 sec, within 0.2 sec, against 65 sec.

4.8 App challenge · 9 min

Run the fifteen-point load rejection on the three architectures, each scenario as shipped, at the default seed and a speed of 10×: scenario 3, Single-element on a load rejection, and scenario 5, Two-element: the steam-flow feedforward, each to 800 sec, and scenario 7, Three-element control, to 860 sec. Each reading is then 710 sec after the load change.

(a) For each run read the IAE tile of the scorecard, in m s, and the Valve travel tile. *Checkpoint:* 11.486, 9.094 and 6.521 m s, each within 0.01 m s, for 36, 39 and 21 %.

(b) From each exported file, find the lowest true level, column `level_m`, after the load change. *Checkpoint:* −29.8, −50.2 and −37.4 mm, each within 0.1 mm.

(c) Explain why three-element control cannot bring the trough below about 38 mm whatever its gain.

4.9 App challenge · 8 min

Load scenario 10, Feedwater pump trip, on two-element control. The pump trips at 90 sec and the discharge pressure falls from 55 to 43 bar. Use the default seed and a speed of 10×, and run each case to 990 sec, 900 sec after the trip.

(a) Run the scenario as shipped. From the file, read the feedwater flow `fw_flow_kgps` and the valve stem `fw_valve_pct` at 100 sec, and the lowest true level `level_m` and its time. Read the IAE tile at 990 sec. *Checkpoint:* 10.4 kg/s at 69 % open; −149 mm at 340 sec; IAE 86.51 m s, within 0.05 m s.

(b) Reload the scenario, set Control architecture to Three-element before Run, and repeat. Compute the ratio of the two IAE readings. *Checkpoint:* −7.6 mm, IAE 1.459 m s, within 0.01 m s; a ratio of 59.

4.10 Application area · 3 min

After the pump trip of scenario 10 the discharge pressure is 43 bar against a drum at 40 bar. The valve passes $q_{f,\max} = 58.6\,\mathrm{kg/s}$ wide open at the design differential of 15 bar. What is the most feedwater the valve can now deliver, and what does that mean for a level loop at full load? Name one practical way to keep authority after such a trip.

4.11 Application area · 3 min

A unit was commissioned on two-element control as a base-load machine and is now cycled down to 40 % load every night. Give two results of this chapter that argue for a different architecture at low load, and one behavior of a real boiler at very low load that the model cannot show.

4.12 Review the agent · 5 min

An AI assistant reviews scenario 11, in which the turbine picks up from 25 to 120 % at 5 %/s on single-element control, and replies as follows. “The danger is the swell. The indicated level peaks at +228.5 mm, only 21.5 mm below the high-high trip. Limit the governor to 0.5 %/s and the peak falls to +138.3 mm, so the drum rides through the pickup with no trip and no change to the controller.” Find the flaw in the proposal and say how it could be checked.

CHAPTER 5

Nuclear Steam Dump Simulator for Feedforward and Override on a Load Rejection

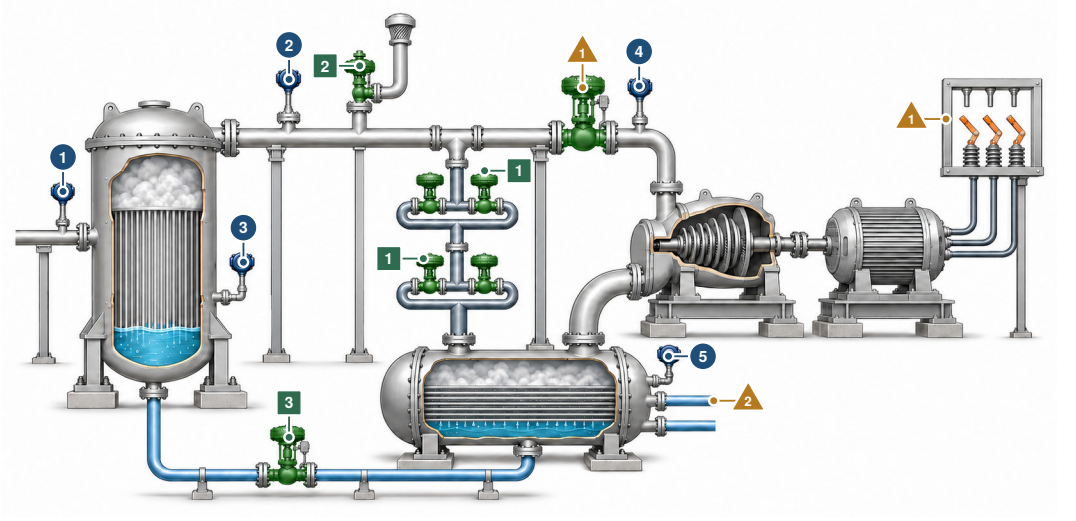

The controller moves 1 steam dump valves to the condenser 2 relief valve to the air 3 feedwater valve

The controller reads 1 average coolant temperature 2 steam pressure 3 steam generator water level 4 steam pressure inside the turbine 5 condenser pressure

What disturbs it 1 breaker opens, turbine valve shuts 2 cooling water temperature

When the breaker opens and the turbine valve shuts, the steam has to go through the dump valves into the condenser, and the controller learns of it first from the pressure inside the turbine. Illustration, not to scale.

IN PLAIN TERMS Where the Steam Goes When the Grid Is Lost

The main idea. A nuclear plant makes steam that spins a turbine, which drives a generator. When the link to the power grid suddenly opens, the turbine valves shut, but the reactor keeps making steam with nowhere to go.

The challenge. The pressure reaches the safety valves in under fifteen seconds, too fast for a thermometer in the water to warn in time. So valves that send the steam to a cooler that turns it back into water must open at once, on a signal from the turbine.

Why it matters. Each time the safety valves open, clean steam is lost to the air and the valves must be tested.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the fact box of Section 5.1, the history box of Section 5.2, the try-it boxes of Sections 5.4 and 5.6, the fact box of Section 5.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	25 min	The first page, Section 5.6, Section 5.7, Section 5.8 and the open challenge

The control question When the generator breaker opens, the steam a pressurized water reactor is still making has nowhere to go, and the secondary reaches its first safety valve in under fifteen seconds. Can a temperature loop on the steam dump be tuned to hold that margin, or does the dump need a signal from the turbine and overrides on its constraints?

Learning objectives

- Compute the heating rate of a shut-in secondary from its energy balance and compare it with the stroke time and the sensor lag.
- Measure the process gain of the dump and explain why a full-load program step saturates any reasonable gain.
- Separate what a load-programmed reference and an explicit lead-lag feedforward each contribute to the anticipation.
- Explain how low-selected overrides cap a shared actuator, and why a controller in front of an interlock beats the interlock.
- Judge when a reactor runback, the reliefs or a larger dump is the right remedy, including for failures that nothing measures.

Prerequisites Sections 1.4.2 (prompt jump and reactivity feedback), 1.4.4 (shrink and swell), 1.4.5 (the clock to a limit), 1.6.2 (velocity form, back-calculation) and 1.6.4 (scenarios and data). Chapter 2 for rod control and the temperature program.

The app https://apmonitor.com/apps/turbine-trip/

See also Volume 3, Chapter 2 for a structure that rejects only the disturbance entering its measurement path, and Volume 3, Chapter 6 for feedforward from an upstream measurement with a slow trim. Volume 4, Chapter 2 for a protective override on a fast transient, and Volume 4, Chapter 3 for a selector over a governor and its limiters.

5.1 Introduction

Almost every disturbance a first control course presents arrives slowly enough to be measured. A feed temperature drifts, a load steps, a composition wanders, and in every case the controlled variable moves, the transmitter reports it, and the loop responds before anything is lost. The lesson a student draws is that a loop which rejects a disturbance badly needs a higher gain or a faster measurement. A generator load rejection on a large steam plant refutes it. The breaker opens, the turbine overspeed protection controller shuts the governor valves inside a second, and about 1600 kilograms per second of steam that was going through a machine is going nowhere, while the reactor making it has no idea anything has changed. The controlled variable does not move for several seconds and the transmitter that reads it has a response time of four.

> **DID YOU KNOW? Heat That Outlasts the Trip**
>
> When a reactor trips, fission stops within a second, but the decaying fission products keep heating the coolant. At Three Mile Island in 1979 that heat was just 6 percent of the fission heat, still large enough that it had to be removed to keep the core from overheating [1].

The misconception this chapter corrects is that one. Feedback performance against a disturbance is bounded by the sensitivity function of the loop, and that bound is set by the phase and the delay rather than by the gain, so a disturbance whose spectrum lies above the achievable bandwidth cannot be attenuated at any gain at all [2]. A block diagram states that result and a transfer function proves it, but neither shows what it costs, because the cost here is a discrete list of consequences a plant either avoids or does not. Measured on the shipped constants, seven times the proportional gain on the dump temperature loop with a fivefold shorter integral time moves the peak steam pressure from 83.37 to 81.70 bar, and an average-temperature RTD made eight times faster than any real instrument gives 82.92 bar with 454 kg of steam still through the safety valves. The same controller handed one additional input reaches 79.74 bar and lifts nothing.

The class of problem is general and is not confined to steam plants. A disturbance arrives fully formed, faster than any measurement of its effect can develop, so the only remedy is anticipation from a signal upstream of the effect, and the constraints that were comfortable in normal operation close in during the transient and have to be handled by something other than the loop that caused them. Tuning rules for feedforward from a measured disturbance exist for exactly this situation and are written in terms of the disturbance path rather than the process path [3], and selective or override structures are the standard answer to a constraint that binds only transiently [4]. Compressor anti-surge is the same shape, because a driver trip collapses the flow through the machine faster than the surge line can be approached gracefully and the anti-surge valve must be opened from a signal that is not the surge margin itself [5], and Volume 4, Chapter 2 treats that machine with a protective override. Generator load rejection on a fossil unit or a gas turbine is the same shape by construction, and so is a batch reactor whose exotherm runs away faster than a jacket temperature can report it. In every case the plant carries a fast anticipatory path, a slow trim, and a set of overrides that speak only when a constraint moves.

Simulators of nuclear steam supply systems are ordinarily built to answer a transient-analysis question rather than a control-architecture one [6, 7], and those built for control work usually stop at the reactor and the temperature program [8]. However, a plant model carrying one control structure cannot separate what the physics decided from what the engineer decided, and on a load rejection that separation is the whole question. This chapter puts a pressurized water reactor secondary side under a generator load rejection in front of a control engineer as an experiment rather than a calculation, and prices each rung of a four-step control ladder in bar of peak steam pressure and in tonnes of steam through the safety valves. The plant reduces published steam generator and station dynamics to teaching scale [8, 9], and steam dump control is a structure Westinghouse units have carried for decades [10]. Four named architectures share one actuator, one sample period, one set of limits and one disturbance, so that the difference between them is wiring and nothing else, and the feedback only structure generates its own reference from a measurement downstream of the event. The scenarios form a sequence in which the failures are structural rather than arithmetic and two of them are out of reach of tuning altogether.

Section 5.2 traces steam bypass from a startup device to a licensed transient and names the control structures the vendors fixed. Section 5.3 gives the balances, the constants and the integration scheme in the detail a reader needs to rebuild the plant. Section 5.4 measures the clock with no controller and the process gain of the dump. Section 5.5 sets out the temperature program, the anticipation, the dump demand with its two overrides, and the reliefs, rods and feedwater. Section 5.6 compares the four architectures on one rejection, shows that tuning does not close the gap, locates the anticipation, and follows the constraint that moves with the weather, the two dump modes and the failures nothing reports. Section 5.7 sets the loop in its plant hardware, and Sections 5.8 to 5.10 state the limitations, discuss what the results mean for a design and summarize the chapter.

5.2 Turbine Bypass and Steam Dump Control

Steam bypass around a turbine is older than nuclear power and began as a startup device rather than a transient one. A boiler brought up to pressure has to put its steam somewhere while the turbine is still on the turning gear, and the earliest arrangements were manual valves to the condenser or to atmosphere sized for whatever the boiler made during warmup. What turned bypass into a control problem was the arrival of units large enough that losing the grid was a plant event rather than an inconvenience. A drum boiler can be fired down in a minute or two. A reactor cannot be switched off, its decay heat continues for days, and the buffer between it and the turbine is only the inventory of the steam generators, so the disposal path has to be sized against the transient rather than against the startup.

The sizing question settled into a split that survives in current practice. Ordinary pressurized water units carry a steam dump of roughly 40 % of rated steam flow, enough to hold the plant through a

turbine trip and bring the average coolant temperature down to its no-load program without lifting the main steam safety valves, and units specified to shed the grid and remain at power carry up to 85 % [10]. Vendor treatments of bypass valve design frame the choice in the same terms and add that a plant expected to cycle needs the larger valve for flexibility rather than for protection [11]. The constraint deciding whether the larger valve is worth buying sits downstream of it, because a condenser sized for turbine exhaust is loaded far beyond its design duty by a large bypass, and guidelines for bypass operation are written around the condenser rather than the valve [12]. Section 5.6 measures that constraint on the shipped model and reaches the same conclusion.

Licensing turned the transient into a required calculation. Loss of external load and turbine trip are grouped together in the standard review plan as decrease-in-heat-removal events, and the acceptance criteria are limits on primary and secondary pressure, on safety valve actuations and on the departure-from-nucleate-boiling ratio, with the applicant required to show the consequences both with and without credit for the steam dump and the reactor control system [13]. That requirement is why a load rejection analysis exists for any unit and why the steam dump is documented as a system with arming permissives and interlocks rather than as a valve. Analyses of real full load rejections followed the requirement rather than preceding it, and the early ones were code-to-plant comparisons on units that had actually shed the grid [7]. Modern work on the same transient is largely about coordination, because a plant asked to ride out a rejection has to move the dump, the rods and the feedwater together rather than sequentially [6], and the pressure to do so comes from grid duty rather than from safety [14].

> **HISTORY** **Riding Out a Turbine Trip Below Half Power (1979)**
>
> In December 1979 a national laboratory working for the United States regulator reviewed a request from the owner of one pressurized water reactor. Its protection system tripped the reactor whenever the turbine tripped above 10 percent power. Because both the reactor and the turbine would accept a 50 percent load rejection, the owner asked to trip the reactor on a turbine trip only at or above 50 percent power. A new permissive, P-9, formed from the power range neutron flux channels, would block the reactor trip below that power. The turbine trip would then ride out as a load rejection, with the reactor power dissipated by a steam dump able to bypass up to 50 percent of rated steam. The review noted that the safety analysis took no credit for the steam dump and recommended approval [15].

The control structure itself was fixed by the vendors and is unusually well documented. A Westinghouse steam dump control system carries two operator selected modes. The load rejection controller works on average coolant temperature against a program derived from turbine first-stage impulse pressure, so it acts continuously and in proportion to the load actually lost. The turbine trip controller works against the no-load program instead, is armed by the P-4 turbine trip permissive rather than by a continuous measurement, and is blocked entirely until that permissive appears. Condenser back pressure and steam line pressure interlock the dump valves shut, and the rod control system carries its own runback logic driven from the same load rejection signal [10]. Every one of those is a control-structure decision rather than a tuning decision, and every one is reproduced in the app of this chapter.

The theory those structures rest on was developed in the process industries under different names. A constraint controller that publishes a limit and is low-selected against a working demand is a selective or override control, and the treatment establishing that such a controller caps a demand rather than adding to it is a process control one [4]. Handover between a selected and an unselected controller is the same problem as anti-windup, and the conditioning technique that solves both by back-calculating the integral against the limit actually applied is the standard method [16]. Systematic procedures for designing a set of selectors against moving active constraints are more recent and treat the selector as a control element in its own right [17]. A steam dump controller is a worked example of all three, which is why it belongs in front of a process control student and not only a nuclear one.

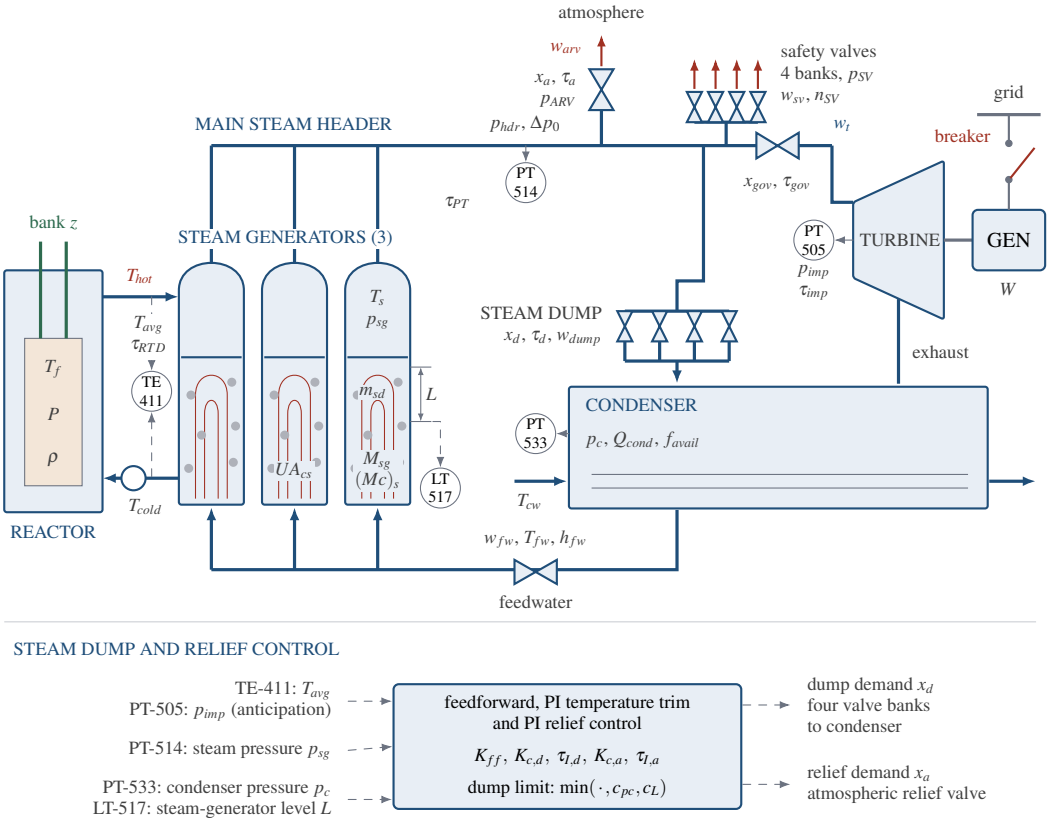

Figure 5.1 Secondary side of unit U-1 during a generator load rejection. The core delivers P across three steam generator tube bundles at conductance UA_{cs}, boiling the inventory M_{sg} at saturation temperature T_s and pressure p_{sg} with steam mass m_{sd} circulating below the water surface. Steam leaves through the header at p_{hdr} into the governor valve x_{gov}, through the four dump banks x_d to the condenser at back pressure p_c, and to atmosphere through the reliefs x_a and the four safety valve banks set at p_{SV}. The dump controller reads turbine impulse pressure p_{imp} as its anticipatory signal, average coolant temperature T_{avg} through RTD TE-411, and condenser pressure and generator level as the two caps c_{pc} and c_L that low-select against its own demand. Dashed lines are signals and the breaker is drawn open.

5.3 Plant Model

Unit U-1 is a Westinghouse-type three-loop pressurized water reactor rated at 2900 MW thermal and 940 MW electrical gross. Three steam generators pass the whole core duty across their U-tube bundles, boiling 1595.0 kg/s of saturated steam at 61.48 bar into a main steam header that runs to one turbine. Four banks of steam dump valves discharge from that header to the condenser, atmospheric reliefs and four banks of main steam safety valves discharge from it to atmosphere, and feedwater returns from the condenser hotwell at 226 °C. That plant is drawn with its instrument tags in Figure 5.1, which also carries every symbol the equations use. The generator breaker is the disturbance and the steam dump valve position is the manipulated variable.

Table 5.1 separates what the loops control from what they move, what upsets them and what the integrator carries. Nineteen states are integrated, and the working units are those of nuclear steam plant practice, power in MW, heat capacity in MW s/°C, temperature in °C, pressure in bar except the condenser end which is in mbar, mass in kg, flow in kg/s, enthalpy in kJ/kg, reactivity in

Table 5.1 Classification of Variables. The Nominal Column Is the Full-Power Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
T_{avg} (TE-411)	CV, measured	310.00 °C	291.7–310.0 °C SP
p_{sg} (PT-514)	CV, measured	61.48 bar	68–84 bar SP
L (LT-517)	CV, measured	50.0 %	35–65 % SP
x_d	MV	0 %	0–100 %
x_a	MV	0 %	0–100 %
v_z rod velocity	MV	0 %/s	±0.5 %/s
w_{fw}	MV	100 %	0–130 %
breaker	DV, measured	closed	open, closed
turbine demand	DV, measured	100 %	0–105 %
T_{cw}	DV, measured	21 °C	4–34 °C
dump banks avail.	DV, unmeasured	100 %	0–100 %
x_a stuck open	DV, unmeasured	0 %	0–100 %
$C_1 \dots C_6$	state	—	≥ 0
D_1, D_2	state	0.028, 0.040	≥ 0
T_f	state	610.0 °C	—
T_{avg}	state	310.00 °C	—
T_s	state	277.01 °C	180–330 °C
M_{sg}	state	135 000 kg	$\geq$ 1000 kg
m_{sd}	state	3190 kg	≥ 0
z	state	68.0 %	0–100 %
x_{gov}	state	0.90	0–1.05
x_d, x_a	state	0 %	0–100 %
p_c	state	53.1 mbar	10–1200 mbar
w_{fw}	state	1595.0 kg/s	≥ 0

pcm ($10^{-5}\,\Delta k/k$) and time in sec. Every controller works in percent of valve travel and percent of transmitter span.

5.3.1 Saturated Properties

The secondary visits 40 to 90 bar and nothing else, so the saturation line is carried as low-order fits rather than as a property library, as shown in Equation (5.1).

$$\begin{aligned} T_{sat}(p) &= 185.0 + 1.88912\,p - 0.006384\,p^2 \\ h_{fg}(p) &= 1930 - 6.0\,p, \qquad h_g(p) = 2831.7 - 0.8\,p \\ h_f(p) &= 1213.7 + 5.2\,(p - 60) \\ \rho_s(p) &= 0.4956 + 0.45663\,p + 8.306\times 10^{-4}\,p^2 \end{aligned} \tag{5.1}$$

where p is in bar, temperature in °C, enthalpy in kJ/kg and steam density in kg/m^3. Saturation temperature is good to about 0.5 °C and latent heat to about 1 % over that band, and $p_{sat}(T)$ is the positive root of the first line rather than a second fit, so the two are exact inverses. The fits carry no pressure clamp, only the saturation temperature state is held between 180 and 330 °C as a numerical guard, and no scenario leaves the fitted band. Feedwater enthalpy is $h_{fw} = 4.19\,T_{fw} + 3.4\times 10^{-4} T_{fw}^2$, which returns 964.3 kJ/kg at 226 °C, and the condenser end uses the Antoine expression $p_c = 10^{8.07131 - 1730.63/(233.426+T)}/750.062$ in bar for T clamped to 5 to 99 °C.

5.3.2 Core Neutronics and Thermal Model

The prompt neutron mode of a light water reactor is about 3 ms and the transient of interest runs for 300 s, so the prompt jump approximation of Section 1.4.2 is applied and only the six delayed

precursor groups are integrated [18], as shown in Equations (5.2) and (5.3).

$$\frac{\mathrm{d}C_i}{\mathrm{d}t} = \beta_i n - \lambda_i C_i, \qquad i = 1 \ldots 6 \tag{5.2}$$

$$n = \frac{\sum_i \lambda_i C_i}{\max(\beta - \rho,\ 0.05\,\beta)} \tag{5.3}$$

where n is fission power as a fraction of rated, C_i is the precursor concentration of group i in the same units, β_i and λ_i are the thermal-fission delayed-neutron fractions and decay constants of U-235, and ρ is net reactivity in $\Delta k/k$. The six fractions are 2.15, 14.24, 12.74, 25.68, 7.48 and 2.73×10^{-4} summing to $\beta = 0.006502$, against decay constants of 0.0124, 0.0305, 0.111, 0.301, 1.14 and 3.01/s [18]. The denominator floor in (5.3) keeps the expression finite on a deep negative excursion. Reactivity is the sum of bank worth and two temperature feedbacks, as shown in Equations (5.4) and (5.5).

$$\rho = \rho_z(z) - \rho_z(z_{ref}) + \alpha_F \Delta T_f + \alpha_M \Delta T_{avg} - \rho_{scram} \tag{5.4}$$

$$\rho_z(z) = W_{tot}\left[f - \frac{\sin 2\pi f}{2\pi}\right], \qquad f = \frac{z}{100} \tag{5.5}$$

where $\Delta T_f = T_f - T_{f,0}$ and $\Delta T_{avg} = T_{avg} - T_{avg,0}$ are departures from the design point, z is control bank position in percent withdrawn, W_{tot} is total bank worth, α_F is the Doppler coefficient, α_M is the moderator temperature coefficient and ρ_{scram} ramps to the scram rod worth over t_{scram} after a trip. Equation (5.5) is the integral S curve, zero fully inserted and W_{tot} fully withdrawn, and its derivative $\frac{\mathrm{d}\rho_z}{\mathrm{d}z} = (W_{tot}/100)(1 - \cos 2\pi f)$ is 42.8 pcm/% at the full-power position $z_{ref} = 68.0\,\%$, where bank worth is 2472 pcm of the 3000 pcm total. That derivative sets the authority of the runback, because 0.5 %/s of bank travel is 21.4 pcm/s of reactivity.

Decay heat is two lumped groups and total core power is their sum with the prompt fission term, as shown in Equations (5.6) and (5.7).

$$\frac{\mathrm{d}D_j}{\mathrm{d}t} = \frac{f_{D,j}\, n - D_j}{\tau_{D,j}}, \qquad j = 1, 2 \tag{5.6}$$

$$P = P_0\left[(1 - f_{D,1} - f_{D,2})\, n + D_1 + D_2\right] \tag{5.7}$$

where $f_{D,j}$ and $\tau_{D,j}$ are the group fraction and time constant, and the weighting in (5.7) makes the equilibrium at $n = 1$ exactly P_0, so the two groups redistribute power rather than adding it. Two thermal nodes carry the heat to the secondary, as shown in Equations (5.8) and (5.9).

$$m_f c_f \frac{\mathrm{d}T_f}{\mathrm{d}t} = f_P P - UA_{fc}\,(T_f - T_{avg}) \tag{5.8}$$

$$\begin{aligned} m_c c_c \frac{\mathrm{d}T_{avg}}{\mathrm{d}t} &= (1 - f_P)P + UA_{fc}(T_f - T_{avg}) \\ &\quad - UA_{cs}(T_{avg} - T_s) \end{aligned} \tag{5.9}$$

where T_f is average fuel temperature, T_{avg} is average coolant temperature, f_P is the fission power fraction that lands in the fuel, and UA_{fc} and UA_{cs} are the fuel-to-coolant and primary-to-secondary conductances. Both sides carry MW. Loop temperatures follow from the core rise, $T_{hot,cold} = T_{avg} \pm \frac{1}{2}\Delta T_{core} P/P_0$, which give 326.50 and 293.50 °C at full power. The primary is incompressible, so there is no pressurizer state.

5.3.3 The Secondary Energy Balance

Equation (5.10) is the whole transient in one line, and every measured number in this chapter follows from it. It is the plant's own form of the clock of Equation (1.8).

$$(Mc)_s \frac{\mathrm{d}T_s}{\mathrm{d}t} = UA_{cs}(T_{avg} - T_s) - \frac{w_{tot} h_{fg}}{1000} - \frac{w_{fw}[h_f - h_{fw}]}{1000} \tag{5.10}$$

where $(Mc)_s$ is the lumped secondary heat capacity of the saturated water, the tube bundles and the shells of all three generators [9, 19] in MW s/°C, w_{tot} is the sum of every steam flow leaving the generators in kg/s, and the factors of 1000 convert kJ/s to MW. The first term is 2900 MW at full power, the second is 1595.0 kg/s at 1561 kJ/kg of latent heat and the third is the 410 MW of sensible heat the feedwater absorbs on its way to saturation, and the three balance to zero. Secondary pressure follows the state through the property fit, $p_{sg} = p_{sat}(T_s)$, and the slope of that fit at the design point is 0.906 bar/°C.

Equation (5.10) is written in saturation temperature because the driving force for the tube bundle is a temperature difference and because writing it this way makes the arithmetic legible. With w_{tot} driven to zero and everything else at the design point the right-hand side is $(2900-410)/1200 = 2.075\,^{\circ}\mathrm{C/s}$, and one degree is 0.906 bar. The rest of the chapter is a consequence of that number.

Pause. Try Exercise 5.1 before reading on.

5.3.4 The Steam Header and Its Valves

Every valve on the header passes flow proportional to header pressure, and the header pressure falls below generator pressure as the square of the total flow. Substituting one relation into the other gives a quadratic in the flow ratio $\xi = w/w_0$, solved in closed form as shown in Equations (5.11) and (5.12).

$$w = K w_0 \frac{p_{hdr}}{p_{hdr,0}}, \qquad p_{hdr} = p_{sg} - \Delta p_0\, \xi^2 \tag{5.11}$$

$$\xi = \frac{-1+\sqrt{1+4ac}}{2a}, \qquad a = \frac{\Delta p_0 K}{p_{hdr,0}}, \quad c = \frac{K p_{sg}}{p_{hdr,0}} \tag{5.12}$$

where $K = K_t + K_d$ is the total flow coefficient of the valves on the header as a multiple of the design value, w_0 is rated steam flow, Δp_0 is the main steam line drop at rated flow and $p_{hdr,0} = p_{sg,0} - \Delta p_0 = 59.98$ bar. Solving (5.11) in closed form matters for more than tidiness. A fourth-order Runge-Kutta stage set asks for the state derivatives four separate times inside one step, and an iterative header solve buried in them would answer each of the four slightly differently.

The two coefficients and the flow split are given by Equation (5.13), and the reliefs and safety valves discharge from the steam lines upstream of the header restriction, so they see generator pressure directly, as shown in Equation (5.14).

$$K_t = \frac{x_{gov}}{x_{gov,0}}, \qquad K_d = c_{av}\frac{r_d\, x_d}{100}$$

$$w_t = w_0 \xi \frac{K_t}{K_t+K_d}, \qquad w_{dump} = w_0 \xi \frac{K_d}{K_t+K_d} \tag{5.13}$$

$$w_{arv} = w_0\, r_a \frac{x_a}{100}\frac{p_{sg}}{p_{sg,0}}, \qquad w_{sv} = w_0\, r_{sv}\, n_{SV}\frac{p_{sg}}{p_{sg,0}} \tag{5.14}$$

where r_d, r_a and r_{sv} are the dump, relief and per-bank safety valve capacities as fractions of rated steam flow, $x_{gov,0} = 0.90$ is the governor position at full load, and c_{av} is the product of the fraction of dump banks in service and the condenser interlock of Equation (5.17). At the design capacities the full dump bank passes 638 kg/s, the reliefs 159 kg/s and each safety valve bank 478 kg/s. The safety valves are four banks staggered by Δp_{SV}, each ramping linearly to full lift one accumulation above its own setpoint, as shown in Equation (5.15).

$$n_{SV} = \sum_{k=0}^{3} \mathrm{clip}\left(\frac{p_{sg} - (p_{SV} + k\,\Delta p_{SV})}{p_{acc}},\ 0,\ 1\right) \tag{5.15}$$

where $p_{SV} = 82.7$ bar is the lowest bank and p_{acc} is the accumulation to full lift. Equation (5.15) carries no latch and no blowdown, which keeps the right-hand side continuous for the integrator at the cost of understating the discharge immediately after a real valve pops and overstating it during reseating.

5.3.5 The Condenser and the Interlock

The condenser is the steam dump's only customer, and it is modeled as a back pressure correlation rather than as a heat exchanger, as shown in Equation (5.16).

$$p_c^* = p_{sat}\left(T_{cw} + \frac{Q}{\dot{m}_{cw}c_p} + TTD_0\left(\frac{Q}{Q_0}\right)^{0.4} + \frac{TTD_d\, Q_{dump}}{Q_0}\right) \tag{5.16}$$

where $Q = Q_0 w_t / w_0 + Q_{dump}$ is the total duty, $Q_{dump} = w_{dump}[h_g(p_{sg,0}) - 140]/1000$ is the dump contribution with condensate leaving at 140 kJ/kg, T_{cw} is the circulating water inlet temperature, $\dot{m}_{cw}c_p$ is its heat capacity rate, TTD_0 is the terminal temperature difference at design duty and TTD_d is an extra terminal difference charged against dump steam. The state follows that target through a first-order lag, $\tau_{cond}\frac{\mathrm{d}p_c}{\mathrm{d}t} = p_c^* - p_c$.

The exponent and the extra terminal difference exist because turbine exhaust and dump steam are not the same commodity. Exhaust has already given up its work and arrives spread across the tube bundle, while dump steam arrives at full enthalpy through a handful of nozzles [12]. Measured on the shipped constants, 40 % of rated flow through the dump is 1686 MW into a condenser rated for 1940 MW, which is 87 % of design duty from a plant making no power at all, and 85 % of rated flow is 3583 MW, or 1.85 times design. The plant protects the condenser with an interlock rather than a controller, and the interlock tightens as it is approached, as shown in Equation (5.17).

$$f_{avail} = \mathrm{clip}\left(\frac{p_{block} - p_c}{p_{block} - p_{derate}},\ 0,\ 1\right) \tag{5.17}$$

where $p_{derate} = 115$ mbar is the back pressure at which the dump begins to be cut back and $p_{block} = 150$ mbar is where it is shut entirely, and a separate latch blocks the dump permanently at $p_{trip} = 165$ mbar, five inches of mercury absolute, where a real unit would also trip the turbine. The feedback runs the wrong way for the controller, because more dump raises Q_{dump} in (5.16), which raises p_c, which lowers f_{avail} and takes the dump away.

5.3.6 Inventory, Circulating Steam and Indicated Level

The generators hold a water inventory and, below the water surface, a mass of circulating steam occupying real volume. The circulation carries gross evaporation less the condensation on the incoming subcooled feedwater for one riser residence time τ_{res}, as Equation (5.18) states [9].

$$m_{sd}^* = \tau_{res}\left[\underbrace{\frac{1000\,UA_{cs}(T_{avg} - T_s)}{h_{fg}}}_{w_{boil}} - \underbrace{\frac{w_{fw}[h_f - h_{fw}]}{h_{fg}}}_{w_{cond}}\right] \tag{5.18}$$

where both bracketed terms are in kg/s. The two mass balances and the level indication follow, as shown in Equations (5.19) and (5.20).

$$\tau_{void}\frac{\mathrm{d}m_{sd}}{\mathrm{d}t} = m_{sd}^* - m_{sd} \tag{5.19}$$

$$\frac{\mathrm{d}M_{sg}}{\mathrm{d}t} = w_{fw} - w_{tot} - \dot{m}_{sd}$$

$$L = \mathrm{clip}\left(50 + \frac{M_{sg} - M_{sg,0}}{k_M} + k_{void}\Delta V_{sd},\ 0,\ 100\right) \tag{5.20}$$

where $\dot{m}_{sd}$ is the derivative in the first line, τ_{void} is the void relaxation time, k_M is inventory per percent of narrow-range span, k_{void} converts sub-surface steam volume to indicated level and $\Delta V_{sd} =$

$m_{sd}/\rho_s(p_{sg}) - V_{sd,0}$ is the departure of that volume from its design value. The transmitter clips at its taps and the protection system reads the same clipped signal, because both are wired to the same instrument. Equation (5.20) is the reason indicated level is not the water inventory. Raising the pressure compresses the voids and the level shrinks on a plant that has not lost a kilogram, and lowering it does the reverse. This is the shrink and swell of Section 1.4.4, which Chapter 4 studies on a drum. Section 5.6 reports which direction this transient actually takes.

5.3.7 Actuators and Instruments

Four valve positions and the rod drive are states, as shown in Equation (5.21).

$$\tau_{gov}\frac{\mathrm{d}x_{gov}}{\mathrm{d}t} = u_{gov} - x_{gov}, \quad \tau_d\frac{\mathrm{d}x_d}{\mathrm{d}t} = u_d - x_d, \quad \frac{\mathrm{d}z}{\mathrm{d}t} = v_z \tag{5.21}$$

with the atmospheric reliefs and the feedwater demand following the same first-order form at τ_a and τ_{fw}. A valve derivative is zeroed whenever its state rests on a travel limit and the command still points outward, which stops a fixed step from carrying the position through a stop and back, and each position is clipped once more after the step. The rod drive is a velocity actuator and is blocked at the insertion limit and at full withdrawal the same way. A dump stroke lag of $\tau_d = 0.9\,\mathrm{s}$ is three seconds to full open, which is the number to compare against the clock.

Five instruments carry first-order lags, each advanced by the zero-order-hold form $\hat{y}_k = \alpha\hat{y}_{k-1} + (1-\alpha)y_k$ with $\alpha = \exp(-\Delta t_c/\tau)$, which cannot go unstable at any sample period. The average coolant temperature RTD at $\tau_{RTD} = 4.0\,\mathrm{s}$ is the most consequential parameter in the application, because it is longer than the margin the transient has, and the turbine impulse pressure transmitter at $\tau_{imp} = 0.30\,\mathrm{s}$ is the one that makes anticipation possible. The other three are the steam pressure transmitter at $\tau_{PT} = 0.50\,\mathrm{s}$, the generator level transmitter at $\tau_{LT} = 2.0\,\mathrm{s}$ and the excore neutron detector filter at $\tau_{NI} = 0.10\,\mathrm{s}$.

5.3.8 Protection

Four protective actions are modeled and none is a control action, in the layering of Section 1.4.5. The reactor trip acts on high neutron flux above 118 % of rated, on overtemperature ΔT when the loop rise exceeds the setpoint of Equation (5.22), or on low-low generator level below 17 %.

$$\Delta T_{set} = \Delta T_{core}\left[K_1 - K_2(T_{avg} - T_{avg,0})\right] \tag{5.22}$$

where $K_1 = 1.13$ and $K_2 = 0.030/^\circ\mathrm{C}$, so the allowed loop rise falls as the average temperature climbs. That is the protection system's proxy for margin to departure from nucleate boiling, and it is the trip a secondary side failing to take heat away eventually reaches. High-high generator level above 78 % isolates feedwater (the turbine trip that accompanies it on a real unit is not modeled), a reactor trip runs automatic main feedwater back to 20 % of rated, and auxiliary feedwater appears as a floor of 6 % of rated flow on low-low level.

5.3.9 Parameters and Initialization

Table 5.2 carries the constants.

No scenario starts from a nominal condition. Given a turbine load demand, the plant settles where the steam the turbine passes at the resulting pressure equals the steam the reactor's own temperature feedback makes at the programd average temperature, which is a fixed point rather than a formula. The initialization iterates forty passes of damped substitution on the load fraction, derives the fuel temperature from Equation (5.8) at zero derivative, solves Equation (5.5) by bisection for the bank position that makes $\rho = 0$, seeds the precursors at $C_i = \beta_i n/\lambda_i$, and sets the inventory so that Equation (5.20) reads exactly 50.0 %. The result at full power is $T_{avg} = 310.00\,^\circ\mathrm{C}$, $T_f = 610.0\,^\circ\mathrm{C}$, $T_s = 277.01\,^\circ\mathrm{C}$, $p_{sg} = 61.48\,\mathrm{bar}$, $p_{hdr} = 59.98\,\mathrm{bar}$, $w_0 = 1595.0\,\mathrm{kg/s}$, $z = 68.0\,\%$, $M_{sg} = 135\,000\,\mathrm{kg}$ and $p_c = 53.1\,\mathrm{mbar}$, and the largest state derivative there is below 10^{-8} in its own units.

Table 5.2 Model Parameters. Five Constants Are Hidden From the Operator and the Remainder Are Visible in the Model Panel With a Slider.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
P_0	rated thermal power	2900	MW	Δp_0	steam line drop at rated flow	1.5	bar
$W_{e,0}$	gross generator output	940	MW	$p_{imp,0}$	first-stage pressure, full load	45.0	bar
f_P	fission power into the fuel	0.974	—	r_d	steam dump capacity	0.40	—
$m_f c_f$	fuel heat capacity	47.5	MW s/°C	τ_d	dump valve stroke lag	0.9	s
UA_{fc}	fuel-to-coolant conductance	9.415	MW/°C	r_a	atmospheric relief capacity	0.10	—
$m_c c_c$	primary heat capacity	850	MW s/°C	τ_a	relief stroke lag	2.5	s
ΔT_{core}	core rise at full power	33.0	°C	p_{SV}	first safety valve bank	82.7	bar
α_F	Doppler coefficient	−2.8	pcm/°C	Δp_{SV}	safety valve bank spacing	0.85	bar
α_M	moderator coefficient	−30.0	pcm/°C	p_{acc}	accumulation to full lift	2.3	bar
W_{tot}	control bank worth	3000	pcm	r_{sv}	capacity of each SV bank	0.30	—
v_{max}	maximum rod speed	0.50	%/s	$x_{gov,0}$	governor position, full load	0.90	—
z_{lim}	rod insertion limit	18	%	τ_{gov}	governor closing lag	0.30	s
ρ_{scram}	scram rod worth	0.045	$\Delta k/k$	Q_0	condenser duty at full load	1940	MW
t_{scram}	scram rod drop time	2.2	s	$\dot{m}_{cw} c_p$	circulating water capacity rate	194	MW/°C
UA_{cs}	primary-to-secondary, 3 SG	87.9	MW/°C	TTD_0	terminal difference at design	3.0	°C
$(Mc)_s$	secondary heat capacity, 3 SG	1200	MW s/°C	TTD_d	extra difference on dump steam	6.0	°C
$M_{sg,0}$	inventory at full power	135 000	kg	τ_{cond}	condenser response time	8.0	s
k_M	inventory per percent span	666	kg/%	p_{derate}	dump cutback back pressure	115	mbar
k_{void}	void volume to level	0.62	%/m^3	p_{block}	dump block back pressure	150	mbar
τ_{res}	riser residence time	2.0	s	p_{trip}	condenser high-pressure trip	165	mbar
τ_{void}	void relaxation time	8.0	s	τ_{fw}	feedwater response time	4.0	s
T_{fw}	feedwater temperature	226	°C	K_{fwl}	feedwater level trim gain	2.2	%/%
T_{nl}	no-load programmed T_{avg}	291.7	°C	τ_{RTD}	T_{avg} RTD response	4.00	s
T_{fl}	full-load programd T_{avg}	310.0	°C	τ_{imp}	impulse pressure transmitter	0.30	s
K_1	overtemperature ΔT constant	1.13	—	τ_{PT}	steam pressure transmitter	0.50	s
K_2	overtemperature ΔT coefficient	0.030	1/°C	τ_{LT}	level transmitter	2.00	s
n_{trip}	high neutron flux trip	118	%	τ_{NI}	excore detector filter	0.10	s

Hidden constants. Bank position at full power $z_{ref} = 68\,\%$. Fast decay-heat group $f_{D,1} = 0.028$ with $\tau_{D,1} = 25\,$s. Slow decay-heat group $f_{D,2} = 0.040$ with $\tau_{D,2} = 1500\,$s. *Visible, not listed above.* Level trips $L_{hi,trip} = 78\,\%$ and $L_{lo,trip} = 17\,\%$, and the 20 % feedwater limit after a reactor trip.

The engine of Section 1.6.1 advances the states by classical fourth-order Runge-Kutta on a fixed 10 ms step, the controller runs on a 50 ms period and the historian keeps twenty samples for every simulated second. The fastest modeled dynamics are the governor valve at 0.30 s and the sixth delayed-neutron group at 0.33 s, which leaves the step well inside its stability limit, and the model draws no random numbers, so a scenario replays identically.

5.3.10 Assumptions

Every simplification below carries the condition under which it fails.

Point kinetics with the prompt jump taken. Equation (5.3) is correct for a transient measured in seconds to minutes and wrong wherever the 3 ms prompt mode matters, such as a rod ejection. There is no spatial kinetics, so the axial offset transient a rapid rod insertion produces is absent, and there is no xenon.

Two thermal nodes and constant reactivity coefficients. A real core carries a radial and axial power distribution, a hot channel and a departure-from-nucleate-boiling margin that depends on all of it, so Equation (5.22) is the correct functional form fed by a lumped model. The moderator coefficient runs from near zero at beginning of cycle to about −50 pcm/°C at end of cycle and the whole character of the transient changes with it, but the model has no burnup.

No pressurizer. The primary is incompressible, so the pressurizer level and pressure transient, a large part of what a control room actually watches during a load rejection, is absent along with the spray and heater actions that respond to it.

One lumped secondary for three steam generators. Equation (5.10) carries a single saturation temperature, while a real three-loop plant has three generators seeing slightly different feedwater

temperatures and steam line resistances, and loop-to-loop asymmetry is a genuine feature of a real event.

The header is quasi-steady. Equation (5.12) has no steam line inventory, no acoustic transient and no choked flow, so it cannot reproduce the pressure wave an instantaneous valve closure sends up a main steam line.

Safety valves are continuous ramps. Equation (5.15) has no latch and no blowdown, so the masses quoted here are indicative rather than exact, where a real spring valve snaps open and stays open until the pressure falls well below its setpoint.

The condenser is a correlation. Equation (5.16) captures why the dump loads the condenser and not how a condenser condenses. The extra terminal difference on dump steam is a fitted allowance, and air in-leakage, tube fouling and spray attemperation are absent.

Feedwater is a lag on a demand. There is no feed pump, no three-element flow loop and no tube-side dynamics, and auxiliary feedwater appears only as a flow floor, so a main feedwater pump trip cannot be studied here.

The numerics are more repeatable than a control system. Freezing the step and the controller period gives the same trajectory on every machine, where a distributed control system delivers scan jitter, a coarser signal resolution and an execution order that shifts.

5.4 The Clock

The whole control problem follows from how fast the secondary heats up with nowhere to put its steam, so that number is measured first with no controller in the loop. The plant is initialized at the full power design point, the governor command is set to zero, the dump and the reliefs are held shut, the rods are frozen and the feedwater is held at its initial flow, and Equations (5.2) through (5.21) are integrated on a 2 ms explicit step until the generator pressure reaches the first safety valve bank.

The instantaneous heating rate at the design point with the governor already shut is 2.075 °C/s, exactly the arithmetic of Equation (5.10), and the peak rate during the real shut-in is 1.899 °C/s at $t = 1.18$ s, lower because the governor valve takes 0.30 s to close and the bundle duty is already falling as T_s rises. In pressure the peak rate is 1.755 bar/s at $t = 1.29$ s, the generator pressure passes the no-load value of 76.00 bar at $t = 9.62$ s, and it reaches the first safety valve bank at 82.7 bar at $t = 14.80$ s.

Three numbers make that clock harder than it looks. The dump valves stroke in 2.7 s to full open, so the valves alone consume a fifth of the budget after they are told. The average coolant temperature RTD has a response of 4.0 s, so a controller reading it consumes another quarter of the budget before its measurement is correct. And the no-load steam pressure is itself 76.00 bar, because the no-load programmed average temperature of 291.7 °C has that saturation pressure, leaving only 6.70 bar between where the plant has to end up and where the safety valves open. The margin fought over is not the 21.2 bar between full power and the setpoint. It is the overshoot above a hot standby condition the plant has to reach anyway.

The reactor is the one comfortable part of the plant. Nothing tells it the turbine has gone and its only route to the news is T_{avg} in Equation (5.4). With a moderator coefficient of −30 pcm/°C on a coolant that is heating up and Doppler feedback of −2.8 pcm/°C on a fuel node with a five second time constant, core power falls from 100 to 83 % of rated by $t = 9.6$ s and to 73 % by $t = 14.8$ s without a rod moving, which is why the unattended case does not trip on flux. Reactivity feedback buys real time here and it is not enough, and the gap it leaves is why the steam dump exists.

> **TRY IT IN THE APP** **1 · Full Power, and Then the Breaker Opens**
>
> Load scenario 1 and press Run. Nobody acts: the dump stays shut and the feedwater stays on manual. When the breaker opens, watch the steam pressure climb from 61.48 bar toward the safety valve line at 82.7 bar, which it reaches in under fifteen seconds. Keep running and watch the steam generator level rise as the feedwater fills them.

5.4.1 The Process Gain of the Dump

The dump has to be tuned against something, and the steady-state gain from valve position to average coolant temperature is measured with a differenced experiment, because the inventory dynamics drift over the minutes the measurement needs and a single run would attribute that drift to the valve. The turbine is run back from 100 to 85 % of load with everything in manual, a stepped run at a fixed dump position is compared against a companion run held with the dump shut, and the difference is read over $60 \le t \le 180$ s, after which the companion run reaches the high-high level trip and is no longer a valid reference.

Read at $t = 120$ s the settled average temperatures are 312.00 °C with the dump shut, 311.18 at 15 % open, 310.39 at 30 and 308.92 °C at 60 %, so the average gain over the first 60 % of travel is −0.0513 °C/%. Solving the same steady state directly with the bank frozen gives an incremental gain falling from −0.0595 °C/% near the seat to −0.0469 °C/% wide open, because header pressure falls as the flow rises. About 0.052 °C/% is therefore the right number to tune against, and its reciprocal of 19 %/°C is unity loop gain, so the shipped $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$ is conservative by any ordinary standard.

The consequence of that gain is a mismatch of scale rather than a tuning number. A hundred percent of dump valve is worth 5.2 °C of average coolant temperature in the steady state, while the temperature program of Equation (5.23) spans 18.3 °C. The actuator is a factor of 3.5 too small for its own reference program, so on a full rejection the dump is saturated by the setpoint change alone and the only question left is how quickly the controller knows to saturate it.

5.5 Controller Design

Four architectures are selectable and share everything except which signals are wired in. Rod control and feedwater control are on automatic in all four, so the comparison isolates the dump, and only the reactor runback belongs to the coordinated structure. A manual mode drives the dump valves, the reliefs, the rod velocity and the feedwater demand directly, and the transfer between manual and automatic is bumpless in both directions, as Section 1.6.2 describes.

5.5.1 The Temperature Program and What Drives It

A pressurized water reactor programs its average coolant temperature against turbine load, as shown in Equation (5.23).

$$T_{ref} = T_{nl} + (T_{fl} - T_{nl})\,\mathrm{clip}\left(\frac{\ell}{100}, 0, 1.05\right) \tag{5.23}$$

where ℓ is a load signal in percent, $T_{nl} = 291.7\,^{\circ}\mathrm{C}$ and $T_{fl} = 310.0\,^{\circ}\mathrm{C}$. The question that decides this whole application is what supplies ℓ.

The two architectures carrying a turbine instrument take ℓ from first-stage impulse pressure, $\ell = 100\,\hat{p}_{imp}/p_{imp,0}$, where $\hat{p}_{imp}$ is the transmitter output at $\tau_{imp} = 0.30$ s. Impulse pressure measures how much steam the machine is passing, so it moves at the speed of the event because it is the event, and the moment the governor valves shut Equation (5.23) falls 18.3 °C. At $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$ that displacement alone asks for 110 % of a valve that stops at a hundred, and the measured peak of the proportional term over the transient is 132 %. The reference is therefore itself a feedforward, and it is the larger half of the anticipation in the plant.

The feedback only architecture has no turbine instrument anywhere, which is the point of it, so it generates ℓ from the only other load-like measurement available. That is the reactor's own power from the excore neutron detectors, $\ell = 100\,\hat{n}$ filtered at $\tau_{NI} = 0.10$ s, and the reactor sits downstream of the disturbance. Equation (5.4) does not move until the coolant has already heated enough to feed reactivity back, so the reference stands still while the plant heats, the loop has nothing to act on but a four second RTD, and the valves arrive late. Building the comparison this way rather than freezing the reference is what makes it a real four-way comparison, because a reference programd from turbine load is indistinguishable from feedforward and a feedback only mode that kept it would not be feedback only.

In turbine trip mode the reference is the no-load program T_{nl} outright, and the mode is armed by the P-4 turbine trip permissive, which requires both that the machine has tripped and that $x_{gov} < 0.05\,x_{gov,0}$. Until P-4 appears the dump is blocked and its controller tracks zero, so it comes off the seat from the right place the instant the permissive arrives [10].

5.5.2 Anticipation

The explicit anticipatory term is a lead-lag on the load signal rather than a proportional term, as shown in Equation (5.24).

$$x_{ff} = \mathrm{clip}\Big(K_{ff}\left[\ell_{lag} - \ell\right]^{+}, 0, 100\Big), \qquad \tau_{ff}\frac{\mathrm{d}\ell_{lag}}{\mathrm{d}t} = \ell - \ell_{lag} \tag{5.24}$$

where ℓ_{lag} is a first-order lag of the same load signal at $\tau_{ff} = 20\,\mathrm{s}$, $[\cdot]^{+}$ is the positive part and $K_{ff} = 1.20\,\%/\%$, which drives the banks fully open for a load loss of 83 % or more. The washout is not a refinement. After the transient the only heat left in the secondary is what the reactor is still making, so a feedforward proportional to load lost rather than to the rate of losing it would hold the valves open indefinitely and freeze the plant. In turbine trip mode the same term degenerates to a fixed bias armed by P-4 and decaying on the same constant, $x_{ff} = b_d \exp[-(t - t_{P4})/\tau_{ff}]$, because that mode knows only that the machine tripped.

5.5.3 The Dump Demand and Its Two Overrides

Every loop uses the velocity-form PID of Section 1.6.2, written in this chapter's symbols in Equation (5.25).

$$\begin{aligned} e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\ D_k &= -aK_c\tau_D\frac{\mathrm{dPV}^{f}}{\mathrm{d}t} \\ s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias} + s_k + I_{k-1}, \qquad u_k = \mathrm{clip}\,(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}e_k + \frac{u_k - c_k}{T_t}\right]\Delta t_c \end{aligned} \tag{5.25}$$

where $a = -1$ selects direct action, in which a rising measurement raises the output, and $a = +1$ selects reverse action. The running sum s_k is S_k of Equation (1.14), the unsaturated demand c_k is v_k of Equation (1.15) and u_{bias} is u_b. Only I_k sees the clip, through back-calculation with tracking constant $T_t = \tau_I$ [16]. No loop here has a slew limit. Derivative action is zero on every loop, because on a transient decided in six seconds a derivative on a four second measurement contributes noise and no anticipation, and the controller period is $\Delta t_c = 50\,\mathrm{ms}$.

The steam dump loop is direct acting ($a = -1$), because hotter coolant calls for more dump, as Section 1.3 notes, and its reference is deliberately offset above the temperature program, as shown in Equation (5.26).

$$\mathrm{SP}_d = T_{ref} + \Delta_d, \qquad \Delta_d = 1.5\,^{\circ}\mathrm{C} \tag{5.26}$$

Rod control and the steam dump are two integrating controllers acting on one measurement. Offsetting the dump's reference by a degree and a half hands the steady state to the rods, because once the bank has brought T_{avg} back to program the dump's error is negative and its output falls on its low limit, while leaving the transient untouched, since eighteen degrees of error does not care about one and a half. Section 5.6 reports what happens with the offset removed.

Two constraint controllers watch the same actuator and each publishes a cap rather than a contribution, as shown in Equation (5.27).

$$\begin{aligned} c_{pc} &= \mathrm{PID}_{+1}(p_{c,\mathrm{SP}},\ p_c;\ K_{c,pc},\ \tau_{I,pc};\ 0,\ \bar{c}) \\ c_L &= \mathrm{PID}_{+1}\big(L_{hi},\ \hat{L};\ K_{c,L},\ \tau_{I,L};\ 0,\ \bar{c}\big) \\ x_d &= \mathrm{clip}\big(x_{ff} + u_{fb},\ 0,\ \min(100,\ c_{pc},\ c_L)\big) \end{aligned} \tag{5.27}$$

where both constraint loops are reverse acting ($a = +1$) because a rising back pressure or a rising level must lower the cap being published, $p_{c,\mathrm{SP}} = 108\,\mathrm{mbar}$, $L_{hi} = 70\,\%$, and u_{fb} is the temperature trim from Equation (5.25). A cap wins whenever it needs to and is invisible otherwise, where a constraint controller that added to the demand would fight the loop it is protecting [4, 17]. Volume 4, Chapter 3 builds a larger selector of the same kind, with every integrator tracking the output actually applied.

Two implementation details in (5.27) matter more than the gains. The first is that the losing controller is never hard-tracked to a fixed value. Both constraint loops keep executing with their own back-calculation against the ceiling they are given, so their integrals stay consistent with the plant and a re-selection neither steps nor stalls. The second is the ceiling itself, given by Equation (5.28).

$$\bar{c} = \mathrm{clip}\big(\max(x_d,\ x_{ff},\ x_d^{unsat}) + \Delta_{ov},\ 25,\ 100\big) \tag{5.28}$$

where $\Delta_{ov} = 25\,\%$, x_d and x_d^{unsat} come from the previous controller scan and x_{ff} from the current one. An unselected constraint controller parked on its own hundred percent clamp while the working demand is twenty percent has eighty percent of dead travel to wind down through before it can bite, so it arrives as a relay rather than as a controller. Equation (5.28) holds it a stated margin above the working point instead, so it is always within one integral time of being useful, and the application reports an override as selected only when its cap is actually holding the unsaturated demand down. The dump limits are transferred into the feedback controller rather than applied after it, with $u_{min} = -x_{ff}$ and $u_{max} = \max(u_{min}, \min(100, c_{pc}, c_L) - x_{ff})$, so the anti-windup term in (5.25) sees the limit the sum will respect. Without that transfer the trim would wind up against a cap it cannot see.

5.5.4 The Reliefs, the Rods and the Feedwater

The atmospheric reliefs are a direct-acting loop ($a = -1$) on generator pressure against $p_{ARV} = 78.0\,\mathrm{bar}$ at $K_{c,a} = 30\,\%/\mathrm{bar}$ and $\tau_{I,a} = 10\,\mathrm{s}$, which places them between the working pressure and the first safety valve bank at 82.7 bar as the last modulating path before spring valves, and they are present only in the coordinated architecture. In the other three the relief loop tracks its manual station. Settings selected before Run reinitialize the plant and controllers exactly as Reset does. With Coordinated selected on either path, the relief integral starts at zero and the reliefs first open near 78.25 bar in the summer, failed-bank and suggested-challenge cases. A switch during an advancing run retains the manual-tracking history. A relief that has failed open is applied as a floor on position rather than as a demand, so the controller can wind fully down and the steam still leaves.

Rod control is a velocity station with a deadband, as in Chapter 2, and the runback is a bias on it, as shown in Equation (5.29).

$$\begin{aligned} \varepsilon &= \hat{T}_{avg} - T_{ref}, \qquad \delta = \mathrm{sgn}(\varepsilon)\max(0,\ |\varepsilon| - \delta_z) \\ v_z &= \mathrm{clip}(-K_z\delta + b_{rb},\ -v_{lim},\ v_{lim}) \\ b_{rb} &= -v_{max}\,\mathrm{clip}\left(\frac{\Delta\ell}{100}, 0, 1\right), \quad v_{lim} = v_{max} \\ &\qquad \text{whenever } \Delta\ell > \Delta\ell_{arm} \end{aligned}$$

where $\delta_z = 1.0\,^{\circ}\mathrm{C}$ is the deadband that stops the bank walking all day, $K_z = 0.060\,\%/\mathrm{s}$ per °C, $\Delta\ell = [\ell_{lag} - \ell]^+$ is the same load-loss signal that drives Equation (5.24), $\Delta\ell_{arm} = 10\,\%$ and $v_{lim} = 0.20\,\%/\mathrm{s}$ otherwise. The runback is the rod half of the feedforward, lifting the speed limit to the mechanical maximum of 0.50 %/s and biasing the drive inward in proportion to the load lost, so the bank moves before T_{avg} has changed by a degree. It stops at the insertion limit $z_{lim} = 18\,\%$ because going further eats into the shutdown margin the safety analysis assumes. Feedwater is a feedforward on total steam flow with a level trim, $w_{fw} = \mathrm{clip}(w_{tot} + K_{fwl}(L_{\mathrm{SP}} - \hat{L})w_0/100,\ 0,\ w_{cap})$, where w_{cap} is 130 % of rated flow normally and 20 % after a reactor trip.

Table 5.3 Scenarios and What Each Exercises

#	Scenario	Kind	What it exercises
1	Breaker opens, nobody acting	Demonstration	the clock, and reactivity feedback alone
2	Dump by hand on a 15 % runback	Experiment	the process gain of the actuator
3	Operator opens the dump late	Demonstration	reaction time against a stroke time
4	Feedback only	Demonstration	a fast disturbance and a slow measurement
5	Feedforward plus feedback	Experiment	anticipation from an upstream signal
6	Coordinated	Experiment	overrides, reliefs and the runback
7	50 % rejection, two dump modes	Experiment	armed versus continuous anticipation
8	Circulating water at 32 °C	Demonstration	a constraint that moves with the weather
9	One dump bank does not open	Demonstration	unmeasured loss of actuator authority
10	A relief sticks open afterwards	Experiment	a heat sink nobody is controlling
11	85 % dump capacity	Experiment	where the binding constraint moves to
12	Graded challenge	Challenge	the whole structure under three faults

Table 5.4 Measured Outcome of a Full Generator Load Rejection at $t = 5\,\text{s}$, Read Over $0 \le t \le 120\,\text{s}$. The Four Columns Differ Only in Which Signals Are Wired to the Steam Dump.

	Manual	Feedback	FF + FB	Coordinated
Peak p_{sg} (bar)	84.73	83.37	79.74	76.39
Margin to p_{SV} (bar)	−2.03	−0.67	+2.96	+6.31
Safety valve steam (t)	64.69	1.53	0.00	0.00
Relief steam (t)	0.00	0.00	0.00	0.00
Dump steam (t)	0.0	85.1	90.2	64.7
Time to 50 % open (s)	—	12.1	0.9	0.9
Peak T_{avg} (°C)	320.2	319.5	315.4	312.6
Peak p_c (mbar)	53	86	84	80
Time off program (s)	111	105	115	95
Reactor trip	no	no	no	no

5.6 Results

The twelve scenarios appear in Table 5.3 in the order a laboratory session would take them, with the mechanism each exercises. Each one pins down the starting mode, the architecture, the dump mode, the tuning, the disturbances and any timed events, and each carries a completion rule of the kind Section 1.6.4 describes. The challenge is judged on the case as given, so changing a model parameter or disturbance, or firing an extra learner disturbance event, voids the attempt. Every figure below comes from driving the shipped model at its physics step with no user interface attached and reading its own published outputs, with each configuration set after the case is loaded and before Run. Settings made before Run and carried through Reset give identical results. Every accumulated quantity carries the window over which it accumulated.

5.6.1 Four Architectures on One Rejection

Table 5.4 is the central result. The four architectures meet an identical disturbance, scenario 5 with only the architecture changed, with the breaker opening at $t = 5\,\text{s}$ from a settled full-power start, the same $K_{c,d} = 6.0\,\%/^{\circ}\text{C}$ and $\tau_{I,d} = 25\,\text{s}$, the same 40 % dump capacity, the same 4.0 s RTD and rod control on automatic in all four, read over $0 \le t \le 120\,\text{s}$.

Reading Table 5.4 across the row of safety valve masses gives the whole argument. The unattended dump discharges 64.69 t to atmosphere. A properly tuned proportional-integral controller on the same actuator reduces that to 1.53 t, a factor of 42 and still a reportable lift with a valve test attached to it. Handing the identical controller one additional input takes it to zero with 2.96 bar of margin, and the coordinated structure doubles that margin again. The mechanism is one row higher. Feedback alone

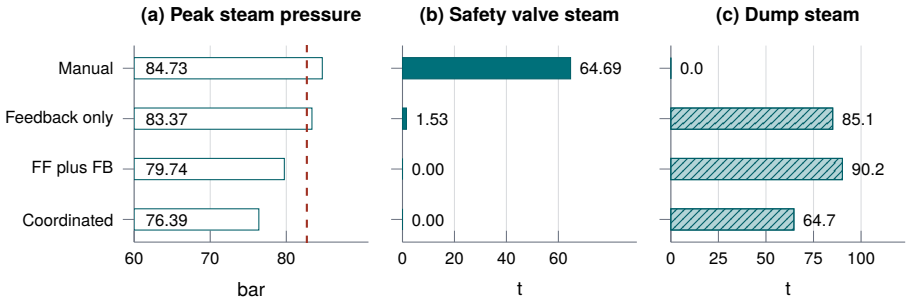

Figure 5.2 Table 5.4 drawn as bars: the full load rejection under the four architectures, read over $0 \leq t \leq 120$ s. (a) Peak steam generator pressure, with the first safety valve bank at 82.7 bar as the dashed line. (b) Steam through the safety valves. (c) Steam through the dump to the condenser. Notice that the two architectures carrying a turbine signal stay left of the dashed line and lift nothing, and that the coordinated one disposes of less dump steam because its runback makes less steam.

does not get the dump banks half open until 12.1 s after the breaker, which is almost exactly as late as a human operator, while both structures carrying a turbine signal are half open at 0.9 s. Figure 5.2 draws three rows of the table as bars.

Two rows read against the intuition they break. The dump steam row is largest for feedforward plus feedback at 90.2 t and smaller for the coordinated architecture at 64.7 t, because the coordinated structure runs the reactor back and has less steam to dispose of. The time off program row is worst for feedforward plus feedback at 115 s, which is not a defect of the measurement, because the counter accumulates whenever $|T_{avg} - T_{ref}|$ exceeds 2 °C in either direction and each architecture is measured against its own reference. A reference programmed from turbine load falls to the no-load value the moment the governor valves shut, so the count starts at once and runs until the coolant has cooled to within 2 °C of hot standby. Over this window every second of it is spent above program. The feedback only reference follows excore power down and stays higher, so the same measure of its run is not a like-for-like comparison. The coordinated architecture is the one that dips below program inside the window, and overcooling is a real failure mode on a plant with a moderator coefficient of –30 pcm/°C, because a colder moderator is positive reactivity.

The safety valve masses in the first two columns depend on the window. In the manual column the discharge is 11.0 t by $t = 30$ s, 34.7 t by 60 s, 64.7 t by 120 s and 75.7 t by 180 s, and it is still accruing at 99.5 t after ten minutes because the safety valves remain the only disposal path in service. The two architectures with a turbine signal discharge exactly nothing at any window.

Figure 5.3 puts the same result in the form the application keeps permanently on screen, because every kilogram per second of steam produced and not disposed of is stored in the secondary as pressure through Equation (5.10). The feedback only panel shows a disposal path collapsed to 60 kg/s one second after the breaker and not exceeding production again until 17.5 s after it, and the safety valve trace shows what pays for the intervening sixteen and a half seconds.

5.6.2 Tuning Does Not Reach It

The strongest form of the lesson is a negative result, because it establishes that the remedy is structural rather than a gain. Table 5.5 sweeps the feedback only architecture over the two things an engineer would ordinarily reach for, with every figure read over $0 \leq t \leq 90$ s.

Seven times the proportional gain buys 1.64 bar and a fivefold shorter integral time buys 0.48 bar. Replacing the four second RTD with an instrument eight times faster than anything a bypass manifold produces buys 0.45 bar, and making the measurement instantaneous buys 0.50 bar, which is the honest asymptote of the instrument remedy. Both remedies together, the gain of 40 with $\tau_I = 5$ s and the 0.5 s RTD, reach 81.30 bar, still 1.56 bar worse than the shipped gain with one wire added.

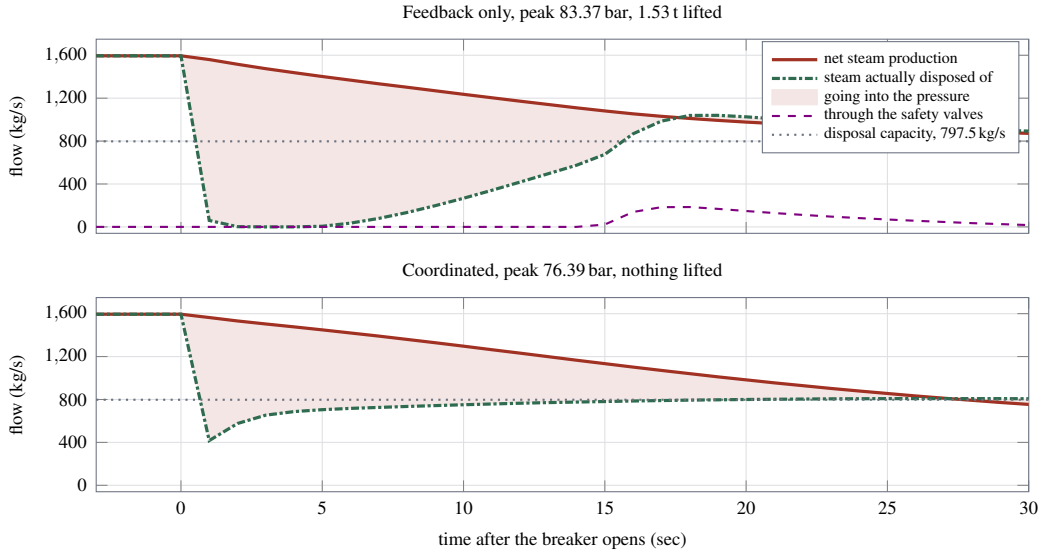

Figure 5.3 The gap between what the steam generators produce and what the plant disposes of, which is the entire transient. Both panels are the same full load rejection with rod control on automatic. Net production is the thick solid line, the steam disposed of the thick dash-dotted line, the flow through the safety valves the dashed line and the disposal capacity the dotted line, and the shaded band between the two thick lines is going into the pressure. Net production is gross evaporation less the condensation on the incoming subcooled feedwater, so it equals disposal exactly at the design point. Feedback only leaves a gap of 1500 kg/s one second after the breaker and does not close it until 17.5 s, 2.6 s after the safety valves begin to lift, while the coordinated structure closes it to 1147 kg/s in the first second and to zero at 27 s. Values are taken at the physics step.

Table 5.5 Feedback Only Under Every Available Remedy. The Bottom Row Is the Same Loop With One Additional Input and No Change of Gain.

Configuration	K_c	Peak p_{sg} (bar)	SV steam (kg)
As shipped, $\tau_{RTD} = 4.0$ s	6	83.37	1529
τ_I raised to 100 s	6	83.65	2226
τ_I cut to 5 s	6	82.89	385
Gain raised	15	82.31	0
Gain raised further	40	81.73	0
Gain raised, $\tau_I = 5$ s	40	81.70	0
RTD eight times faster, 0.5 s	6	82.92	454
RTD instantaneous, 0 s	6	82.87	311
Both remedies together	40	81.30	0
Turbine signal added, gain unchanged	6	79.74	0

The asymptote comes from Equation (5.23) rather than Equation (5.25), because the feedback only reference is generated from excore power and excore power is downstream of the disturbance, so the loop's setpoint stands still for the first several seconds no matter how good its measurement is. What is missing is not speed. It is a signal that moves when the disturbance moves. Volume 3, Chapter 2 reaches the same conclusion on a tank, where each structure rejects the disturbance that enters its own measurement path and leaves the other nearly untouched.

🕮 **Pause.** Try Exercise 5.2 before reading on.

The operator is the same result in human units. With the manual architecture and every dump bank driven fully open at a fixed delay after the breaker, the peak steam pressure is 80.97, 81.54 and 82.47 bar at delays of 1, 5 and 8 s with nothing lifted, then 82.97 bar with 613 kg lifted at 10 s, 83.47 bar with 1997 kg at 12 s and 84.73 bar with 8777 kg at 20 s. The entire difference between a normal day and a reportable event is the two seconds between eight and ten, and no training program makes a two second window reliable. Even acting at one second the peak is 80.97 bar, so a 40 % dump does not make this transient go away. It keeps it under the setpoint.

> ✦ **DID YOU KNOW? One Wire Beats Seven Times the Gain**
>
> From the simulator: raising the gain of the feedback only loop from 6 to 40 lowers the peak steam pressure only to 81.73 bar. Keeping the gain at 6 and wiring in the turbine signal lowers it to 79.74 bar, with nothing through the safety valves (Table 5.5).

5.6.3 Where the Anticipation Actually Lives

The result that surprises practitioners is that the explicit anticipatory term of Equation (5.24) is worth almost nothing at the design tuning. Removing it entirely, with $K_{ff} = 0$ against 1.2 at $K_{c,d} = 6.0\,\%/^{\circ}\mathrm{C}$, changes the peak steam pressure from 79.74 to 79.79 bar and lifts nothing in either case, read over $0 \le t \le 120\,\mathrm{s}$. The reason is the scale mismatch measured in Section 5.4. An 18.3 °C setpoint displacement at 6 %/°C asks for 110 % of a valve that stops at a hundred, so the trim saturates on the reference alone and the summing block has nothing left to add. Detune the trim to $K_{c,d} = 2.0\,\%/^{\circ}\mathrm{C}$, where the same displacement asks for 37 % of valve and the unsaturated demand without the term is only 83 % when the pressure peaks at $t = 27.5\,\mathrm{s}$, and the same term is the difference between 83.40 bar with 1554 kg lifted and 79.75 bar with nothing. The explicit feedforward is not what rejects this disturbance. It is what makes the structure insensitive to its own tuning, which is exactly the property to want on a transient nobody rehearses.

The same term earns its keep outright on a partial event, where the reference step is too small to saturate anything. On the 50 % rejection at $K_{c,d} = 2$, the banks never get past 24 % open with $K_{ff} = 0$ and the plant spends 61 s off program, against 65 % open and 42 s with $K_{ff} = 1.2$, read over $0 \le t \le 300\,\mathrm{s}$.

5.6.4 Attacking the Source Rather Than the Consequence

Of the four things the coordinated architecture adds, the reactor runback of Equation (5.29) does most of the work on a temperate day. Disabling it by raising the arming threshold to 95 % of load lost, the top of its range and above the 92.3 % the load-loss signal reaches, with everything else identical, moves the peak steam pressure from 76.39 bar to 79.10 bar and forces the atmospheric reliefs to discharge 2.3 t where they discharged nothing, read over $0 \le t \le 150\,\mathrm{s}$. The runback is worth 2.71 bar of the 3.35 bar the coordinated architecture gains, it leaves the reactor at 16.5 % of rated power against 31.1 % with the runback disabled, and alone among the four additions it reduces the steam rather than finding somewhere to put it. The condenser override does nothing at all on either run, which is correct behavior rather than a defect, because the peak back pressure is 80 mbar (83 mbar without the runback) against a setpoint of 108 mbar and an interlock that begins cutting the dump back at 115 mbar. An override visible during a normal transient has been set wrongly.

🕮 **Pause.** Try Exercise 5.3 before reading on.

5.6.5 The Constraint That Moves With the Weather

Scenario 8 is the feedforward plus feedback architecture, exactly the configuration with 2.96 bar of margin in Table 5.4, on a day when the circulating water arrives at 32 °C instead of 21 °C. Nothing has failed and nothing alarms. Read over $0 \leq t \leq 150$ s the peak steam pressure is 83.03 bar and 1.2 t goes through the safety valves.

The mechanism is entirely in Equations (5.16) and (5.17). Back pressure peaks at 123 mbar, the interlock takes hold, and at the worst moment the plant is allowed only 78 % of the dump it thinks it has. The controller has its valves at a hundred percent and is getting three quarters of the flow, and there is no flow measurement on a steam dump line to tell it. Switching to the coordinated architecture on the same day gives 79.98 bar with nothing lifted, because the condenser override caps the demand so that back pressure holds at 110 mbar and the plant never enters the interlock band, and because the atmospheric reliefs modulate and take 3.5 t out of the roof. Both initialization paths give these results. The override allows the plant *less* dump than the interlock would have permitted at the worst moment and produces a better transient anyway, because a controller modulates where an interlock ratchets.

> **TRY IT IN THE APP** **8 · A Hot Afternoon on the Intake**
>
> Load scenario 8, with the circulating water at 32 °C, and press Run. Watch the condenser back pressure: it peaks at 123 mbar, the interlock takes hold, and 1.2 t goes through the safety valves. Reset, select the coordinated architecture before the breaker opens and run again: the peak is 79.98 bar with nothing lifted.

Losing a dump bank is worse than losing eleven degrees of circulating water, for the same structural reason. With one of the four banks failing to open, the feedforward plus feedback architecture peaks at 83.31 bar with 3.3 t lifted over the same window, because the controller drives its output to a hundred percent, receives 478 kg/s of capacity instead of 638 kg/s, and cannot tell. The coordinated architecture recovers it to 80.01 bar with nothing lifted, and the recovery comes from the runback and from 3.5 t through the reliefs rather than from dump capacity, which did not exist.

Scenario 11 fits the same unit with an 85 % dump and makes the point from the other direction. The steam pressure becomes a non-event at 68.59 bar, fourteen bar clear of the safety valves. The condenser does not. Solving the fixed point of Equations (5.16) and (5.17) with the banks demanded wide open at the design header pressure gives a settled usable dump fraction of 100 % on 12 °C circulating water, 85 % on 21 °C, 71 % on 26 °C and 55 % on 32 °C, with back pressures of 97, 120, 125 and 131 mbar respectively. A 40 % dump loses at most 0.5 % of its capacity in that range.

At 32 °C an 85 % dump delivers 744 kg/s, only 17 % more than the 638 kg/s a 40 % dump delivers unconstrained on a cool day. The duty explains it. Demanded wide open a 40 % dump puts 1686 MW into a condenser rated for 1940 MW, which is 87 % of design duty from a plant generating nothing, and an 85 % dump puts in 3583 MW, or 1.85 times design. Above roughly 40 % of rated flow the binding constraint stops being valve capacity and becomes condenser duty, and buying more dump valves buys nothing. The coordinated architecture on the same big dump holds back pressure at 112 mbar so the interlock never engages.

5.6.6 The Two Dump Modes and a Partial Rejection

Scenario 7 sheds half the turbine load with the machine still synchronized, so the P-4 permissive never appears. Read over $0 \leq t \leq 200$ s, load rejection mode follows the program for the load actually carried, opens the banks to 99 % and settles T_{avg} at 302.5 °C against a program of 302.0 °C, with a peak of 68.08 bar and 27 s off program. Turbine trip mode never opens the dump at all, peaks at 72.42 bar and spends 69 s off program, with the reactor's own temperature feedback and the rods absorbing the entire event. Nothing is broken, because the mode is doing what it is designed to do, which is nothing until the turbine has actually tripped, and that is why a plant carries both [10].

5.6.7 Two Structures Easier to Assert Than to Measure

Two of the structures described above carry a standard justification that does not survive measurement on this plant, and both are more useful once measured. The application states the measured behavior in its own control hints, its model dialog and its instructional guide, and asserts the first in its test suite.

The first is the level override. Equation (5.20) makes indicated level move with pressure, and the coordinated architecture carries a level cap on the dump for that reason. That cap never selects in any of the twelve scenarios. Measured over $0 \le t \le 300$ s on the full rejection, indicated level falls to 30.6 % in the manual architecture and 34.2 % in the coordinated one and never exceeds 50.2 % in any of the four, with the void term reaching −58.5 % of span while the inventory rises to 174 t. The reason is structural and it is a better teaching point than a swell would have been. Indicated level swells when pressure falls, and on this transient the pressure never falls below where it started, because the no-load steam pressure of 76.00 bar is above the full-power value of 61.48 bar. A load rejection on this plant can only shrink the level, so the nearer of the two level limits is the low-low trip at 17 %, which the manual case approaches to within 13.6 % of span. Re-aiming the same cap at that limit would be worse than leaving it where it is, because a cap on the dump is the wrong actuator for a falling level. Shutting the dump raises the pressure and collapses the voids further, which drives indicated level down rather than up, and low level on this plant belongs to feedwater and to the auxiliary feedwater floor that starts below the low-low trip. An override protecting against a direction the transient cannot travel is correct engineering that happens to be inert on this event, and saying so is more useful than implying that it fires. The application keeps the loop in service, says in the control hint that a load rejection cannot select it, and asserts both halves of that in its test suite: the cap must not select on any of the four architectures, and it must still select and hold when its limit is set below the level being controlled to.

The two scenarios that do reach the high-high level trip at 78 %, 1 and the shut-dump run of 2 (at $t = 182$ s), reach it by filling rather than by swelling, and the distinction matters because the two have different cures. In scenario 1 feedwater is on manual at 100 % of rated while gross evaporation collapses to 1325 kg/s by $t = 40$ s, so the inventory rises from 135 to 179 t by $t = 62$ s, indicated level peaks at 79.4 % at $t = 63$ s and the trip isolates feedwater at $t = 60.4$ s. The void term stays between 22 and 39 % of span negative from $t = 30$ s to 120 s, which is the signature of a fill and the opposite of a swell, and the dump is shut in both runs anyway, so a cap on the dump could not have helped.

The second is the reference offset Δ_d of Equation (5.26). Two integrating controllers on one measurement is the classical setup for a limit cycle, and the offset is easy to justify as the cure for one, but no cycle occurs at either setting on this plant. Driven over $300 \le t \le 2400$ s on the coordinated architecture, $\Delta_d = 1.5\,^{\circ}\mathrm{C}$ gives a dump position range of 12.2 % and a steam pressure range of 2.73 bar with one crossing of the mean, and $\Delta_d = 0$ gives 2.6 % and 0.50 bar, also with one crossing. Neither is a cycle. What the offset does do is exactly the structural thing it was put in for, because after forty minutes the dump settles at 1.6 % of valve with $\Delta_d = 1.5\,^{\circ}\mathrm{C}$ and at 49.1 % with $\Delta_d = 0$. The offset is what makes the valves shut and hands the steady state to rod control, at the price of a 1.5 °C standing offset in T_{avg} and a bank driven onto its 18 % insertion limit, where $\Delta_d = 0$ leaves the average temperature exactly on program with the bank at 45.3 % and the dump holding half open indefinitely. That is a real trade between a shut valve and an on-program temperature, and it is a question of which controller owns the steady state rather than a question of stability. Neither of the two settings is pathological and the choice between them is an engineering one, which is a better teaching point than a pathology.

5.6.8 Failures Nothing Reports, and the Challenge

Scenario 10 sticks an atmospheric relief at 40 % open forty seconds after a coordinated rejection has already been handled. The dump closes to 8.0 % and lets the relief do the work, and the average coolant temperature comes back to its program and stays there, which is exactly the robustness that

hides the problem, because the temperature loop cannot tell one steam path from another. The bound depends on where the window starts. At $t = 50\,\text{s}$ the deviation is 11.5 °C, which belongs to the tail of the rejection that began at $t = 5\,\text{s}$ rather than to the relief that stuck at $t = 45\,\text{s}$. From $t = 120\,\text{s}$ to $t = 300\,\text{s}$ the deviation stays within 2.35 °C, worst at $t = 141.5\,\text{s}$, from $t = 180\,\text{s}$ onward within 1.25 °C, and at $t = 300\,\text{s}$ it is 1.04 °C. Over the same five minutes 18.4 t of chemistry-controlled secondary water leaves through the roof and the reactor is driven down to 9.1 % of rated power by a valve nobody is controlling, with nothing out of limits and nothing alarming. The indication that would tell an operator to gag the valve is a mass balance, not a temperature. Tripling the dump valve stroke lag to 2.7 s, so the banks need nine seconds to reach full open instead of three, moves the feedforward plus feedback peak from 79.74 to 80.06 bar and the coordinated peak from 76.39 to 77.02 bar, because the demand arrived at one second.

The graded challenge combines a full rejection with circulating water at 28 °C and one dump bank out of service, and hands the reader turbine trip mode with $K_{c,d} = 2$, $\tau_{I,d} = 90\,\text{s}$ and half the feedforward gain. It is judged at $t = 360\,\text{s}$, and its temperature band applies over the last two minutes. As handed it peaks at 83.93 bar and puts 23.9 t through the safety valves over $0 \le t \le 360\,\text{s}$, and its worst T_{avg} error over the last two minutes is 4.23 °C, so it fails on both counts. The suggested configuration (coordinated, load-rejection mode, $K_{ff} = 1.2$, $\tau_{ff} = 12\,\text{s}$, $K_{c,d} = 8$, $\tau_{I,d} = 20\,\text{s}$, a firmer condenser override, $K_z = 0.08$ and the runback armed at 8 %) finishes at 80.01 bar with nothing through the safety valves, no reactor trip and no condenser trip, with T_{avg} within 1.27 °C of program at six minutes and a peak back pressure of 93 mbar. It passes on both initialization paths with a worst error of 1.98 °C over the judging window, 0.52 °C inside the band. It costs 3.5 t of secondary water through the reliefs and runs the reactor back to 3.8 % of rated power, and both are the price of a dump bank that did not open.

Figure 5.4 shows the feedback-only architecture fourteen seconds into the same rejection Table 5.4 scores. Reproducing that instant numerically gives $p_{sg} = 82.09\,\text{bar}$ against a first-bank setting of 82.70 bar, net production of 1110 kg/s against disposal of 575 kg/s, a dump position of 66.0 % on a demand of 73.5 % with the anticipatory term at zero, $T_{avg} = 317.88\,^{\circ}\text{C}$ with the RTD reading 315.46 °C, core power at 67.2 % and back pressure at 42 mbar, which agrees with every value legible in the image. The first bank begins to lift 0.9 s later, at $t = 19.9\,\text{s}$, and the deficit the inset shades is the quantity Figure 5.3 tracks through the whole event.

5.7 Technology in Practice

The tags of Figure 5.1 are, in a plant, resistance thermometers in the reactor coolant loops, pressure transmitters on the steam lines, the condenser and the turbine's first stage, banks of air-operated dump valves, and the controllers of a non-safety control system. Figure 5.5 draws the dump loop with that hardware in place of the schematic's symbols.

What Measures It

Average coolant temperature comes from resistance temperature detectors, which with thermocouples are the most common temperature sensors in nuclear plants [20]. Steam pressure, condenser back pressure and turbine first-stage pressure come from pressure transmitters, which convert the pressure force to an analog signal such as 4 to 20 mA or, in newer smart transmitters, to a digital serial bus signal [20]. The chapter carries TE-411 as a 4.0 s lag, the impulse pressure transmitter as a 0.30 s lag and PT-514 as a 0.50 s lag. The gap between the first two is the lesson of Section 5.4.

> **TECHNOLOGY** **Sensor: The Resistance Thermometer in the Coolant**
>
> A resistance temperature detector is a fine wire of pure metal, usually platinum or copper, whose resistance rises with temperature. It is not very sensitive, but it is very linear [21]. With the thermocouple it is one of the two most common temperature sensors in nuclear plants, set in contact with the fluid or in a thermowell [20]. A thermowell lets the element be changed without opening the process, but its metal adds a first-order lag,

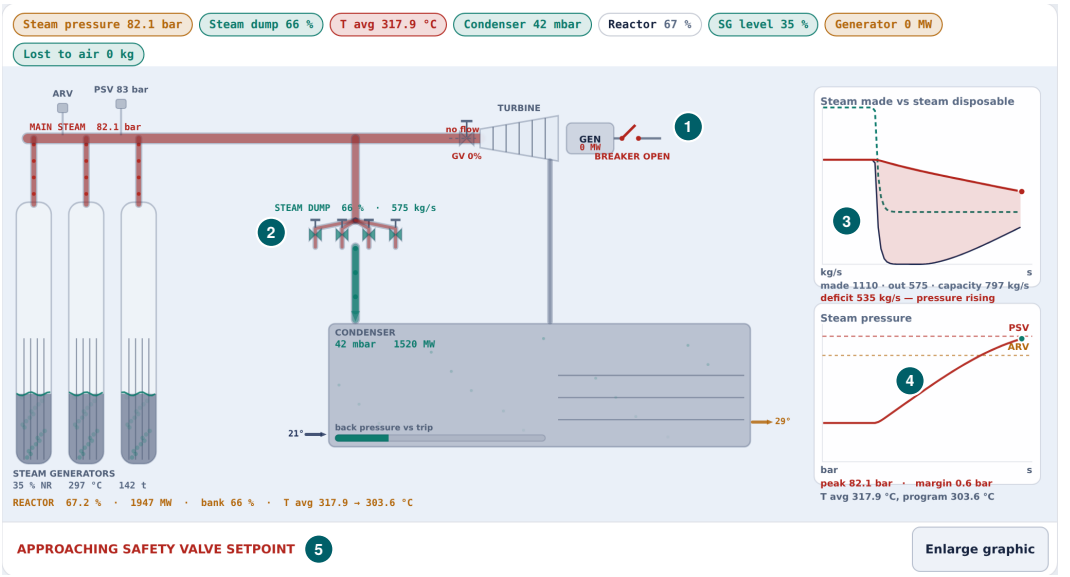

Figure 5.4 The feedback-only architecture at $t = 19$ s, fourteen seconds into a full load rejection and under a second before the first safety valve bank begins to lift, in the plant view of the app. The breaker is open and the governor valves read zero (1), so no steam is reaching the turbine. The dump banks (2) stand at 66.0 %, and the secondary is making 1110 kg/s net and disposing of 575 kg/s of it against a total capacity of 797 kg/s. The inset (3) shades that 535 kg/s deficit, and the pressure is taking it: the steam generators stand at 82.09 bar with the first safety valve bank set at 82.70, a margin of 0.6 bar (4), and the annunciator (5) reads the approach. T_{avg} is 317.88 °C, the reactor has run back to 67.2 %, the condenser is absorbing 1520 MW at 42 mbar, and nothing has yet been lost to atmosphere.

> and an air gap at the bottom of the well adds a second one. On a trend the reading trails the process, and a lag severe enough can make a feedback loop oscillate [21]. TE-411 stands for all of this as a single 4.0 s lag.

What Moves It

The final elements are the four dump banks, the atmospheric reliefs and, through the runback, the control rods. Steam dump to the condenser is listed among the non-safety control functions of a pressurized water reactor, beside turbine control, the main steam atmospheric relief valves and rod control [22]. The main steam safety valves sit outside every controller, which is why Equation (5.15) gives them no command. In the chapter each dump bank strokes as a first-order lag of 0.9 s, an idealization with no dead band and no stiction.

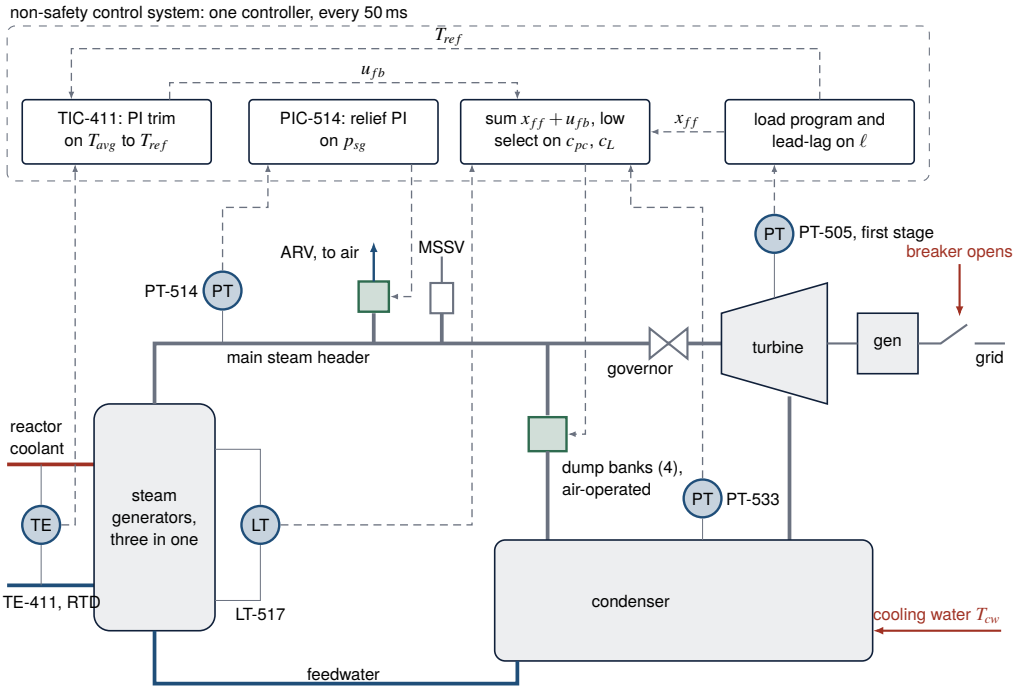

Figure 5.5 The steam dump loop of Figure 5.1 with plant hardware in place of the symbols. Sensors are circles and the two final elements the controller moves, the dump banks and the atmospheric relief, are green squares. Dashed arrows are signals. The temperature trim reads TE-411 between the hot and cold legs, the load program reads PT-505 on the turbine's first stage, and the low selector caps the dump from PT-533 and LT-517. The governor valve and the main steam safety valves (MSSV) take no command from this controller, and the rod runback, the third final element, is not drawn. Notice that PT-505, on the turbine, is the one sensor that moves the instant the breaker opens.

TECHNOLOGY **Final Element: The Air-Operated Steam Dump Valve**

In an air-operated valve, air on a diaphragm pushes the stem against a spring. A converter turns the controller's 4 to 20 mA into air pressure, and a positioner adjusts the air until the stem reaches the demanded position. In some safety-related designs, solenoid valves dump the air and the spring returns the valve to its fail-safe position [22]. Dump valves to the condenser must stroke very quickly on a turbine load rejection, yet hold stable pressure control during startup and shutdown. At one unit, digital positioners on these valves had more air capacity than the analog ones they replaced, so the volume boosters could be removed [22]. The chapter's dump banks need three seconds from shut to full open.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, under a supervisory layer that sends setpoints to the field controllers [23]. The chapter's controller executes every 50 ms over a 10 ms physics step. The structure is a set of function blocks in one controller: a temperature PI whose reference is the load program, a lead-lag on the load signal summed with it, two constraint PI blocks that publish caps, and a low selector that applies the smallest, with the applied limit fed back to every integrator. The runback is a bias on the rod velocity station, and turbine trip mode is armed by the P-4 permissive from the protection system.

> **IN PRACTICE** **A Rejection That Ends at House Load**
>
> Most plants are designed to ride out a limited, rapid load rejection by cutting back reactor power without a reactor trip, on the provision that such an event is rare. A unit able to drop to house load, with the grid disconnected and the generator supplying only the plant's own auxiliaries, typically brings the reactor down to about 20 to 30 percent of rated power and sends the excess heat to the condenser. Such a period is usually limited to a few hours, by constraints that include the capacity of the condensation system [24]. The runback of Equation (5.29) is the reactor's half of that cutback.

What Is Modernizing It

Three developments reach this loop. Digital positioners on control valves diagnose air leakage, friction and dead band, so that degradation is found before an upset, and they allow partial stroke tests on line [22]. On-line monitoring mines the data the control system already records for underlying changes in the relationships between measured parameters. It can show that a sensor has not drifted since its last calibration and so extend the calibration period [20]. And control systems increasingly resemble office computing as they adopt its computers, operating systems and network protocols, which gives them far less isolation from the outside world and a greater need to be secured [23]. Operating nuclear plants are moving from analog to digital control through individual system upgrades [20], so each upgraded dump controller joins that network.

5.8 Limitations

Nothing here is a safety analysis and nothing here is a plant model. Point kinetics, an absent pressurizer, one lumped secondary standing for three steam generators, fitted saturation properties and a correlation where a condenser belongs together rule out any statement about a real unit's peak secondary pressure, its safety valve sizing, its margin to departure from nucleate boiling, or whether some particular plant would ride out some particular load rejection. A licensing calculation for this transient uses a qualified system code against acceptance criteria written for the purpose [13], and the correct use of the numbers in Section 5.6 is as comparisons between control structures on one model, never as absolute predictions.

Three specific consequences follow. The safety valve masses are indicative rather than exact, because Equation (5.15) is a continuous ramp with no latch and no blowdown. The condenser results state the shape of the constraint correctly and its magnitude only to the accuracy of a fitted terminal temperature difference, so they establish that condenser duty becomes binding above roughly 40 % of rated dump flow without fixing the crossover point [12]. And the whole character of the transient depends on the moderator temperature coefficient, which is a constant here and runs from near zero at beginning of cycle to about −50 pcm/°C at end of cycle in a real core, so the reactivity feedback that buys the first ten seconds in Section 5.4 is worth much less early in a fuel cycle than the shipped value suggests.

5.9 Discussion

The engineering conclusion of Table 5.4 is not that feedforward is good. It is that the peak steam pressure on a load rejection is set by the instrument list rather than by the controller settings, and that the decision fixing it is made at design time when somebody chooses whether to run a cable from the turbine first-stage pressure transmitter to the steam dump controller. A factor of seven in gain is worth 1.64 bar, an instrument eight times faster than physically available is worth 0.45 bar, and the cable is worth 3.63 bar. Every disposal path in the plant addresses the consequence of the disturbance and the runback addresses its source, and on the shipped constants the source is worth more than any of the consequences at 2.71 bar of the 3.35 bar the coordinated architecture gains, because the same load-loss signal that opens the dump valves can drive the control rods and the anticipatory measurement is reused rather than duplicated.

Condenser back pressure is the quietest constraint in a steam plant. It appears in no control narrative, it varies with the weather and with tube cleanliness, and through Equation (5.17) it silently rescales the authority of the one actuator the whole transient depends on. Eleven degrees of circulating water temperature turned a transient with 2.96 bar of margin into one that lifted 1.2 t, with nothing failed and nothing alarmed. The remedy is not more valve, because an 85 % dump delivers only 744 kg/s on 32 °C water against 638 kg/s from a 40 % dump on a cool day. The remedy is a controller in front of the interlock, holding the plant off protection so that a ratcheting constraint is never approached, and that principle transfers to any plant where protection and control share an actuator. Methods beyond the implemented loops fit the same model, since a nineteen-state balance carrying a hard dump capacity, a moving condenser constraint and a rod insertion limit is a constrained model predictive control problem in which the low selects of Equation (5.27) turn into inequality constraints and the anticipation of Equation (5.24) turns into an explicit disturbance forecast from the impulse pressure measurement [17].

OPEN CHALLENGE Steam Bypass as a Daily Tool

A plant asked to cut its electrical output for a short time can bypass the turbine and send steam to the condenser while the reactor stays near full power, provided its feedwater, condenser and valves are designed for it [24]. The price is wear. Redirected steam loads the condenser shell and tubes, the bypass lines, the turbine stop and throttle valves and the steam spargers inside the condenser, and more so when bypass is the only means of flexible operation [24]. Progress would be a dump controller that weighs this wear against moving the reactor, the trade the runback of this chapter makes on a single event.

5.10 Conclusion

This chapter studied a pressurized water reactor secondary side under a generator load rejection as a working laboratory for a disturbance that arrives faster than any measurement of its effect can develop. Nineteen states are integrated against saturation property fits and a main steam header solved in closed form, covering six delayed-neutron precursor groups under a prompt jump, two decay-heat groups, fuel and coolant temperatures, a lumped secondary saturation temperature with its water inventory and its sub-surface steam holdup, a control bank position, four valve actuators and a condenser back pressure. Shut in at full power the secondary heats at 2.075 °C/s, passes the no-load pressure of 76.00 bar at 9.62 s and reaches the first main steam safety valve at 82.7 bar at 14.80 s, against dump valves that stroke in three seconds and an average-temperature RTD with a four second response. On one identical rejection the unattended dump peaks at 84.73 bar and discharges 64.69 t to atmosphere within two minutes, a properly tuned temperature loop with no turbine instrument peaks at 83.37 bar and discharges 1.53 t, the same loop with turbine impulse pressure wired to it peaks at 79.74 bar and lifts nothing, and a coordinated structure with two low-selected overrides and a reactor runback peaks at 76.39 bar. Seven times the gain reaches 81.73 bar and an instantaneous RTD reaches 82.87 bar, so neither tuning nor instrumentation closes a gap that only a measurement of the disturbance closes. The reactor runback is worth 2.71 bar of the coordinated architecture’s 3.35 bar advantage, and on 32 °C circulating water the plant’s own condenser interlock removes 22 % of the dump capacity at the moment it is needed.

A fuller model would add a pressurizer so that the primary pressure transient and its spray and heater actions appear, three separately modeled steam generators so that loop asymmetry is visible, latching safety valves with blowdown so that the discharge masses become quantitative, a real three-element feedwater system, and a burnup coordinate so that the moderator coefficient moves through a fuel cycle. What the steam dump plant gives a control engineer is fifteen seconds in which structure beats gain.

KEY POINTS

- When the turbine valves shut, the steam generators keep making steam, and the pressure reaches the first safety valve in under fifteen seconds.
- Seven times the gain on the four second RTD reaches 81.73 bar, while the same loop with turbine impulse pressure wired to it reaches 79.74 bar and lifts nothing.
- The reactor runback is worth 2.71 bar of the coordinated architecture's 3.35 bar advantage, because it reduces the steam rather than finding somewhere to put it.
- A controller in front of the condenser interlock beats the interlock: on 32 °C circulating water the coordinated structure lifts nothing where feedforward plus feedback lifts 1.2 t.

Where else this applies. A protective override on a fast transient also guards the compressor of Volume 4, Chapter 2, and Volume 4, Chapter 3 low-selects limiters over a governor as the dump's overrides do here.

References

[1] 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.

[2] K. J. Åström, "Limitations on control system performance," *European Journal of Control*, vol. 6, no. 1, pp. 2–20, 2000.

[3] J. L. Guzmán and T. Hägglund, "Tuning rules for feedforward control from measurable disturbances combined with PID control: a review," *International Journal of Control*, vol. 97, no. 1, pp. 2–15, 2024.

[4] F. G. Shinskey and J. P. Shunta, "Selective, override, and limit controls," in *Process Control*. Elsevier, 1995, pp. 112–118.

[5] J. T. Gravdahl and O. Egeland, "Modelling and control of surge for a centrifugal compressor with non-constant speed," in *Compressor Surge and Rotating Stall*, ser. Advances in Industrial Control. Springer London, 1999, pp. 143–176.

[6] L. Wang, C. Chen, L. Chen, Z. Li, and W. Zeng, "A coordinated control methodology for small pressurized water reactor with steam dump control system," *Energy*, vol. 282, p. 129008, 2023.

[7] E. Lin, Y. H. Cheng, and J. K. Hsiue, "Analysis of chinshan unit 1 full load rejection transient using RETRAN," *Nuclear Technology*, vol. 54, no. 3, pp. 380–384, 1981.

[8] M. D. Murrell, "Boiler simulation, PWR station dynamics," in *Boiler Dynamics and Control in Nuclear Power Stations*. Emerald Publishing Limited, 1973, pp. 4.1–4.9.

[9] T. W. Kerlin, E. M. Katz, J. Freels, and J. Thakkar, "Dynamic modelling of nuclear steam generators," in *Boiler Dynamics and Control in Nuclear Power Stations 2*. Emerald Publishing Limited, 1980, pp. 249–258.

[10] U.S. Nuclear Regulatory Commission, Human Resources Training and Development, "Westinghouse technology systems manual, section 11.2: Steam dump control system," USNRC HRTD, Rev. 0403, pp. 11.2–1–11.2–8, 2003, aDAMS Accession No. ML11223A294. [Online]. Available: https://www.nrc.gov/docs/ML1122/ML11223A294.pdf

[11] A. Logar, T. Depolt, and E. Gobrecht, "Advanced steam turbine bypass valve design for flexible power plants," in *2002 International Joint Power Generation Conference*. ASMEDC, 2002, pp. 43–49.

[12] D. M. Nightingale, "Design guidelines for the safe operation of steam surface condenser turbine bypass on combined cycle power plants," in *Proceedings of the ASME 2017 Power Conference*. American Society of Mechanical Engineers, 2017, p. V001T05A001.

[13] U.S. Nuclear Regulatory Commission, "Standard review plan 15.2.1–15.2.5: Loss of external load; turbine trip; loss of condenser vacuum; closure of main steam isolation valve; and steam pressure regulator failure," U.S. Nuclear Regulatory Commission, Washington, DC, Tech. Rep. NUREG-0800, Rev. 2, Mar. 2007, aDAMS Accession No. ML070300702. [Online]. Available: https://www.nrc.gov/docs/ML0703/ML070300702.pdf

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

[15] W. E. Reeves, "Technical evaluation of the proposed deletion of a reactor trip on a turbine trip below 50-percent power for the Beaver Valley nuclear power plant, Unit 1," Lawrence Livermore Laboratory, Livermore, CA, Tech. Rep., Dec. 1979.

[16] R. Hanus, M. Kinnaert, and J.-L. Henrotte, "Conditioning technique, a general anti-windup and bumpless transfer method," *Automatica*, vol. 23, no. 6, pp. 729–739, 1987.

[17] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

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

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

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

[24] 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.

5.11 Exercises

5.1 Concept check · 1 min

Equation (5.10) sets the heat capacity $(Mc)_s$ times the rate of change of T_s equal to the bundle duty less two steam terms. What does the equation describe?

(a) A balance on one delayed-neutron precursor group of the core

(b) An energy balance on the saturated secondary side of the steam generators

(c) An energy balance on the fuel node that receives the fission power

(d) The first-order lag of the steam dump valve position on its demand

5.2 Concept check · 2 min

On the full rejection, the feedback only architecture with $K_{c,d} = 40\,\%/^{\circ}\mathrm{C}$ and $\tau_{I,d} = 5\,\mathrm{s}$ peaks at 81.70 bar, while feedforward plus feedback at the shipped $K_{c,d} = 6\,\%/^{\circ}\mathrm{C}$ peaks at 79.74 bar. Which statement explains the gap?

(a) A high gain on a four second RTD makes the loop oscillate, so its peak cannot fall further

(b) The integral winds up against the full-open limit and holds the valves back after the breaker

(c) Its reference follows excore power, which lags the coolant, so no input moves when the breaker opens

(d) The dump valves stroke more slowly when the demand changes faster, so a higher gain cannot open them sooner

5.3 Concept check · 2 min

On a 21 °C day the coordinated architecture of scenario 6 peaks at 76.39 bar, against 79.74 bar for feedforward plus feedback. Which of the four additions does most of that work?

(a) The condenser back-pressure override, which caps the dump as the back pressure rises

(b) The steam generator level override, which holds the indicated level off its high limit

(c) The modulating atmospheric reliefs, which open above their 78.0 bar setpoint

(d) The reactor runback, which drives the bank in on the same load-loss signal

5.4 Theory and calculation · 6 min

Unit U-1 runs at its design point: 1595.0 kg/s of steam at $p_{sg} = 61.48$ bar, feedwater at 964.3 kJ/kg, a bundle duty of 2900 MW and $(Mc)_s = 1200\,\mathrm{MW\,s}/^{\circ}\mathrm{C}$. Use the saturation fits of Equation (5.1).

(a) Compute the latent and sensible terms of Equation (5.10) and show that the right-hand side is zero. *Checkpoint:* 2490 MW and 410 MW, summing to 2900 MW.

(b) The governor valves shut and every steam flow stops, with everything else held at the design point. Find the heating rate in degrees and, with the slope 0.906 bar/°C, in bar per second. *Checkpoint:* 2.075 °C/s and 1.88 bar/s.

(c) Find the saturation temperature at the first safety valve bank, 82.7 bar, and the time to reach it at the rate of (b). Compare it with the measured 14.80 s and give two reasons for the difference. *Checkpoint:* 297.57 °C, a rise of 20.56 °C, about 9.9 s.

5.5 Theory and calculation · 5 min

Section 5.4 measures the settled average coolant temperature at 85 % turbine load as 312.00 °C with the dump shut and 308.92 °C with it 60 % open.

(a) Find the average process gain of the dump over that travel. *Checkpoint:* −0.0513 °C/%.

(b) Taking the gain magnitude as 0.052 °C/%, find the controller gain that gives unity loop gain, and the loop gain of the shipped $K_{c,d} = 6\,\%/°\mathrm{C}$. *Checkpoint:* 19.2 %/°C and 0.31.

(c) The program of Equation (5.23) falls 18.3 °C when the impulse pressure collapses. Find the proportional demand this asks for at $K_{c,d} = 6$ and at $K_{c,d} = 2$, and the steady temperature change that a fully open dump can make. *Checkpoint:* 110 %, 37 % and 5.2 °C.

5.6 Theory and calculation · 6 min

The control bank of U-1 has a total worth $W_{tot} = 3000$ pcm with the S curve of Equation (5.5), sits at $z_{ref} = 68\,\%$ at full power, and the runback drives it at $v_{max} = 0.5\,\%/\mathrm{s}$. The moderator coefficient is −30 pcm/°C.

(a) Find the bank worth and the differential worth at z_{ref}. *Checkpoint:* 2472 pcm and 42.8 pcm/%.

(b) Find the reactivity rate of a full-speed runback at that position. *Checkpoint:* 21.4 pcm/s.

(c) Turbine-trip mode is armed on a machine that still carries half its load, so the dump aims at the no-load program of 291.7 °C instead of the half-load value of Equation (5.23). Find that half-load value, the reactivity the overcooling adds, and how long the runback would take to insert the same worth. *Checkpoint:* 300.85 °C, +274.5 pcm and about 12.8 s.

5.7 App challenge · 8 min

Load scenario 4, *Feedback only*, in Auto with Steam dump architecture on Feedback only. The breaker opens at 5 s. Use speed 2 and run each case for 60 s, then export the CSV file.

(a) Run the scenario as shipped. Read the Peak steam pressure and Safety valve steam tiles, and find in the file the first row at which `dump_pos_pct` reaches 50 %. *Checkpoint:* 83.4 bar and 1529 kg, within 5 kg. Half open at 17.15 s, 12.1 s after the breaker, within 0.1 s.

(b) Press Reset, set Dump temperature gain to 40 and Dump integral time to 5 s before Run, and run again. Read the same tiles and the last value of `metric_peak_psg_bar`. *Checkpoint:* 81.70 bar, within 0.01 bar, and 0 kg.

(c) The shipped feedforward plus feedback architecture peaks at 79.74 bar. Explain why seven times the gain leaves the loop 1.96 bar short of it. *Checkpoint:* The reference does not move with the breaker.

5.8 App challenge · 8 min

Load scenario 5, *Feedforward plus feedback*, in Auto with Steam dump architecture on Feedforward + FB. Run four cases of 60 s at speed 5, each set before Run from a reload of the scenario, with

Feedforward gain (Advanced) and Dump temperature gain as follows: 1.2 and 6, 0 and 6, 0 and 2, 1.2 and 2. Export the CSV file after each run.

(a) Read the last value of `metric_peak_psg_bar` and the Safety valve steam tile for each case. *Checkpoint:* 79.74 bar, 79.79 bar, 83.40 bar with 1554 kg lifted, 79.75 bar, each within 0.01 bar.

(b) Find the peak pressure the explicit feedforward term saves at each gain. *Checkpoint:* 0.05 bar at $K_{c,d} = 6$ and 3.65 bar at $K_{c,d} = 2$.

(c) Explain the difference with the proportional demand that the program step asks for at each gain. *Checkpoint:* 110 % against 37 % of valve.

5.9 App challenge · 8 min

Load scenario 8, *A hot afternoon on the intake*, in Auto. The circulating water is at 32 °C and the breaker opens at 5 s. Run each case to 150 s at speed 5 and export the CSV file.

(a) Run with Steam dump architecture on Feedforward + FB as shipped. Read Peak steam pressure, Safety valve steam and Peak back pressure, and the lowest value of `cond_derate` in the file. *Checkpoint:* 83.0 bar, 1223 kg within 5 kg, 123 mbar and 0.78.

(b) Press Reset, set Steam dump architecture to Coordinated before Run and run again. Read the same three tiles, Lost to atmosphere and the lowest `cond_derate`. *Checkpoint:* 80.0 bar, 0 kg, 110 mbar, 3519 kg within 10 kg, and 1.00.

(c) Explain why the coordinated structure does better while it lets the condenser take less steam than the interlock allowed. *Checkpoint:* A controller modulates in front of the interlock.

5.10 Application area · 4 min

A plant considers replacing its 40 % dump with one of 85 % of rated steam flow, $w_0 = 1595\,\mathrm{kg/s}$. Dump steam enters the condenser at $h_g(p_{sg,0})$ and leaves at 140 kJ/kg, and the condenser is rated for 1940 MW.

(a) Find the duty of each dump wide open at the design pressure, as a fraction of the condenser rating. *Checkpoint:* 1686 MW, 87 %, and 3583 MW, 1.85 times the rating.

(b) On 32 °C water the larger dump delivers 744 kg/s once the interlock settles. Find the usable fraction and compare the flow with the 638 kg/s of the smaller dump on a cool day. *Checkpoint:* 0.549 of its capacity, 17 % more flow.

5.11 Application area · 3 min

In the manual architecture every dump bank is opened fully at a fixed delay after the breaker. A delay of 8 s peaks at 82.47 bar with nothing lifted, 10 s lifts 613 kg and 12 s lifts 1997 kg. A procedure writer proposes to rely on the control room operator. Give two reasons from the chapter why the dump must be automatic, and why an automatic temperature loop alone is not the answer either.

5.12 Review the agent · 5 min

An AI assistant reviews Table 5.4 and writes: “Feedback only spends 105 s off program against 115 s for feedforward plus feedback, so it holds the temperature program better. Its safety valve lift is only 1.53 t. A plant that values temperature control should prefer feedback only.” Find the flaw in the argument and say how to check it on the app.

CHAPTER 6

Boiling Water Reactor Simulator for Coupled Loops With Sign-Reversing Dynamics

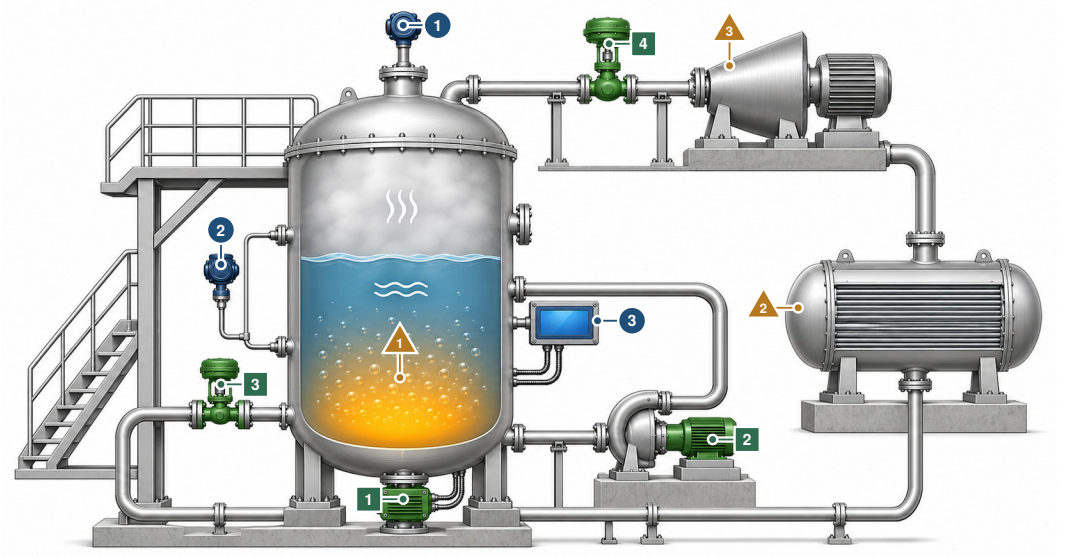

The controller moves 1 control rod drives 2 recirculation pump speed 3 feedwater valve 4 turbine control valves

The controller reads 1 steam dome pressure 2 vessel water level 3 reactor power

What disturbs it 1 boiling noise in the core 2 feedwater heating 3 turbine steam demand

Notice that all four controls, from the rod drives under the vessel to the turbine valve above it, act on the same boiling water whose bubbles also lift the level the transmitter reads. Illustration, not to scale.

IN PLAIN TERMS **Boiling Water That Steers Its Own Reactor**

The main idea. In this reactor, water boils inside the core and the steam drives the turbine directly. The steam bubbles slow the chain reaction that makes the heat, so anything that changes the boiling changes the power.

The challenge. Adding cooler water collapses some bubbles, so the level gauge first reads low while the vessel is filling. A controller that trusts it feeds the wrong way, and the cure is to measure the steam leaving and replace it as it goes.

Why it matters. Four controls push on the same bubbles, so how they are wired together matters more than how hard each one pushes.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 6.2, the try-it and fact boxes of Section 6.4, the try-it and fact boxes of Section 6.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 6.6, Section 6.7, Section 6.8 and the open challenge

The control question The steam voids in a boiling water reactor core are at once the missing moderator, the strongest reactivity feedback and a false part of the level reading, so every actuator moves every loop. How should the four loops be paired, and how much of the result comes from structure rather than from tuning?

Learning objectives

- Compute the rod worth, the integrating gain of the pressure path and the decay ratio from the model's equations.
- Explain shrink and swell from the split of the vessel inventory into liquid and sub-surface steam.
- Pair the four loops from the measured self-regulation, pressure coefficient and flow-to-power gain.
- Compare single-element and three-element feedwater on one load change and say why structure wins.
- Judge a gain-scheduling rule against a measured process gain.

Prerequisites Sections 1.3, 1.4.2, 1.4.4, 1.6.2 and 1.6.4, and Chapter 4. Section 1.6.3 for the predictive controller only.

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

See also Volume 6, Chapter 2 (a pairing decided before any loop is tuned) and Volume 6, Chapter 7 (a structure choice that changes what the plant does). Volume 7, Chapter 2 for a self-regulating plant that ignites, and Volume 7, Chapter 13 for a delay-driven instability whose boundary moves with the operating point.

6.1 Introduction

A block diagram hides the one fact that decides how a boiling water reactor is controlled. Water boils in the fuel channels themselves, and the resulting steam voids are simultaneously the missing moderator, the strongest reactivity coefficient in the core, the mechanism by which recirculation flow sets power, and the volume that a differential-pressure level transmitter counts as though it were water. Four transfer functions drawn side by side present that as four plants. It is one state. Moving the turbine valves, the feedwater valve, the recirculation pumps or the control rods moves the void fraction, so each of the four loops is a load on the other three, and the coupling runs through the physics rather than through a matrix that could be inverted away.

The misconception this simulator corrects is specific and common. When the vessel level indication falls after feedwater is added, or rises after feedwater is lost, the reading is taken for an instrument fault. It is not. A differential-pressure cell straddling a two-phase mixture reports the height of the mixture, and in the first seconds of a feedwater change the mixture and the water inventory move in opposite directions. The application computes both from the same inventory and draws them together, so the wrong-way motion appears as a measured swollen surface against a computed collapsed surface rather than as a numerator coefficient chosen by the author of a textbook problem. A measured feedwater step of +10 % drives the indication −40.9 mm while the vessel is gaining 193 kg, and no gain and no filter changes that.

The class of plant is wider than reactors. Three features define it: a measurement that is the sum of two states responding to the same actuator with opposite signs, a feedback path that the operator never closed and cannot open, and loop gains that migrate with operating point. A natural-circulation steam drum has the first [1], distillation composition control with an inverse-responding tray has the first and the third, and any plant with a right-half-plane zero has the bandwidth limit that follows from it [2, 3]. What a boiling water reactor adds is that the second feature dominates the other two. The plant is its own controller, at a measured feedback gain of 30.6 pcm per percent of power, so interaction between the loops is not a nuisance to be decoupled but the mechanism that delivers most of the control action [4, 5].

Reduced-order boiling water reactor models exist in quantity, written to answer stability questions [5] or optimal control questions [6], and feedwater control algorithms for these plants have been documented against measurements [7]. However, a model published as a set of state equations gives the student the couplings already identified and the signs already assigned, which is the part of the

problem that decides the control structure. This chapter builds the void fraction into an interactive boiling water reactor as a state rather than a parameter, so that the four loop pairings, the two sign reversals and the migration of the gains can each be measured on the running plant, and it quantifies what each of five control architectures achieves on the same disturbances. The point kinetics, the drift-flux void correlation [8, 9] and three-element feedwater control are all standard. The last has been ordinary practice on drum boilers and on these reactors since long before either was simulated [7]. The simulator carries five loop architectures over one actuator set, one sample period and one set of limits, with the switch made by re-seeding the controllers rather than by a transfer routine, so a comparison between structures compares information flow and nothing else. The scenarios are ordered so that the plant, not the controller, is the thing that goes unstable. Every balance, property fit, constant and numerical choice needed to rebuild the simulator appears in Section 6.3.

Section 6.2 traces how boiling water reactor control arrived at its loop assignments. Section 6.3 states the plant model, its parameters and its assumptions. Section 6.4 measures the couplings and the two sign reversals on the running plant and places them on the operating map. Section 6.5 gives the loop pairing, the PID loops, the six architectures, the gain scheduling, the protection and the predictive controller. Section 6.6 compares single-element and three-element feedwater, follows a runback into the exclusion region, takes the graded challenge apart one change at a time and sets the predictive controller against the coordinated loops. Section 6.7 sets the loops in their plant hardware. Sections 6.8 and 6.9 state where the model stops and what the results mean beyond this plant, and Section 6.10 sums up. The exercises follow in Section 6.11.

6.2 Control of Boiling Water Reactors

Boiling in a reactor core was originally regarded as a hazard rather than a design feature, on the argument that a moving steam-water interface would make reactivity unpredictable. The experimental reactors of the 1950s settled the question the other way. Boiling proved to be self-limiting because the voids themselves remove moderator, and the control system that followed was built around a single-loop direct cycle in which the steam raised in the core drives the turbine with no secondary side. The earliest control account of such a plant describes the arrangement that survived: the turbine admission valves are assigned to reactor pressure rather than to power, because the void feedback already regulates power and the pressure is what must be held for that feedback to be predictable [10]. The assignment is still counterintuitive to anyone who has run a fired boiler, where the firing rate answers load and the throttle answers steam demand.

> ⧖ HISTORY **Boiling on Purpose, From Test Reactor to Power Plant (1954 to 1956)**
>
> The first test boiling water reactor showed that a reactor with water boiling in its core regulates itself and is inherently safe. It was destroyed on purpose in 1954, to show that extraordinary measures were needed to produce a damaging reaction. A pressurized successor was given a turbine generator, and on the night of July 17, 1955, its steam lit a nearby town. The Experimental Boiling Water Reactor followed, built in 19 months between May 1955 and December 1956 to prove the direct cycle, in which steam raised in the vessel goes straight to the turbine. It made 20 megawatts of heat and 5 of electricity and ran in parallel with the utility grid. Swings in demand went to an automatic steam valve that bypassed the turbine, so the reactor power did not have to follow them [11].

The void coefficient is the reason and it cuts both ways. It gives the plant strong inherent load-following, and it makes every quantity that changes the boiling into a reactivity input. In this model the coefficient is -1.2×10^{-3} in $\Delta k/k$ per percent of void, which at a delayed-neutron fraction of 0.0056 is 21.4 ¢ per percent of void. Because the void fraction is set by the flow through the channels, recirculation flow became the load-following actuator. Slowing the recirculation pumps raises the void fraction, subtracts reactivity and takes the reactor down with no rod motion at all, which covers roughly the top third of the power range in seconds rather than in the minutes a rod bank needs. Optimal-control studies of boiling water reactor load following are written around exactly that

actuator and the rod bank behind it [6]. The cost is that the recirculation system rather than the rod drive becomes the primary control element, and that the flow-to-power gain depends on where the plant sits on the operating map.

The same coefficient produced the industry's defining control event. A boiling channel is a feedback loop in its own right, because the pressure drop across it depends on the void fraction, the void fraction depends on the enthalpy the fluid absorbed one transit time earlier, and the neutronics close the loop through the void coefficient. Density-wave oscillations of this kind were catalogued for two-phase systems generally in the early 1970s and analyzed for the coupled neutronic case through the 1980s [5]. The event of March 1988 at LaSalle Unit 2 is told in the history box below [12]. The response was an exclusion region drawn on the operating map itself, entered by reducing flow rather than by raising power, and a requirement to leave it immediately by adding flow. Because the mode announces itself as ordinary neutron noise that stops settling rather than as an alarm, monitors that estimate the decay ratio from the neutron noise time series were developed for the purpose [13].

> ⧖ HISTORY **Two Pumps Trip, and the Core Oscillates (1988)**
>
> In March 1988 a boiling water reactor in the United States was running at about 85 percent power when a valving error in a routine test tripped both recirculation pumps. Power fell sharply as the core went over to natural circulation, to 40 percent within half a minute [14, 15]. Less steam meant less extraction steam for the feedwater heaters, some heaters were isolated, and the colder feedwater added reactivity [15]. Within five minutes the power began to oscillate, and about seven minutes after the pump trip the reactor scrammed on overpower at about 118 percent of full power [14]. A laboratory simulation traced the instability to low core flow, moderately high power, low feedwater temperature and a power shape peaked toward the bottom of the core [15].

Vessel level control took a separate path and arrived at the same place as the fossil boiler industry. Level alone onto the feedwater valve is the obvious loop and it is adequate at steady load. On a load change it is not, because the first several seconds of the measurement have the wrong sign, and a controller that reacts to them makes the excursion worse. The answer is three-element control: the steam flow measurement is fed forward to a feedwater flow controller so that every kilogram leaving the vessel is replaced as it leaves, and the level controller is reduced to a slow trim on the residual inventory error. Feedwater control algorithms for boiling water reactors have been documented in that form and compared against plant data [7], and drum boiler practice reached the same conclusion once the steam held below the water surface was recognized as the source of the reversal [1]. The steam generator of Chapter 4 carries the same three structures on a drum. The fix is structural rather than a matter of tuning because the right-half-plane zero the balances produce cannot be canceled, and it fixes a ceiling on closed-loop bandwidth that raising the gain only worsens [3]. Section 1.4.4 derives that ceiling for a drum.

6.3 Reactor Model

The unit is a BWR/4-class plant rated at 3330 MW thermal and 1100 MW electric. Saturated water leaves the steam separators at the top of the vessel, flows down the downcomer annulus where it mixes with the feedwater, is driven through the core by the jet pumps, boils in the fuel channels, and the steam is separated and dried and passes to the turbine through the control valves. Figure 6.1 shows the arrangement, the four manipulated variables and every symbol used in the equations that follow. The narrow-range level transmitter straddles the water surface, so it reports the swollen level while the water inventory sits at the collapsed level, and both are drawn.

Table 6.1 classifies what the four loops interact with. Seventeen states are integrated. The plant is expressed in seconds, kilograms, cubic meters, MPa, kJ/kg and °C, with level carried in meters and displayed in millimeters.

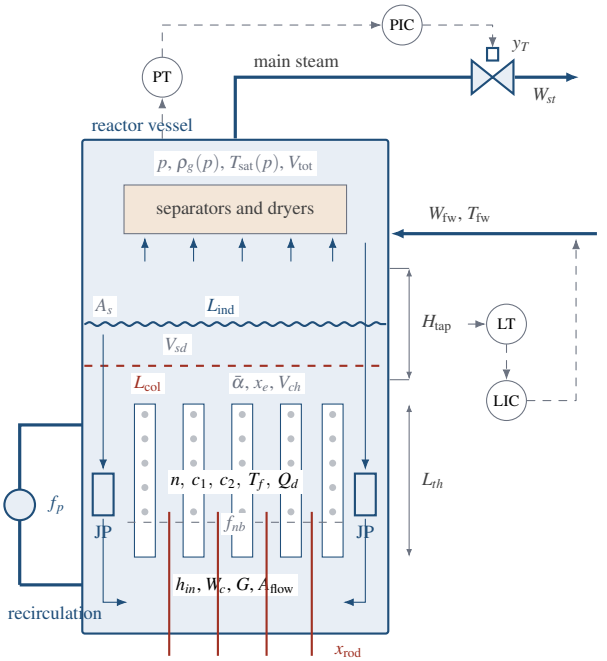

Figure 6.1 Reactor vessel, core channels and the four manipulated variables. Saturated water separated at the top of the vessel flows down the downcomer annulus, mixes with feedwater W_{fw} at temperature T_{fw}, enters the core at enthalpy h_{in} and is driven at flow W_c and mass flux G through channels of thermal height L_{th} by jet pumps at drive fraction f_p. Boiling begins above the non-boiling fraction f_{nb} and produces core-average void $\bar{\alpha}$ and exit quality x_e against the point-kinetics states n, c_1, c_2, the fuel temperature T_f, the decay-heat fraction Q_d and the rod bank x_{rod}. Steam leaves through the turbine control valves at position y_T, which PIC holds the dome pressure p with. The transmitter LT spans H_{tap} across the surface and reports the swollen level L_{ind}, while the water inventory sits at the collapsed level L_{col} and the difference is the sub-surface steam volume V_{sd}.

Table 6.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range	Symbol	Role	Nominal	Range
p dome pressure	CV, meas.	7.170 MPa	6.85–7.45 MPa SP	$\bar{\alpha}$ core void	state	40.32 %	0–97 %
L_{ind} level	CV, meas.	0 mm	±250 mm SP	h_{in} core inlet	state	1227.7 kJ/kg	—
Q_{th} power	CV, meas.	100 %	40–105 % SP	W_c core flow	state	13 100 kg/s	600–24 000 kg/s
y_T turbine valves	MV	85.0 %	0–100 %	f_p drive fraction	state	0.950	0–1.02
W_{fw} demand	MV	100.0 %	0–130 %	p dome pressure	state	7.170 MPa	4.6–9.4 MPa
f_p pump speed	MV	95.0 %	0–100 %	M_v vessel inventory	state	227 209 kg	—
x_{rod} bank	MV	35.0 % in	0–100 %	W_{fw} flow	state	1800.2 kg/s	—
T_{fw}	DV, meas.	216 °C	150–225 °C	T_{fw} state	state	216 °C	—
FW pump capacity	DV, meas.	130 %	0–140 %	y_T valve position	state	0.850	0–1
turbine admission	DV, meas.	1.00×	0.30–1.10×	x_{rod} state	state	35.0 %	0–100 %
ρ_{ext}	DV, unmeas.	0 pcm	±200 pcm	n neutron density	state	1.000	—
boiling noise	DV, unmeas.	0×	0–4×	c_1, c_2 precursors	state	124.4, 1716.5	—
T_f fuel temperature	state	699.5 °C	—	z_o void oscillation	state	0	±0.35
Q_d decay heat	state	0.0650	—	z_v its rate	state	0	±3

6.3.1 Saturated Properties

Saturation properties are quadratics in pressure fitted through IAPWS values at 5, 7 and 9 MPa, as shown in Equation (6.1).

$$\begin{aligned} \rho_g &= 0.13825\,p^2 + 3.9285\,p + 2.25225 \\ \rho_f &= -0.04875\,p^2 - 18.24\,p + 869.78875 \\ T_{\text{sat}} &= -0.54625\,p^2 + 17.5\,p + 190.09625 \\ h_f &= -2.1125\,p^2 + 81.9\,p + 797.9125 \\ h_{fg} &= 1.075\,p^2 - 80.15\,p + 2013.575 \end{aligned} \tag{6.1}$$

where p is in MPa, densities are in kg/m^3, temperature in °C and enthalpies in kJ/kg. The fits reproduce 6 and 8 MPa to better than 0.1 % on every property. The flashing coefficient $\mathrm{d}h_f/\mathrm{d}p = -4.225\,p + 81.9$ follows by differentiation and is used for the swell term. Feedwater enthalpy is $h_{\text{fw}} = 4.19\,T + 3.4 \times 10^{-4}\,T^2$ with T in °C, giving 920.9 kJ/kg at the design 216 °C. Property arguments are clamped to 4.5–9.5 MPa rather than extrapolated. At the design pressure of 7.170 MPa the fits give $\rho_g = 37.53\,\text{kg/m}^3$, $\rho_f = 736.50\,\text{kg/m}^3$ and $T_{\text{sat}} = 287.49\,°\text{C}$, so the steam is a twentieth as dense as the water it displaces and carries essentially none of the moderation.

6.3.2 Boiling Channel and Void Fraction

The channel closes an enthalpy balance on the heat the fuel is handing over, as shown in Equation (6.2).

$$h_e = h_{in} + \frac{Q_c}{W_c}, \qquad x_e = \frac{h_e - h_f}{h_{fg}} \tag{6.2}$$

where h_e is the exit enthalpy in kJ/kg, Q_c is the heat crossing into the coolant in kW, W_c is the core flow in kg/s and x_e is the exit quality. Voids appear before the bulk reaches saturation, so the non-boiling fraction of the channel is shortened from the thermodynamic value by the factor K_{nb}, as shown in Equation (6.3), which is the device that carries the subcooled-boiling behavior [9].

$$f_{nb} = \frac{K_{nb}\,\Delta h_{\text{sub}}}{\Delta h_{\text{sub}} + x_e h_{fg}}, \qquad \Delta h_{\text{sub}} = h_f - h_{in} \tag{6.3}$$

where Δh_{sub} is the core inlet subcooling in kJ/kg and f_{nb} is clamped to 0–0.95. The local void fraction at quality x follows the Zuber-Findlay drift-flux correlation [8, 16] and is averaged in closed form over a boiling length whose quality runs linearly from 0 to x_e, as shown in Equation (6.4).

$$\begin{aligned} \alpha(x) &= \frac{x}{C_0\left[x + \frac{\rho_g}{\rho_f}(1-x)\right] + \frac{\rho_g V_{gj}}{G}} \\ \bar{\alpha}_b &= \frac{1}{A} - \frac{B}{A^2 x_e}\ln\left(1 + \frac{A x_e}{B}\right) \\ A &= C_0\left(1 - \frac{\rho_g}{\rho_f}\right), \qquad B = C_0\frac{\rho_g}{\rho_f} + \frac{\rho_g V_{gj}}{G} \end{aligned} \tag{6.4}$$

where C_0 is the distribution parameter, V_{gj} the drift velocity in m/s and $G = W_c/A_{\text{flow}}$ the mass flux in kg/m^2/s. The core-average void fraction that the neutronics see is the boiling-length average diluted over the whole channel, as shown in Equation (6.5).

$$\alpha_{\text{ss}} = (1 - f_{nb})\,\bar{\alpha}_b, \qquad \tau_v \frac{\mathrm{d}\bar{\alpha}}{\mathrm{d}t} = \alpha_{\text{ss}} - \bar{\alpha} \tag{6.5}$$

where τ_v is the void relaxation time in seconds. At the design point (6.2) to (6.5) return $x_e = 13.742\,\%$, an exit void fraction of 65.23 % from (6.4), a non-boiling fraction of 10.57 % and a core-average void fraction of 40.32 %.

6.3.3 Point Kinetics and the Reactivity Balance

Reactivity is the sum of a rod term on an S-shaped integral worth curve, a Doppler term on the fuel temperature and a void term, all referenced to the design point so that the design state closes exactly, as shown in Equation (6.6).

$$\begin{aligned} \rho &= \rho_{\text{rod}} + \rho_D + \rho_\alpha + \rho_{\text{ext}} \\ \rho_{\text{rod}} &= W_{\text{rod}} \left[S(x_{\text{rod},0}) - S(x_{\text{rod}}) \right] \\ S(u) &= \frac{u}{100} - \frac{1}{2\pi} \sin\left(\frac{2\pi u}{100} \right) \\ \rho_D &= \alpha_D \left(T_f - T_{f,0} \right), \qquad \rho_\alpha = \alpha_v \left(\bar{\alpha}_{\text{eff}} - \bar{\alpha}_0 \right) \end{aligned} \tag{6.6}$$

where W_{rod} is the total bank worth in $\Delta k/k$, S is the normalized integral worth, α_D is the Doppler coefficient in $\Delta k/k$ per °C, α_v is the void coefficient in $\Delta k/k$ per unit void, $\bar{\alpha}_{\text{eff}} = \bar{\alpha} + z_o$ adds the density-wave oscillation and ρ_{ext} is an external reactivity disturbance. Equation (6.6) is the reactivity ledger of Equation (1.3) with the moderator term carried by the voids. Differentiating S gives a differential rod worth of $W_{\text{rod}} \left[1 - \cos(2\pi u/100) \right] /100$, which at the design insertion of 35 % is 47.6 pcm per percent of bank travel, or 8.51 ¢ per percent. Two delayed-neutron groups carry the kinetics, as shown in Equation (6.7). It keeps the prompt term that the prompt jump approximation of Equation (1.1) drops, and it carries two precursor groups in place of one.

$$\begin{aligned} \Lambda \frac{\mathrm{d}n}{\mathrm{d}t} &= (\rho - \beta_1 - \beta_2)\, n + \Lambda (\lambda_1 c_1 + \lambda_2 c_2) \\ \Lambda \frac{\mathrm{d}c_1}{\mathrm{d}t} &= \beta_1 n - \Lambda \lambda_1 c_1, \qquad \Lambda \frac{\mathrm{d}c_2}{\mathrm{d}t} = \beta_2 n - \Lambda \lambda_2 c_2 \end{aligned} \tag{6.7}$$

where n is the neutron density normalized to unity at rated power, c_i are the precursor concentrations, Λ is the prompt neutron generation time in seconds, β_i are the delayed fractions and λ_i their decay constants in s. One dollar is $\beta_1 + \beta_2 = 0.0056$. Heat passes through a single average fuel node and a one-group decay-heat state, as shown in Equation (6.8).

$$\begin{aligned} Q_{\text{th}} &= Q_0 \left[n \left(1 - f_{dk} \right) + Q_d \right], \qquad Q_c = \frac{T_f - T_{\text{sat}}}{R_f} \\ \frac{\tau_f}{R_f} \frac{\mathrm{d}T_f}{\mathrm{d}t} &= Q_{\text{th}} - Q_c, \qquad \tau_{dk} \frac{\mathrm{d}Q_d}{\mathrm{d}t} = f_{dk}\, n - Q_d \end{aligned} \tag{6.8}$$

where Q_0 is the rated thermal power in kW, f_{dk} is the decay-heat fraction, $R_f = \Delta T_{f,0}/Q_0$ is the fuel-to-coolant thermal resistance in °C/kW, τ_f is the fuel time constant and τ_{dk} the short-term decay-heat time constant. Thermal power is the reported controlled variable and neutron flux is reported separately, because the flux answers a reactivity insertion in milliseconds while the thermal power follows it through $\tau_f = 6.0\,\text{s}$.

6.3.4 Recirculation Flow

Core flow is driven by the jet-pump head plus the buoyancy of the voids themselves against a quadratic loop resistance, as shown in Equation (6.9).

$$\begin{aligned} \Delta p_{\text{nat}} &= \bar{\alpha} \left(\rho_f - \rho_g \right) g L_{th} \\ W_{c,\text{ss}} &= \sqrt{\frac{f_p^2\, \Delta p_{d,0} + \Delta p_{\text{nat}}}{K_{\text{loop}}}} \\ \tau_p \frac{\mathrm{d}f_p}{\mathrm{d}t} &= f_{p,\text{cmd}} - f_p, \qquad \tau_W \frac{\mathrm{d}W_c}{\mathrm{d}t} = W_{c,\text{ss}} - W_c \end{aligned} \tag{6.9}$$

where f_p is the pump drive fraction, $\Delta p_{d,0}$ is the jet-pump head at design drive, K_{loop} is the loop resistance in Pa/(kg/s)^2, L_{th} is the thermal driving height in m, τ_p is the drive time constant and τ_W the flow inertia time constant. The two constants are not free. The natural-circulation flow fraction f_{nat} that the plant is quoted at fixes the total loop resistance and the jet pumps supply the remainder, as shown in Equation (6.10).

$$K_{\text{loop}} = \frac{\Delta p_{\text{nat},0}}{f_{\text{nat}}^2 W_{c,0}^2}, \qquad \Delta p_{d,0} = \frac{\Delta p_{\text{nat},0}\left(f_{\text{nat}}^{-2} - 1\right)}{f_{p,0}^2} \tag{6.10}$$

where $\Delta p_{\text{nat},0}$ is (6.9) evaluated at the design void fraction. Feedwater mixes into the recirculating saturated water in the downcomer, and the transport of that mixture to the core inlet is the dead time in every shrink transient, as shown in Equation (6.11).

$$\tau_{DC}\frac{\mathrm{d}h_{in}}{\mathrm{d}t} = h_f - \frac{W_{\text{fw}}}{W_c}\left(h_f - h_{\text{fw}}\right) - h_{in} \tag{6.11}$$

where τ_{DC} is the downcomer mixing transport time in seconds. At the design point (6.11) is stationary at $h_{in} = 1227.7\,\text{kJ/kg}$, which is 48.87 kJ/kg of subcooling and 9.97 °C below saturation.

6.3.5 Pressure, Inventory and the Two Levels

The vessel carries one energy balance and one mass balance, as shown in Equation (6.12).

$$\begin{aligned} C_e\frac{\mathrm{d}p}{\mathrm{d}t} &= Q_{\text{th}} + W_{\text{fw}}h_{\text{fw}} - W_{\text{out}}\left(h_f + h_{fg}\right) \\ \frac{\mathrm{d}M_v}{\mathrm{d}t} &= W_{\text{fw}} - W_{\text{out}} \\ W_{\text{out}} &= K_t\, y_T\, p\, c_{\text{turb}}\, m_{\text{MSIV}} + K_{\text{SRV}}\max(0,\, p - p_{\text{SRV}}) \end{aligned} \tag{6.12}$$

where C_e is the lumped energy storage of water, steam and metal in kJ/MPa, M_v is the vessel inventory in kg, K_t is the choked turbine nozzle coefficient, c_{turb} is the admission capacity multiplier, m_{MSIV} is 1 with the main steam isolation valves open and 0 once they have latched shut, and K_{SRV} is the relief valve capacity in kg/s/MPa. The nozzle coefficient is calibrated so that the design valve position passes the design steam flow, $K_t = W_{s,0}/(y_{T,0}\,p_0) = 295.4\,\text{kg/s}$ per unit valve per MPa. Differentiating the first line of (6.12) at the design point gives an integrating process gain of –4.19 kPa/s per percent of valve travel, so pressure has no self-regulating term of its own in the valve path.

Pause. Try Exercise 6.3 before reading on.

The level follows from splitting the inventory by volume. Sub-surface steam has three parts, the void in the channels and upper plenum, the steam the separators are holding in transit, and the steam that a falling pressure is flashing out of the liquid, as shown in Equation (6.13).

$$\begin{aligned} V_{sd} &= \bar{\alpha}_{\text{eff}}V_{ch} + \frac{\tau_{\text{sep}}W_{sg}}{\rho_g} + \frac{\tau_{\text{sep}}W_{\text{flash}}}{\rho_g} \\ W_{\text{flash}} &= -\frac{V_l\rho_f}{h_{fg}}\frac{\mathrm{d}h_f}{\mathrm{d}p}\frac{\mathrm{d}p}{\mathrm{d}t}, \qquad W_{sg} = W_c\,x_e \end{aligned} \tag{6.13}$$

where V_{ch} is the two-phase volume above the boiling boundary in m^3, τ_{sep} is the separator and standpipe holdup time in seconds, W_{sg} is the boiling rate and W_{flash} the flashing rate, both in kg/s. The inventory then divides exactly between liquid, sub-surface steam and dome steam, as shown in Equation (6.14).

$$\begin{aligned} V_{wr} &= \frac{\left(M_v - \rho_g V_{sd} - \rho_g V_{\text{tot}}\right)/\rho_f + V_{sd}}{1 - \rho_g/\rho_f} \\ L_{\text{sw}} &= \frac{V_{wr} - V_{wr,0}}{A_s}, \qquad L_{\text{col}} = \frac{V_{wr} - V_{sd} - V_{l,0}}{A_s} \end{aligned} \tag{6.14}$$

where V_{wr} is the water-region volume in m^3, V_{tot} is the vessel free volume, A_s is the free surface area in m^2, L_{sw} is the swollen surface height and L_{col} the collapsed level, both in meters from normal water level. The two differ by V_{sd}/A_s, which at the design point is $80.71\ \mathrm{m}^3$ over $32\ \mathrm{m}^2$, or 2522 mm of apparent level that is not water.

The transmitter does not report L_{sw} either. Its measured pressure difference is converted to a height using the two densities present when the cell was ranged, not the two densities present in the vessel, as shown in Equation (6.15).

$$L_{\mathrm{ind}} = \frac{\left(\frac{H_{\mathrm{tap}}}{2} + L_{\mathrm{sw}}\right)\left(\rho_f - \rho_g\right) + H_{\mathrm{tap}}\left(\rho_g - \rho_{g,0}\right)}{\rho_{f,0} - \rho_{g,0}} - \frac{H_{\mathrm{tap}}}{2} \tag{6.15}$$

where H_{tap} is the separation between the taps in m and a zero subscript marks a density at the calibration pressure. Equation (6.15) returns L_{sw} exactly when $p = p_0$ and departs from it as the vessel moves off the ranging pressure, so the same physical surface reads differently at a different dome pressure.

6.3.6 Density-Wave Stability

The decay ratio of the density-wave mode is a fitted function of the power-to-flow ratio and, independently, of absolute flow, which is what places the exclusion region in the upper-left corner of the operating map [5], as shown in Equation (6.16).

$$\mathrm{DR} = \mathrm{DR}_0 \exp\left[a_{\mathrm{DR}}\left(\frac{P}{F} - 1\right) + b_{\mathrm{DR}}\left(1 - F\right)\right] \tag{6.16}$$

where P is thermal power and F core flow, both as fractions of rated, and P is taken from the fuel node rather than from the flux so that the decay ratio does not itself oscillate. The mode is a driven second-order oscillator on the void fraction whose damping comes from the decay ratio, as shown in Equation (6.17).

$$\begin{aligned} \frac{\mathrm{d}z_o}{\mathrm{d}t} &= z_v, \qquad \frac{\mathrm{d}z_v}{\mathrm{d}t} = -\omega^2 z_o - 2\zeta_{\mathrm{eff}}\,\omega\, z_v + d_o \\ \zeta &= \frac{-\ln \mathrm{DR}}{\sqrt{4\pi^2 + \ln^2 \mathrm{DR}}}, \qquad \zeta_{\mathrm{eff}} = \zeta + k_{\mathrm{NL}} z_o^2 \end{aligned} \tag{6.17}$$

where $\omega = 2\pi/T_{\mathrm{osc}}$, d_o is the boiling-noise drive and k_{NL} is the amplitude-dependent term that bounds the limit cycle. The oscillation enters the reactivity balance through $\bar{\alpha}_{\mathrm{eff}}$ in (6.6), so a decay ratio above unity appears on the neutron flux, which is where a stability monitor reads it [13].

6.3.7 Parameters, Initialization and Numerics

All sixty-two constants are listed in Table 6.2. Three are hidden and never displayed, and the remaining fifty-nine are adjustable from the model panel. The three control bands and the exclusion-region boundary appear in no derivative and are read only by the scorecard accumulators, which total the seconds each controlled variable spends outside its band.

Each scenario begins from an equilibrium the application solves for. Three residuals in three unknowns close the algebra, because the void fraction depends on the flow through (6.4), the flow depends on the void fraction through the buoyancy term in (6.9), and the power depends on both through (6.6). The residuals are total reactivity, the pressure or valve constraint, and the flow balance, and they are driven to zero by damped Newton iteration on a numerical Jacobian. The precursor states are then initialized to $c_i = \beta_i n/(\lambda_i \Lambda)$ and the inventory to the value that puts the transmitter of (6.15) on the operator's level setpoint. The equilibrium test in the application evaluates all seventeen derivatives at that state and requires the largest magnitude, scaled by the state, to be below 1×10^{-7}. The measured value at the design point is 1.4×10^{-14}, on the neutron density.

Table 6.2 Model Parameters. The Three Marked h Are Hidden and Never Displayed, and the Rest Are Adjustable.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
Λ	prompt generation time	4.5×10^{-5}	s	p_0	design dome pressure	7.17	MPa
β_1	fast delayed fraction	0.002 24	—	C_e	energy storage	1.40×10^7	kJ/MPa
λ_1	fast decay constant	0.40	s	V_{tot}	vessel free volume	480	m^3
β_2	slow delayed fraction	0.003 36	—	$V_{wr,0}$	water-region volume	380	m^3
λ_2	slow decay constant	0.0435	s	A_s	free surface area	32	m^2
α_v	void coefficient	−0.120	$\Delta k/k$	τ_{sep}	separator holdup	0.8	s
α_D	Doppler coefficient	-2.3×10^{-5}	1/°C	H_{tap}	tap separation	1.6	m
W_{rod}	total bank worth	0.030	$\Delta k/k$	$y_{T,0}$	design valve position	0.85	—
$x_{\text{rod},0}$	design insertion	35	%	τ_{TCV}	valve stroke time	0.25	s
—	rod drive rate	0.55	%/s	k^h_{gen}	generator output	0.6111	MW/(kg/s)
—^{h}	scram insertion rate	35	%/s	$T_{\text{fw},D}$	design feedwater temp.	216	°C
Q_0	rated thermal power	3330	MW	τ_{fw}	feedwater flow constant	3.5	s
$W_{c,0}$	rated core flow	13100	kg/s	$\tau_{\text{fw}T}$	feedwater temp. constant	30	s
A_{flow}	core flow area	8.7	m^2	DR_0	decay ratio at rated	0.25	—
V_{ch}	two-phase volume	105	m^3	a_{DR}	power/flow sensitivity	1.20	—
L_{th}	thermal driving height	5.0	m	b_{DR}	low-flow sensitivity	0.77	—
C_0	distribution parameter	1.13	—	T_{osc}	oscillation period	2.2	s
V_{gj}	drift velocity	0.23	m/s	k^h_{NL}	limit-cycle saturation	400	—
K_{nb}	subcooled-boiling factor	0.55	—	L_{hi}	Level 8 trip	+0.560	m
$\Delta T_{f,0}$	fuel-to-coolant rise	412	°C	L_{lo}	Level 3 scram	−0.450	m
τ_f	fuel time constant	6.0	s	p_{MSIV}	low-pressure isolation	5.60	MPa
τ_v	void relaxation	1.6	s	p_{SRV}	relief valve setpoint	7.86	MPa
τ_{DC}	downcomer transport	3.0	s	K_{SRV}	relief valve capacity	3000	kg/s/MPa
f_{dk}	decay-heat fraction	0.065	—	p_{scram}	high-pressure scram	7.72	MPa
τ_{dk}	decay-heat constant	80	s	n_{scram}	APRM high flux scram	1.20	—
$f_{p,0}$	design pump drive	0.95	—	τ_{Pm}	pressure transmitter	0.15	s
f_{nat}	natural-circulation flow	0.33	—	τ_{Lm}	level transmitter	1.2	s
τ_p	pump drive constant	5.0	s	τ_{Nm}	APRM filter	0.6	s
τ_W	flow inertia constant	2.5	s	τ_{Wm}	flow transmitter	0.8	s
DR_{lim}	exclusion-region boundary	0.80	—	—	level control band	0.100	m
—	pressure control band	0.10	MPa	—	power control band	2.0	%

Integration is classical fourth-order Runge-Kutta on a fixed 0.005 s step. Controller execution is every 0.05 s, a 20 Hz scan chosen to match a real electrohydraulic pressure regulator, and ten samples of every simulated second reach the historian. The noise sequences are drawn from a mulberry32 generator seeded at 20 260 906, which makes a noisy run repeat bit for bit. The engine, the seed and the speed selector, which changes no trajectory, are those of Section 1.6.1.

6.3.8 Assumptions

Detail is spent where a controller has to be designed around the physics and withheld everywhere else. Each simplification fails somewhere.

One point-kinetics node with two delayed groups. A real core has a strongly bottom-peaked axial power shape and can oscillate with its two halves out of phase, which is the regional mode that made the LaSalle event difficult to diagnose [12]. A point model cannot represent a regional mode at all, so only the core-wide in-phase oscillation appears here. The two groups reproduce the prompt jump and the minute-scale precursor tail but not the intermediate decay constants.

One average channel with uniform axial power. Several hundred bundles at different powers are represented by one channel, so every thermal limit that an operator actually manages, the critical power ratio and the linear heat generation rate, is written against a hot channel the model does not have. The application therefore says nothing about margin to dryout.

Constant void and Doppler coefficients. Both vary with exposure, void history and rod pattern, and the void coefficient becomes markedly more negative late in a cycle. There is no xenon, no burnup and no rod pattern, only one ganged bank on one S-curve, with the external reactivity slider standing in.

Drift flux integrated in closed form. Equation (6.4) is sound for rod bundles at 4–9 MPa above roughly 300 kg/m^2/s [8, 16]. Below that the drift term dominates and the correlation degrades, which is precisely the low-flow corner where the instability region sits, so the void fractions reported there are indicative rather than predictive.

Density-wave instability is a calibrated oscillator. Equations (6.16) and (6.17) set the damping of a second-order mode from a fitted decay ratio and drive it with boiling noise. The onset boundary and the 0.45 Hz period are calibrated to published boiling water reactor behavior [5], but the channel momentum equations that produce the phenomenon are not solved, so the limit-cycle amplitude is plausible rather than predicted.

One exponential group of decay heat. Equation (6.8) carries 6.5 % of rated power with an 80 s time constant, which is fair for the first two or three minutes after a scram and increasingly wrong after that.

The turbine is a choked nozzle and the generator is a constant. There is no shaft, no inertia, no frequency, no overspeed protection and no turbine bypass valves, so a turbine trip here is a pressurization rather than the mechanical event it would be, and the bypass system a real plant uses to absorb one is absent. Chapter 5 treats the steam dump that takes that role on a pressurized water reactor.

One control volume for pressure and one for mass. Carry-under, downcomer subcooling gradients and the separator hydraulics are represented only by the holdup time τ_{sep}, so regional level effects and the difference between narrow-range and wide-range indication are outside the model.

A feedwater pump trip halves the capacity at once. A real pump set coasts down. The instant halving is a stylized element that isolates the swell transient from the coastdown, which would otherwise blur it.

No engineered safety systems. Once the plant has scrammed and isolated, the application integrates decay heat into a closed vessel with relief valves. Emergency core cooling, reactor core isolation cooling and standby liquid control are absent, and nothing about the sequences that follow a scram is represented.

6.4 Coupling, Sign Reversal and the Operating Map

The four couplings that decide the control structure are all measurable on the running plant, and the application is arranged so that each can be isolated. The Pressure-only architecture leaves the pressure regulator in service and every other station on the operator, which is the configuration a plant test would use.

6.4.1 Recirculation Flow as the Power Handle

Slowing the recirculation pumps from 95 % to 80 % takes the core flow from 100.0 % to 86.18 %, raises the core-average void fraction from 40.32 % to 41.04 %, which the reported reactivity balance shows as −15.4 ¢, and settles the power at 90.90 %. The gain is 0.659 % of rated power per percent of core flow and 63 % of the change arrives in 3.3 s. The rod bank position at the end of the run differs from its starting value by less than one part in 10^9, so none of it came from a rod. That is the normal load-following actuator on the plant [6].

The gain is not constant. Evaluated as a local derivative along the 35 % rod line it runs from 0.632 % of power per percent of flow at rated conditions to 0.820 % per percent at 52 % flow, because the same fractional flow change removes a larger absolute void inventory when the channels are already voided. The gain from pump speed to core flow moves the other way, from 0.932 % per percent at 95 % drive to 0.790 % per percent at 46.7 % drive, because jet-pump head goes as the

square of the drive. The two effects very nearly cancel: the measured gain from pump speed to power is 0.589 % per percent at 95 % drive and 0.627 % per percent at 46.7 % drive, a change of 6 % across half the pump range. That is not what the jet-pump-squared argument on its own predicts, and Section 6.6.3 reports what it costs the gain-scheduling argument.

6.4.2 What Rods Actually Buy

Driving the bank in 7 % from the design insertion inserts 366 pcm, which is 65.4 ¢. Power falls only 11.46 % once settled, or 1.64 % of rated per percent of insertion, because the feedbacks return almost all of it. The core-average void fraction falls from 40.32 % to 38.17 %, worth +46.0 ¢, the fuel cools by 47.2 °C, worth +19.4 ¢, and total reactivity settles within 0.02 ¢ of zero. The rods did not move the operating point so much as move the rod line the operating point sits on. Timing sharpens the comparison. The bank travels at 0.55 %/s, so 7 % of insertion is 11.5 s to 90 % of the move against the 3.3 s in which the recirculation pumps delivered 63 % of a comparable power change.

Sweeping an external reactivity bias with pressure and level held and the pumps and rods fixed measures the self-regulation directly. Biases of +10 pcm, −10 pcm, +20 pcm and −20 pcm settle at 30.59, 30.66, 30.55 and 30.69 pcm per percent of power, and the symmetric difference at either amplitude is 30.62 pcm per percent, so a controller acting on this plant works in parallel with a proportional regulator of about 30.6 pcm per percent that it did not install and cannot switch off. The pressurized water reactor of Chapter 2 has the same kind of regulator in its two negative temperature coefficients, and there it answers a change in steam demand with no rod moving.

6.4.3 Pressure as a Reactivity Signal

Raising the dome pressure makes the steam denser, collapses the voids and inserts reactivity. Holding power and flow and evaluating (6.1) to (6.5) at a higher pressure gives +21.8 pcm at +0.05 MPa, +43.4 pcm at +0.10 MPa and +86.3 pcm at +0.20 MPa. That is why the turbine control valves are assigned to pressure rather than to power and why the pressure loop is tuned tighter than anything else on the plant.

Taking the pressure regulator out of service with the valves parked four percent further open depressurizes the dome at up to 17.7 kPa/s and takes the power from 100 % to 88.8 % by the time the dome reaches 6.30 MPa, an average of 1.28 % of rated per 0.1 MPa. The excursion does not stop there. Falling pressure lowers the power, lower power makes less steam, and less steam lowers the pressure further, and in this scenario nothing arrests it. The run reaches 5.60 MPa at 243.44 s, the main steam isolation valves latch shut at the next controller scan (243.45 s), and the pressurization that follows collapses voids fast enough to scram the plant on low vessel level in the very next 5 ms integration step. The high pressure trip arrives 7.7 s after that and the dome peaks at 8.366 MPa on the relief valves.

Whether the excursion runs at all turns out to depend on the feedwater rather than on the valve. Scenario 9 holds level on three-element control, so the feedwater tracks the falling steam flow and the core inlet subcooling falls with it, from 43.9 kJ/kg at 50 s after the park to 38.9 kJ/kg at 100 s. Repeating the same parked-valve experiment in MANUAL, where the feedwater is pinned at 100 % of rated, the subcooling instead holds at 44.3 kJ/kg as the steam flow falls away from it, and that reactivity arrests the slide. Parked at 89 % the MANUAL dome settles at 6.552 MPa with the power at 93.9 %, while the same park with the level loop in service is at 6.317 MPa 100 s after the park, still falling at about 7 kPa/s, and diverges from there. Across parked positions from 82 % to 90 % the MANUAL dome reaches a stationary value, 7.694 MPa at 82 %, 7.338 MPa at 84 %, 7.007 MPa at 86 %, 6.774 MPa at 87.5 % and 6.410 MPa at 90 %, each stationary to 0.002 MPa or better over the 50 s preceding its level trip, and a ± 0.15 MPa perturbation at the design position decays back to 7.170 MPa with a fitted time constant of 28 s.

What is not self-regulating in MANUAL is the inventory. A fixed feedwater flow against a changed steam flow is a standing mass imbalance, so every off-design park ends on a level trip and not on a pressure trip: Level 8 at 331 s from an 87.5 % park and Level 3 at 569 s from an 84 % park. The

pressure loop on this plant is held up in part by the level loop, and taking the level loop out changes which of the two diverges.

The reverse direction is faster and is the clearest single demonstration in the application of what the level measurement is worth. Parking the valves at 62 % from full power pressurizes the dome, collapses voids, and drops the indicated level −400 mm within 0.5 s. The Level 3 low-level scram fires 0.835 s after the valve command changes, at which instant the dome has risen only 0.057 MPa, the collapsed level has risen +24 mm and the vessel has gained 276 kg. A reactor is tripped on low water level, in under a second, while the vessel is filling.

The same collapse limits how fast the pressure setpoint may rise. Scenario 12 steps it from 7.17 to 7.22 MPa at 420 s, as shown in Figure 6.3. Under Coordinated PIC takes the step at once and the indication falls to −114.9 mm at 420.95 s, 1.7 s outside the band, with no water lost, while the pressure enters ±0.005 MPa of the setpoint at 423.7 s, overshoots 0.0086 MPa and stays inside from 440.70 s. The MPC of Section 6.5.6 predicts the collapse, walks p^R up (half the step by 427 s, 90 % by 442 s) and trims the feedwater, so the level dips only −17.7 mm, at 425.4 s, and the pressure settles without overshoot at 440.02 s, about when Coordinated does. Its cost is the integrated pressure error, 0.551 against 0.224 MPa s from 420 to 600 s.

6.4.4 Feedwater as a Reactivity Actuator

Losing a feedwater heater string cools the feedwater by 28 °C with nothing on the reactor touched. Core inlet subcooling rises from 48.87 kJ/kg to 67.37 kJ/kg, the boiling boundary moves up the channel, the core-average void fraction falls from 40.32 % to 39.57 %, which is +16.0 ¢, and the power rises to 109.45 %. Half the rise arrives in 17.5 s, 99 % of it in 152 s, and the power first reaches its final value 237 s after the loss, limited by the 30 s feedwater temperature lag and by the fuel node. At the new steady state the Doppler term reads −16.0 ¢ against the void term's +16.0 ¢, the fuel having warmed 38.9 °C, and the rod position is unchanged to five decimal places. A failure two hundred meters from the reactor building raises core power by a tenth, delivered entirely through the condensate system.

6.4.5 Shrink and Swell, Measured

Stepping the feedwater +10 % with the pressure regulator in service and the level loop out admits an extra 180.0 kg/s to a vessel holding 227.2 t. The mass balance alone predicts the level rising at 7.64 mm/s. The indicated level instead falls, reaching −40.9 mm at 3.45 s and not returning through its starting value until 6.10 s. At the instant of the minimum the collapsed level of (6.14) is +12.1 mm and the inventory is 193 kg heavier than it was, and the collapsed level rose monotonically throughout. Nothing about the inventory went the wrong way. Only the measurement did, because the colder inlet enthalpy of (6.11) shortened the boiling length of (6.3) and condensed steam that had been holding the surface up. Section 1.4.4 writes the same split for a drum as a water term and a void term, Equation (1.7). Power rises at the same time and peaks 4.47 % higher 5.5 s after the step, so even the feedwater valve is a reactivity actuator on this plant.

> **TRY IT IN THE APP** **2 · Shrink: The Level Goes the Wrong Way**
>
> Load scenario 2 and press Run on the Vessel view. At 30 s the feedwater steps to 110 %. Watch the solid indicated level against the dashed collapsed level: the indication falls to −40.9 mm at 3.45 s after the step while the collapsed level rises the whole time, and the indication does not return through its starting value until 6.10 s.

The opposite disturbance is larger and more dangerous. Halving the feedwater pump capacity drops the flow to 1170 kg/s against 1800 kg/s of steam, a deficit of 630 kg/s. The indicated level rises to +212.9 mm at 3.15 s, at which instant the collapsed level is −58.8 mm: the two readings are 271.7 mm apart and pointing in opposite directions. Less feedwater is less subcooling, a longer boiling length and more void, and the extra steam lifts the surface faster than the deficit lowers it.

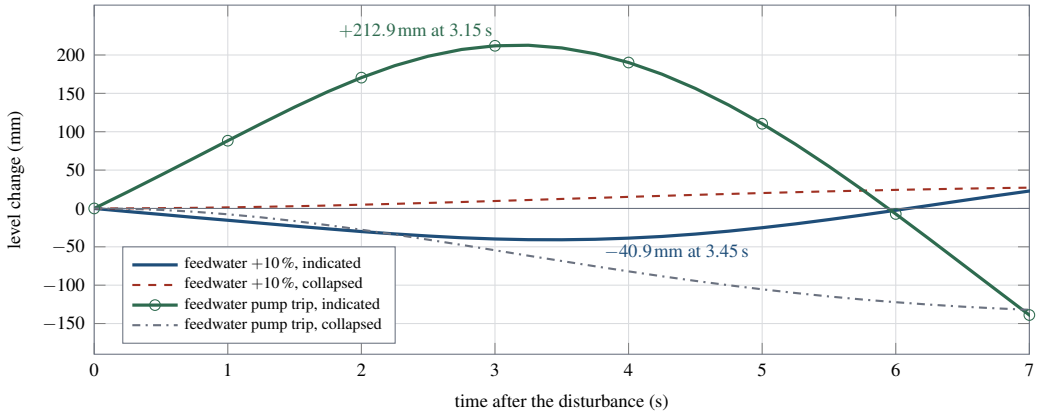

Figure 6.2 Indicated and collapsed level through both sign reversals, recorded from the running application and differenced against the pre-disturbance value. A +10 % feedwater step drives the indication −40.9 mm while the collapsed level rises monotonically, and a feedwater pump trip drives it +212.9 mm while the collapsed level is already 58.8 mm down. Both indications cross back through their starting value within 6.1 s, and in neither case does the inventory move in the direction the indication first takes. The feedwater step is drawn solid (indicated) and dashed (collapsed), the pump trip solid with circles (indicated) and dash-dotted (collapsed).

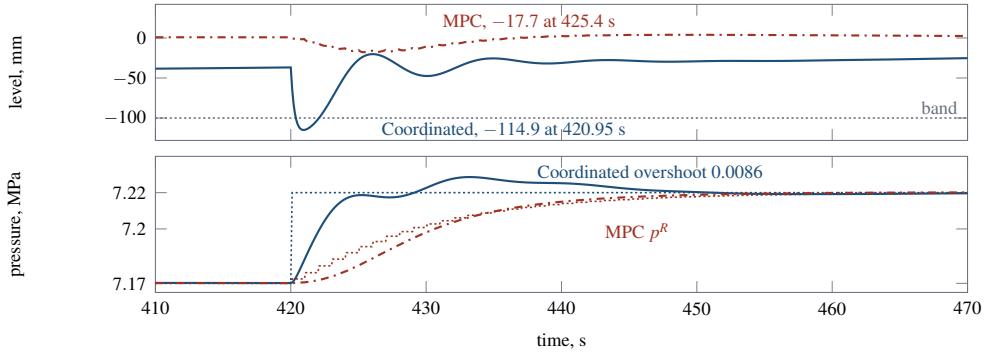

Figure 6.3 Scenario 12's pressure step under Coordinated (blue, solid) and the MPC (red, dash-dotted): level error, 36.8 mm low under Coordinated before the step, and pressure with the PIC setpoint each gives (dotted). The gray dotted line in the upper panel is the edge of the level band.

Then the mass wins. The indication crosses back through its starting value at 5.9 s, reaches −290 mm at 20 s, and the Level 3 scram fires 33.05 s after the trip.

Both disturbances are applied from an exactly stationary plant, so the pre-disturbance value serves as the held companion run and every trace in Figure 6.2 is one variable away from it.

Pause. Try Exercise 6.2 before reading on.

6.4.6 The Operating Map

Power and core flow are not independent, so the plant is licensed against a map rather than against a list of limits. Along the 35 % rod line the decay ratio of (6.16) rises from 0.250 at rated conditions to 0.283 at 90 % flow, 0.371 at 70 % and 0.495 at 52 % flow, at which point the power has fallen to 65.6 %. Stopping both recirculation pumps at the design rod line settles the plant at 50.4 % power on 34.6 % flow with a decay ratio of 0.717. That is just below the exclusion-region boundary $\mathrm{DR}_{\mathrm{lim}}$

of Table 6.2, in the configuration the LaSalle event was in [12]. Sweeping the rod bank from 70 % to 25 % inserted with both pumps stopped traces the natural-circulation line from 20.7 % power at 28.9 % flow to 59.5 % power at 35.5 % flow, so natural circulation alone supports well over half of rated power on about a third of rated flow.

> **DID YOU KNOW? If the Scram Had Failed**
>
> A laboratory simulation of the 1988 oscillation asked what would have happened had the reactor not scrammed. Doppler and void feedback limited the oscillations, but the power peaked at 1,300 percent of full power while its time average rose only to 86 percent [15].

6.5 Controller Design

6.5.1 Loop Pairing

Four loops close on this plant and the pairing follows from the physics rather than from a decoupling calculation. Dome pressure is paired with the turbine control valves because the valve path has the fastest and largest effect on pressure and because pressure is a reactivity term that must be held tightly. Vessel level is paired with feedwater because feedwater is the only inventory handle. Core power is paired with recirculation flow because that path is a factor of three faster than the rod path and costs no fuel duty. The rod bank is paired with the recirculation pump demand rather than with any plant variable, as a valve position controller that walks the pumps back toward the middle of their range so that the fast handle keeps authority in both directions. The couplings that remain are strong, which is why the comparison in Section 6.6 is between structures and not between tunings [4]. The reverse osmosis train of Volume 6, Chapter 2 is a plant on which the same pairing decision is made before any loop is tuned.

Pause. Try Exercise 6.1 before reading on.

6.5.2 Control Algorithm

All five controllers run the velocity-form PID of Section 1.6.2, executed every $\Delta t_c = 0.05\,\mathrm{s}$ and written here in the loops' own symbols. Each scan forms increments, as shown in Equation (6.18).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k) \\
\Delta u_k^{P} &= K_{c,k}\,e_k - K_{c,k-1}\,e_{k-1}, \qquad \Delta u_k^{I} = \frac{K_{c,k}}{\tau_I}\,e_k\,\Delta t_c
\end{aligned}
\tag{6.18}
$$

The output and the integral then follow the recursion of Equation (6.19).

$$
\begin{aligned}
s_k &= s_{k-1} + \Delta u_k^{P}, \qquad c_k = u_{\mathrm{bias}} + s_k + I_{k-1} \\
u_k &= \min(u_{\max},\ \max(u_{\min},\ c_k)) \\
I_k &= I_{k-1} + \Delta u_k^{I} + \frac{\Delta t_c}{T_t}\,(u_k - c_k)
\end{aligned}
\tag{6.19}
$$

where $a = +1$ is reverse and $a = -1$ direct action, $K_{c,k}$ is the proportional gain in force at scan k, τ_I the integral time, T_t the tracking time constant, $u_{\mathrm{bias}} = 0$, and the integral is seeded with the actuator position the plant starts from. The running sum s_k is the S_k of Equation (1.14), and the demanded output c_k is the v_k of Equation (1.15). Because each increment is the difference of two whole proportional terms, s_k is always exactly $K_{c,k}e_k$, so a gain that the scheduling of Section 6.5.4 moves acts on the whole proportional term at once. Anti-windup is the back-calculation in the last line, and s_k is never truncated. Derivative action is not exposed on any of the five loops, so $\tau_D = 0$. Manual and automatic transfer is bumpless both ways, as Section 1.6.2 describes. The one exception is the gain scheduling of Section 6.5.4, which acts only in AUTO, so a switch to AUTO first re-seeds the loops on the gains AUTO will use, as an architecture switch does.

Table 6.3 The Five Implemented Loops as Written, With the Default Settings Every Scenario Loads

Loop	PV → MV	a	K_c	τ_I (s)	T_t (s)
PIC	$p \to y_T$	−1	80 %/MPa	6	2
LIC	$L_{\text{ind}} \to W_{\text{fw}}$	+1	30 %/m	90	40
FIC	$W_{\text{fw}} \to$ valve	+1	1.5 %/%	3	5
RRC	$Q_{\text{th}} \to f_p$	+1	1.2 %/%	15	30
RWM	$f_{p,\text{cmd}} \to x_{\text{rod}}$	+1	0.05 %/%	400	240

Table 6.4 The Six Control Architectures and What Each One Closes

Architecture	Loops in service
Manual	none. Every station on the operator
Pressure only	PIC. The configuration a plant test uses
Single-element FW	PIC, and LIC writing the feedwater demand directly
Three-element FW	PIC, FIC on feedwater flow with steam flow fed forward, and LIC trimming the flow setpoint
Coordinated	the above plus RRC on power, RWM behind the pumps, and gain scheduling on power and pump speed
MPC	one plan each second for the pump and feedwater demands and the PIC setpoint, RWM behind the pumps, Coordinated as fallback

The five loops and their default settings are listed in Table 6.3. The pressure loop is direct acting in the ISA sense, action −1 in the convention of Section 1.3, because opening the valves lowers the pressure. The level, feedwater flow, power and rod loops are reverse acting, action +1. The rod controller reads the pump *demand* rather than the pump speed and is deliberately the slowest controller on the plant.

Output limits are part of the design. PIC spans 0–100 % of valve travel. FIC and LIC in single-element mode span 0–130 % of rated steam flow, further limited to the available feedwater pump capacity and forced to zero once the Level 8 feedwater trip has latched. In three-element mode LIC is instead clamped to ±75 %, because its output is then a trim on the feedforward rather than the demand. RRC spans 28 % to 100 % of pump drive, the lower bound being the speed below which the jet pumps stop controlling the flow, and RWM spans the full rod bank. Two limits act after a controller output, where its integral cannot see them. In three-element mode the sum of the feedforward and the level trim is clipped to 0–130 % and to the pump capacity outside LIC, which knows only its own ±75 %. Rate limiting lives in the plant rather than in the controllers: the rod bank of (6.6) is limited to 0.55 %/s and to 35 %/s on a scram, and the pump drive, feedwater flow and valve position reach their commands through the lags τ_p, τ_{fw} and τ_{TCV} of Table 6.2.

6.5.3 Six Architectures on One Actuator Set

The application selects among six structures, listed in Table 6.4, all of them writing to the same four actuators through the same limits at the same scan.

Three-element control writes the feedwater flow setpoint as shown in Equation (6.20).

$$W_{\text{fw,SP}} = K_{\text{ff}}\,\widehat{W}_{st} + u_{\text{LIC}}, \qquad u_{\text{FIC}} = \text{PID}\left(W_{\text{fw,SP}},\ \widehat{W}_{\text{fw}}\right) \tag{6.20}$$

where $\widehat{W}_{st}$ and $\widehat{W}_{\text{fw}}$ are the filtered steam and feedwater flows as percentages of the rated value, K_{ff} is the feedforward gain and u_{LIC} is the level controller's trim. At $K_{\text{ff}} = 1$ every kilogram leaving the vessel is replaced as it leaves. At $K_{\text{ff}} = 0$ the structure reverts exactly to single element, and any value in between leaves the level controller supplying the balance out of an error signal it measures backwards.

Because the level controller's output means the demand in one architecture and a trim in the other, and the scheduled gains change with the architecture, switching would otherwise bump the actuators. At every change of architecture, feedforward gain, scheduling or rod controller, the application puts the level, flow, power and rod controllers on the gains their next scan will use and makes each track the command the plant already has (less the feedforward for the level trim). Tracking sets s_k to the proportional term of the present error and I_k to the balance, so the next scan adds only the change in error since the switch. From a lined-out full-power state, no transition among the five loop architectures moves any actuator command by more than 0.023 % on its first scan, the same order as the loops' own motion between scans. During a transient the switch is as quiet. Switching from single to three element at 200 s of the single-element load change, at 85.4 % power with the indication 278 mm high, leaves the feedwater command at 86.83 % on the first scan, and three-element control then trims it to no less than 84.00 % over the next 10 s. A seed that cleared s_k instead would add the whole $K_c e_k$ on the first scan and drop the same command to 73.57 %.

6.5.4 Gain Scheduling

The coordinated architecture schedules two gains on measured states, as shown in Equation (6.21).

$$\begin{aligned} K_c^{\mathrm{LIC}} &= K_c \cdot \mathrm{sat}\left(\frac{0.72}{n}, 0.6, 2.2\right) \\ K_c^{\mathrm{RRC}} &= K_c \cdot \mathrm{sat}\left(\frac{f_{p,0}}{f_p}, 0.7, 3.0\right) \end{aligned} \tag{6.21}$$

where $\mathrm{sat}(\cdot, l, h)$ clamps to the bracket, n is the neutron density and f_p the pump drive fraction. The level term detunes the loop as power rises, on the argument that there are more voids to condense and the right-half-plane zero moves toward the origin. The recirculation term raises the controller gain as the pumps slow. Section 6.6.3 measures what each contributes, and the second term does not survive the measurement.

6.5.5 Protection

Five trips and two isolations are implemented and each is armed in at least one scenario. The Level 3 low-level scram at −450 mm and the Level 8 feedwater and turbine trip at +560 mm both fire in the feedwater pump trip scenario. The high-pressure scram at 7.72 MPa and the APRM high flux scram at 120 % both fire in the turbine valve scenario, as do the main steam isolation at 5.60 MPa and the relief valves at 7.86 MPa. The oscillation scram counts the seconds the decay ratio spends at or above unity, drains the count at 0.4 s per second below it, and fires when the count passes 15 s, which happens in the instability scenario and is the first trip to act there.

6.5.6 Model Predictive Control

The sixth architecture replaces RRC, LIC and FIC with one model predictive controller (MPC) that plans, every $T_s = 1$ s, the pump demand $f_{p,\mathrm{cmd}}$, the feedwater demand $W_{\mathrm{fw,cmd}}$ and the setpoint p^R handed to PIC, which still moves the valves at 20 Hz, with RWM behind the pumps. Its model is Section 6.3 at the declared values of Table 6.2 and the design disturbances, never the run's, so an edited parameter is model mismatch, and it reads only the loops' transmitters, commands and setpoints. Its 21 states are the plant's without T_{fw}, z_o and z_v, the three transmitter lags, the PIC integral at its live tuning, and unmeasured inputs d_ρ, d_E and d_M added to the reactivity, energy and mass balances, so the void feedback of (6.6) and the level of (6.15) are in every prediction. Each decision linearizes it about the estimate, with an exact hold and the rods frozen. Against ±5 % pump and feedwater steps and ±0.05 MPa setpoint steps from full power, 120 s each, one decision's linearization stays within 0.062 %, 1.98 mm and 0.000 35 MPa of the plant, close because the model is the plant's own balance.

Each plan minimizes Equation (6.22) over $N = 120$ steps in twelve move blocks, from 1 step to 75.

$$\mathscr{J} = \sum_{k=1}^{N} \left[q_Q e_{Q,k}^2 + q_L e_{L,k}^2 + q_p e_{p,k}^2 \right] + \sum_i \left(a_i \varepsilon_i^2 + b_i \varepsilon_i \right) + \lambda \sum_j \left[\Delta f_{p,j}^2 + \Delta W_{\mathrm{fw},j}^2 + \left(\Delta p_j^R / 0.01 \right)^2 \right] \tag{6.22}$$

The errors are power in %, indicated level in mm and pressure in MPa less their setpoints, with $q_Q = 1$, $q_L = 0.01$ and $q_p = 10^4$, so a percent, 10 mm and 0.01 MPa cost alike, and Δ is a move between blocks, priced at $\lambda = 100$. The bounds are hard: RRC's 28–100 %, FIC's feedwater limit, p^R inside the pressure band, and 2 %/s, 5 %/s and 0.01 MPa/s, never active in scenario 12 (largest moves 0.75 %, 0.27 %, 0.003 75 MPa), where λ alone smooths the moves. One slack ε_i each softens the ± 100 mm level band, the ± 0.10 MPa pressure band and the licensed 100 % power, at a_i of 1 per mm^2, 10^6 per MPa^2 and 1.2×10^7 per percent and b_i of 10 per mm, 10^4 per MPa and 1.2×10^7 per percent, judged every second step. A level or power bound already violated returns from the present value with a 30 s or 10 s time constant. Each decision is one receding-horizon quadratic program [17, 18] of 39 variables and 447 rows, solved by a dual active-set method [19] in a mean 3.7 ms in Chrome, 0.37 % of a core at real time. Section 1.6.3 sets out the shared form of the controller and its estimator.

A steady-state Kalman filter at the declared design point corrects the model every scan under every architecture from the three transmitters (σ of 0.002 MPa, 5 mm and 0.3 %), with d_ρ, d_E and d_M as random walks of variance $\rho_K (0.05\,\mathrm{pcm})^2$, $\rho_K (500\,\mathrm{kW})^2$ and $\rho_K (0.5\,\mathrm{kg/s})^2$ per scan and $\rho_K = 10$, which removes offset [20]. A grid pass on scenarios 12 and 8 chose λ and ρ_K among settings that held both bands without a trip or more travel than Coordinated. The Coordinated loops track the applied commands with no preload and take over, without a bump, whenever a plan is rejected as suboptimal, nonfinite or out of bounds, until the next valid plan. On that or a switch away p^R returns to the operator's setpoint at 0.01 MPa/s. A scram, an isolation or a Level 8 trip, whose logic reads no architecture, stands the MPC down for the run, so it solves and writes nothing after it, and a move of the pump station (runback, pump trip or operator) pins the pumps until Reset.

6.6 Results

Table 6.5 lists all thirteen scenarios in the order they are meant to be run, the mechanism each one shows and what completes it. Each one presets the architecture, all five sets of tuning constants, the disturbance sliders and whatever timed events it needs. Briefings, completion credit, Reset and the challenge rules are those of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance, including firing a disturbance button, voids the attempt, while the prescribed scenario events remain fair.

6.6.1 Single-Element Against Three-Element Control

Scenarios 5 and 6 are the same plant, the same recirculation ramp from 95 % to 70 % pump drive in ten steps over 60 s, and the same level gains of $K_c = 30\,\%/\mathrm{m}$ and $\tau_I = 90\,\mathrm{s}$. The only difference between the two runs is where the steam flow measurement goes. Both end at 77.2 % core flow, the single-element run at 83.9 % power with the steam flow 17.0 % below rated and the three-element run at 84.7 % power with the steam flow 15.4 % below. Table 6.6 reports the outcome.

Single element reaches +287.4 mm and spends 221.6 s outside the control band, against −63.0 mm and 0.0 s for three element, a factor of 4.6 on the same disturbance with the same gains. The mechanism is worth following. Power comes off, so steam flow comes off, so the level should rise, and it does, but slowly at first, because the falling power also means less steam below the surface, and less steam pulls the indication down. The single-element controller therefore sees almost nothing for the first thirty seconds and keeps feeding at full rate, and by the time the mass imbalance appears in the indication it is minutes of overfeeding old. The collapsed level confirms it: the vessel really took on 428.8 mm of water it did not need, against 166.4 mm for the three-element run. Figure 6.5 is that

Table 6.5 Scenarios and Completion Requirements

#	Scenario (kind)	Mechanism	Required to complete
1	Full power by hand *experiment*	four actuators, one internal state	In Manual, rods at 33 % with power above 101.5 % for 10 s, pumps at 90 % with power below 98 % for 10 s, turbine valves parked at 86 or 87.5 % until the dome pressure moves less than 0.2 kPa/s for 30 s, then feedwater at 105 % for 2 min.
2	Shrink: the level goes the wrong way *experiment*	shrink against the inventory	With the pressure regulator in service, the indication below −20 mm while the collapsed level is above zero within 10 s of the 110 % feedwater step at 30 s, then the run continued to 100 s through the return to 100 %.
3	Power without moving a rod *experiment*	flow as the power handle	Pumps at 80 % with no rod travel for 30 s between 150 and 220 s, then the run continued to 300 s through the return to 95 %.
4	The same power change with rods *experiment*	self-regulation and rod speed	The bank seen in transit from 35 to 42 % inserted, then total reactivity within 1 ¢ of zero for 20 s after 250 s.
5	Single-element feedwater through a load change *demonstration*	inverse response inside the loop	The indication above +200 mm after the 95–70 % pump ramp, and the full 400 s run.
6	Three-element feedwater, same load change *experiment*	the structural fix	The same ramp run to at least 300 s in three-element control at $K_{\text{ff}} = 1.0$ with less than 5 s outside ±100 mm.
7	Feedwater pump trip *demonstration*	swell while the inventory falls	The indication above +100 mm while the collapsed level is below zero within 10 s of the trip, the Level 3 scram, and 30 s after it.
8	Loss of feedwater heating *experiment*	reactivity through the condensate system	Power above 108.5 % for 60 s after the heater loss, then a switch to Manual and pumps lowered until power stays below 103 % for 10 s.
9	A pressure step is a reactivity step *demonstration*	pressure as reactivity, no equilibrium with the level loop in service	The parked-valve run until the isolation valves have closed and the reactor has scrammed, then Reset and the regulator-fails-closed event, ending in a Level 3 scram with the isolation valves open and no high-pressure trip.
10	Recirculation pump runback *experiment*	the map, and which way a runback moves	The indication below −150 mm between 40 and 80 s, then the decay ratio above 0.4 for 20 s after 200 s.
11	Into the exclusion region *demonstration*	the plant, not the loop, going unstable	The decay ratio above 0.7 for 20 s on the first plateau (150 to 220 s), then the run continued to 300 s through the second flow reduction.
12	Void feedback in the model *experiment*	three handles planned against the swell	The full 600 s under MPC, then Coordinated, Reset and the 600 s again: the load ramp of the challenge with every loop set up well, and a pressure setpoint step at 420 s.
13	Challenge: 70 to 100 % on recirculation flow *challenge*	structure against tuning	Over 600 s with the model parameters and disturbances as given, the indicated level within ±100 mm and the dome pressure within ±0.10 MPa throughout, the decay ratio below 0.60, and the power within 0.5 % of 100 % from 480 s.[a]

[a] As given, the challenge fails with a worst indicated level of 410 mm, four times the band, a worst pressure error of 0.259 MPa, a peak decay ratio of 0.47 and the power still 26.15 % from 100 % over the last 2 min. With the suggested configuration it passes with 67 mm (33 mm inside the band), 0.021 MPa, 0.45 and 0.07 %.

run at its worst point, with both surfaces drawn on the vessel view and the two readings 141 mm apart. Single element is not badly tuned here. It is badly informed, and the fact that would have solved it was sitting on a transmitter the loop was not connected to.

Table 6.6 Single-Element Against Three-Element Feedwater on an Identical 95–70 % Recirculation Ramp With Identical Level Gains, Read Over 400 s

	Single element	Three element
Worst indicated level (mm)	+287.4	−63.0
Time of that worst value (s)	221	130
Seconds outside ±100 mm	221.6	0.0
Indicated level at 400 s (mm)	+33.2	+2.7
Worst collapsed level (mm)	+428.8	+166.4
Feedwater demand range (%)	80.1–100.2	84.3–100.0

> **TRY IT IN THE APP** **6 · Three-Element Feedwater, Same Load Change**
>
> Load scenario 6, press Run and let the recirculation ramp from 95 % to 70 % play. Read the worst indicated level, −63.0 mm, and the time outside ±100 mm, none. Then load scenario 5, the same ramp on single-element control, and run all 400 s: +287.4 mm and 221.6 s outside the band.

What did not change is instructive. At 395 s the three-element run holds the indication at +3 mm while the collapsed level reads +147 mm, a standing gap of 144 mm, because at 85 % power the vessel needs more water to show the same level. The gap does not close, and it is not supposed to. Three-element control does not remove the non-minimum-phase behavior from the plant. It removes it from the feedback path.

6.6.2 Runback and the Exclusion Region

Running the recirculation drive back to 45 % takes the core flow from 100 % to 56.2 % and the power down the 35 % rod line from 100 % to 68.9 %, with no rod motion. The operating point moves further left than it moves down. The power-to-flow ratio rises from 1.00 to 1.23 and the decay ratio from 0.250 to 0.460, with a transient peak of 0.501. The level dips −247 mm at 14.2 s on the way, while the collapsed level stands above its starting value. The core-average void rises as the flow comes off, but the falling power cuts the steam passing through the separators, and the steam held below the surface falls with it. The feedwater that the three-element controller cuts to match the falling steam flow does not deepen the dip, because the vessel is gaining water at that instant. A runback is normally a protective action that prevents something worse, and this is its price: every runback moves the plant toward the corner of the map where it is least stable, and it gets there by removing power.

Under the MPC the runback pins the pumps at 45 % at 40.005 s, the scan it acts under three-element, and the indication rises to +114.9 mm, 5.7 s outside the band against three-element's 35.3 s. Under Coordinated the runback and the trip of both pumps also pin the pumps on their station until Reset, with the power loop and the rod controller standing aside. Fired at 400 s of the suggested challenge, the runback takes the power to 72.4 % (MPC 70.8 %). The trip of both pumps scrams Coordinated on low level at 412.5 s, while the MPC keeps the level within 241 mm.

Scenario 11 goes the rest of the way, from a high rod line with the ordinary boiling noise switched on. Table 6.7 reports the three plateaus.

At the first plateau the decay ratio is 0.765, the exclusion-region warning has already annunciated at 92 s, and nothing is out of limits. The void oscillation amplitude has grown by a factor of 1.37, from 0.281 % to 0.385 % peak to peak, but the power noise has grown only from 5.99 % to 6.41 % peak to peak and its standard deviation from 1.23 % to 1.24 %, so the growth is visible on the void and not yet on the flux. The second reduction is unambiguous. The decay ratio crosses unity 14 s after the flow reduction, the power oscillation grows from 5.4 % to 13.6 % peak to peak over the following 15 s, and the oscillation scram fires 29 s after the reduction. That is 15 s after the decay ratio first reached unity, because it stayed there and the trip logic counts 15 s of it. The low-level scram follows 4 s later, so the oscillation trip is the one that acts.

Table 6.7 Two Recirculation Reductions From a High Rod Line, With Boiling Noise at 1.5×. Values Over 20–60 s, 150–220 s, and From the Decay Ratio Reaching Unity to the Oscillation Trip (234–249 s).

	Initial	Plateau 1	Plateau 2
Pump demand (%)	62	35	12
Core flow (%)	70.9	49.8	38.3
Thermal power, mean (%)	99.2	80.1	68.4
Decay ratio	0.504	0.765	1.00–1.07
Power, peak to peak (%)	5.99	6.41	13.8
Power, std. dev. (%)	1.23	1.24	3.10
Void oscillation, peak to peak (%)	0.281	0.385	0.952

Table 6.8 The Challenge Ramp With One Change at a Time From the Configuration as Given. Every Row Is the Same 600 s Disturbance Sequence, and Off Is the Time the Indicated Level Spends Outside ±100 mm.

Change from as given	Worst L (mm)	Worst p error (MPa)	Power range (%)	Off (s)
none	410	0.2589	59–113	412
K_{ff} 0.6 to 1.0	384	0.2524	60–113	404
gain scheduling on	3097	1.5051	11–109	540
rod controller on	77	0.1661	67–100	0
six tunings restored	170	0.0216	70–100	199
tunings and K_{ff}	47	0.0217	70–100	0
all four (suggested)	67	0.0206	66–100	0
architecture to MPC	13	0.0459	70–100	0

What made the plant unstable was the flow and not the power. The power fell at every step. That is why the exclusion region occupies the upper-left corner of the map and why it is entered by reducing flow, and it is a control problem in which the plant rather than the controller is the thing that goes unstable [5, 12].

> **✦ DID YOU KNOW? Less Power, Less Stable**
>
> From the simulator: as the recirculation flow is cut twice from a high rod line, the mean power falls from 99.2 to 68.4 percent, yet the decay ratio climbs from 0.504 to between 1.00 and 1.07 and the power swing grows from 5.99 to 13.8 percent peak to peak (Table 6.7).

6.6.3 Graded Challenge: Structure Against Tuning

The challenge ramps the power setpoint from 70 % to 100 % in six steps between 60 s and 260 s and asks for the indicated level held inside ±100 mm, the dome pressure inside ±0.10 MPa and the decay ratio below 0.60 for ten minutes, with the power within 0.5 % of rated over the last two. The coordinated architecture is selected and configured badly on purpose: the feedforward gain is 0.6 instead of 1.0, the level loop runs at 130 %/m with a 20 s reset instead of 30 with 90, the power loop at 4.5 %/% with a 60 s reset instead of 1.2 with 15, the pressure loop at 22 %/MPa with a 30 s reset instead of 80 with 6, and both the gain scheduling and the rod controller are switched off. Table 6.8 changes one thing at a time from that starting point, each applied at $t = 0$ before the run.

The suggested configuration meets the target, with a worst indicated level of 67.0 mm, a worst pressure error of 0.0206 MPa, a peak decay ratio of 0.450 and 100.03 % power reached at 97.0 % core flow. The only rod motion in the run is the valve position controller walking the bank from 35.00 % to 33.73 %, a total travel of 6.07 %, which is what relieves the pumps and leaves the fast handle authority in both directions. Reading the table row by row separates the four fixes, and two of

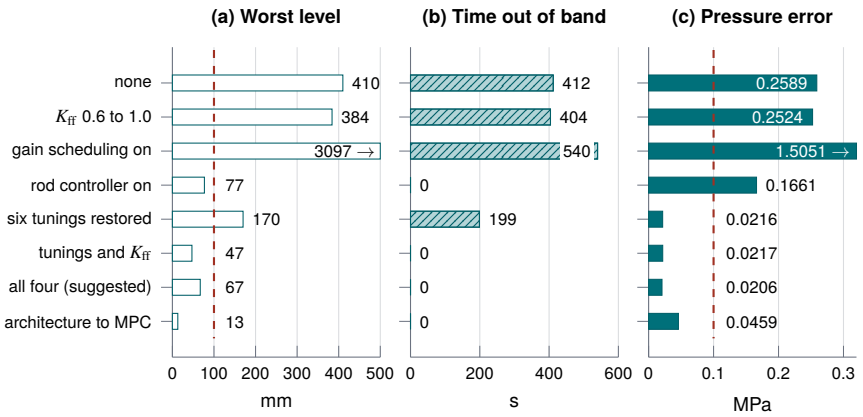

Figure 6.4 Table 6.8 drawn as bars: the challenge ramp with one change at a time from the configuration as given. (a) Worst indicated level, with the 100 mm band as the dashed line. (b) Time the indicated level spends outside ±100 mm. (c) Worst pressure error, with the 0.10 MPa band as the dashed line. The gain scheduling row runs off panels (a) and (c), so its bars end at the edge with the value printed. Notice that the rod controller alone does more for the level than the six restored tunings, and that only the last three rows stay inside both bands.

them do not behave the way the usual argument for them predicts. Figure 6.4 draws the three columns of the table as bars.

Restoring the six tuning constants does the pressure work on its own, taking the worst pressure error from 0.2589 MPa to 0.0216 MPa, and about half the level work, from 410 mm to 170 mm. Adding the feedforward on top finishes the level at 47 mm, which is better than the suggested configuration achieves. The feedforward alone with the original gains buys almost nothing, 384 mm against 410 mm, because a power loop oscillating between the pump stops generates a level disturbance far larger than the feedforward corrects. Turning the rod controller on alone is the single most effective change, 410 mm to 77 mm, because keeping the pumps away from their limits stops the excessive power gain from slamming them.

Turning the gain scheduling on alone is catastrophic. The plant scrams on low level at 103.7 s and trips on Level 8 at 145.9 s, and with the post-scram void collapse included the run records a worst level of 3097 mm, a worst pressure error of 1.5051 MPa and the power falling to 11 %. Equation (6.21) multiplies whatever gain it is given, and at the starting 46.7 % pump drive the recirculation factor is 2.03, so an already excessive K_c of 4.5 becomes 9.15. With the gains and the feedforward correct the scheduling makes no useful difference either, giving 102.6 mm against 46.5 mm with the rod controller off and 67.0 mm against 67.5 mm with it on. That agrees with Section 6.4: the process gain from pump speed to power changes by 6 % across half the pump range, not by the factor of two the scheduling law is scaled for.

Switching only the architecture to the MPC of Section 6.5.6 also meets the target, with a peak decay ratio of 0.450, at a pressure IAE of 9.0 against 2.0 MPa s for the suggested configuration, because the slow pressure loop as given still sits under it. The challenge is unchanged.

6.6.4 The Design Point and What the Application Verifies

Every scenario opens at the state in Table 6.9, which the Newton solve of Section 6.3 reaches to machine precision. Holding the design scenario for 60 s moves no reported quantity in the fifth decimal place.

Thirteen application-specific checks run headlessly on the shipped model, in addition to the framework's own structural, bumpless-transfer, actuator-bound, CSV-consistency and reproducibility

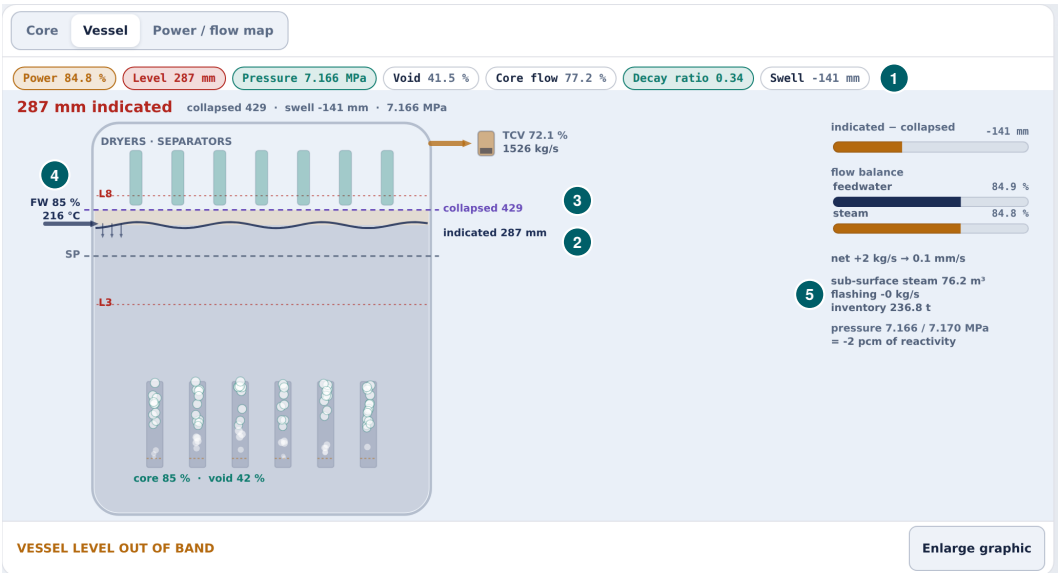

Figure 6.5 Single-element feedwater control in the Vessel view of the app at $t = 3.68$ min, at the worst point of the 95–70 % recirculation ramp of scenario 5, which is the failure the comparison of Section 6.6 rests on. Callouts: (1) the status tiles, (2) the indicated level, (3) the collapsed level, (4) the feedwater inlet with its demand and temperature, (5) the readouts of sub-surface steam and inventory. The vessel drawing carries both surfaces and they disagree by 141 mm: the level transmitter reads 287 mm against a setpoint of zero while the collapsed level, which is what the water inventory actually is, stands at 429 mm. The controller is acting on the smaller of the two numbers, so it is feeding a vessel that is already 236.8 t full, and the scene reports the indication out of band. The plant holds 84.8 % power on 77.2 % core flow at 7.166 MPa, and the void fraction has risen to 41.5 % as the flow came off. The turbine valves have closed to 72.1 % against a steam flow of 1526 kg/s, feedwater and steam are within 2 kg/s so the surface is nearly stationary at 0.1 mm/s, and the 7.166 against 7.170 MPa pressure error is annotated on the panel as the −2 pcm of reactivity it is worth.

Table 6.9 The Design Operating Point, Solved Rather Than Assumed

Quantity	Symbol	Value	Quantity	Symbol	Value
Core thermal power	Q_{th}	3330.0 MW	Core exit void fraction	$\alpha(x_e)$	65.23 %
Generator output	P_e	1100.1 MW	Core-average void fraction	$\bar{\alpha}$	40.32 %
Dome pressure	p	7.170 MPa	Non-boiling fraction	f_{nb}	10.57 %
Saturation temperature	T_{sat}	287.49 °C	Average fuel temperature	T_f	699.49 °C
Core flow	W_c	13 100 kg/s	Rod bank insertion	x_{rod}	35.0 %
Core mass flux	G	1506 kg/m^2/s	Recirculation drive	f_p	95.0 %
Steam flow	W_{st}	1800.2 kg/s	Turbine valve position	y_T	85.0 %
Feedwater enthalpy	h_{fw}	920.9 kJ/kg	Sub-surface steam volume	V_{sd}	80.71 m^3
Core inlet subcooling	Δh_{sub}	48.87 kJ/kg	Sub-surface steam as level	V_{sd}/A_s	2522 mm
Core inlet subcooling	—	9.97 °C	Vessel inventory	M_v	227.2 t
Core exit quality	x_e	13.742 %	Decay ratio	DR	0.250

tests. They verify the equilibrium, the design point, the sign of every actuator path, the pressure-to-reactivity coefficient, the direction of every closed loop, the shrink transient, the flow-to-power gain with the rod bank frozen, the single-element against three-element comparison, the location of the

Table 6.10 Coordinated Loops Against the MPC at Scenario 12's Settings, Same Seed, Events and Noise. IAE in % s, mm s and MPa s and Travel From the 10 Hz Historian

Case and measure	Coordinated	MPC
Scenario 12, ramp, 0–420 s		
Worst level, 60–420 s (mm)	49.3	27.6
IAE power, level, pressure	438.0, 10 124, 2.016	500.7, 5788, 0.985
Pump travel (%)	140.7	47.2
Within 0.5 % of 100 % from (s)	314.6	302.4
Scenario 12, pressure step, 420–600 s		
Level dip (mm)	−114.9	−17.7
IAE pressure	0.224	0.551
IAE power, level, pressure over 600 s		
Scenario 9, regulator out	143.6, 3163, 182.9	93.6, 1820, 177.9
Case 1, admission −5 %	72.9, 1585, 0.481	56.6, 1712, 0.579
Scenario 8, heater loss	194.2, 5885, 0.507	193.9, 4297, 0.765

exclusion region, the equilibrium of the three-element scenario before its ramp, that an architecture or AUTO switch on a running plant is bumpless, and that the challenge target is met by the suggested configuration and missed by the one given. Fourteen more check the MPC. All forty-seven checks pass on the shipped code, as do the fourteen checks that walk each scenario's completion rule.

6.6.5 Model Predictive Control Against the Coordinated Loops

Scenario 12 is the challenge ramp with every loop set up well, then the pressure step of Figure 6.3, and Table 6.10 sets Coordinated against the MPC. No run in it trips or rejects a plan.

On the ramp the MPC cuts the level IAE by 43 %, the pressure IAE by 51 % and pump travel by two thirds, but loses the power IAE, 500.7 against 438.0 % s, because λ prices each pump move. The 67 mm of Table 6.8 is that Coordinated run's opening rod move at 4.7 s, not the ramp. On the admission step and the heater loss it loses pressure IAE, moving p^R to help power and level. With 10 mm of level noise and boiling noise at 1× it loses power IAE, 724.5 against 621.6 % s, and moves the feedwater 274 against 165 %.

Every trip is the plant's own, so an instant moves only with the plant's path. With the regulator failing closed at 150 s of scenario 12 the MPC scrams on high pressure later, at 189.08 s against 179.81 s, because its plan has no pressure term once the valves are parked, and after each of the six trips measured under it the MPC solved no plan. Switched from three-element at $t = 0$ it peaks at 105.29 % power, against 104.39 % into Coordinated, below the 120 % flux scram, so only the power tile turns amber.

6.7 Technology in Practice

The tags of Figure 6.1 are, in a plant, pressure, level and flow transmitters, neutron detectors in the core, and the controllers of a plant control system. Figure 6.6 draws the four loops with that hardware in place of the schematic's symbols.

What Measures It

Pressure transmitters convert the pressure force to an analog signal such as 4 to 20 mA or, in newer smart transmitters, to a digital signal, and the most common level measurement in nuclear plants is a differential pressure between two taps on the vessel. Neutron detectors range from fission and ion chambers to self-powered detectors, and the readings of many in-core detectors are combined to calculate the overall reactor power [21]. Measurements in the vessel of a boiling water reactor that monitor adequate cooling of the core have their own standard, IEC 61343 [21]. The chapter carries the pressure transmitter as a 0.15 s lag, the level transmitter as a 1.2 s lag on Equation (6.15), the

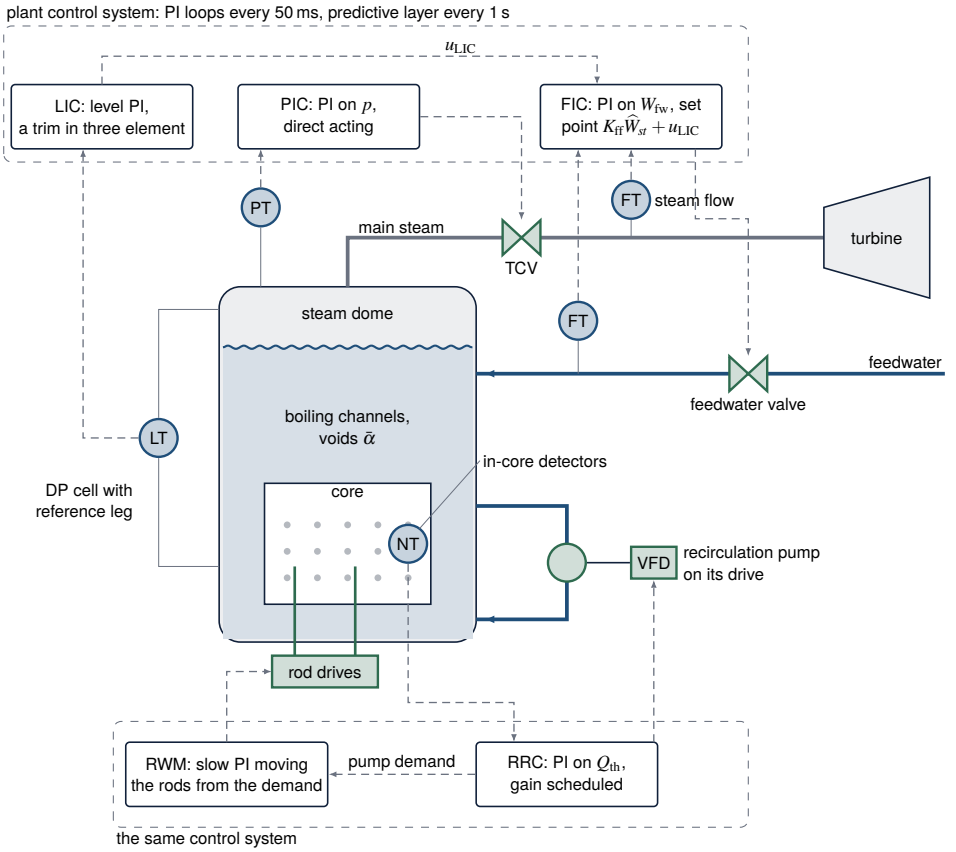

Figure 6.6 The four loops of the chapter with plant hardware in place of the symbols of Figure 6.1. Sensors are blue circles and final elements green: the turbine control valves (TCV) under PIC, the feedwater valve under FIC with steam flow fed forward and LIC trimming its set point, the recirculation pump on its variable-frequency drive (VFD) under RRC, and the rod drives under the slow valve position controller RWM, which reads the pump demand. Notice that the level transmitter's two taps straddle the boiling surface, so the reading moves with the steam below it.

power signal through a 0.6 s filter and the flow transmitters as 0.8 s lags. Figure 6.7 shows why the level reading is a head and not a height.

> **TECHNOLOGY Sensor: The Level Transmitter and Its Reference Leg**
>
> Most level measurements in nuclear plants are differential pressures between a tap low on the vessel and a tap high on it, with a filled reference leg on one side [21]. The pressure of a liquid column gives its height only if the density of the liquid is known and constant. The reference leg holds a constant head only while its own density stays constant, and a zero shift at calibration cancels that head. With the leg on the high port the output falls as the level rises, so the connection is chosen for a fail-safe reading [22]. An impulse line that clogs, traps gas or fouls makes the reading inaccurate or sluggish [21]. The chapter's LT converts with the densities of its ranging pressure, behind a 1.2 s lag.

What Moves It

The final elements are the turbine control valves, the feedwater supply, the two recirculation pumps and the control rod bank. In the chapter the valves stroke as a 0.25 s lag, the feedwater flow follows its demand through a 3.5 s lag and the pump drive through a 5.0 s lag, while the rods are rate limited at

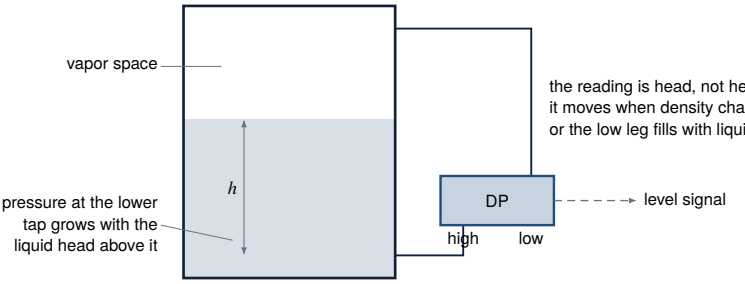

Figure 6.7 Level by differential pressure, from the book's instrument library. Notice that the transmitter reads the head between its two taps, which depends on the density of what lies between them. In this chapter that is a mixture of water and steam, which is why the indicated level of Equation (6.15) can fall while the vessel fills.

0.55 %/s. These are idealizations with no dead band and no slew limit on the pumps or the feedwater, as Section 6.8 states.

> **TECHNOLOGY** **Final Element: The Recirculation Pump on a Variable-Frequency Drive**
>
> Changing the flow through the core is the normal way to move a boiling water reactor's power between about 30 and 100 percent along its highest rod line, because more flow sweeps the voids out faster. The pump speed comes from a variable-frequency drive, an inverter that sets the motor's supply frequency and keeps the voltage in proportion. Such drives have replaced the older motor-generator sets, and one plant reported a fourfold rise in the reliability of its recirculation system after the change. Newer digital drives carry redundant power cells: a drive can lose two and still deliver full power, and a failed cell is swapped in eight cycles with a small loss of speed [23]. The chapter's drive follows its command through the 5.0 s lag τ_p.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, and nested loops in which the set point of one loop comes from the measurement of another [24]. The chapter's loops execute every 0.05 s over a 5 ms integration step, and the predictive controller plans every second and hands its pressure setpoint to PIC. The structure is a set of function blocks: a direct-acting PI on dome pressure, a level PI whose output is added to the steam flow feedforward to set a feedwater flow PI, a power PI on the pump demand, a slow valve position PI that moves the rods from that demand, and a gain scheduler on two of them.

> **IN PRACTICE** **Power Changes Ride on the Recirculation Flow**
>
> A boiling water reactor that operates flexibly can follow load on recirculation flow alone, along a curve of power against core flow at a fixed rod position. This suits the upper range, about 60 to 100 percent of rated power, where the power changes faster and with fewer changes in its distribution across the core. With no rods moving, the fuel sees smaller temperature swings, and the pressure and temperature in the main steam system stay constant. Changes beyond the recirculation range are made with the control rods [25]. The chapter's rod controller moves the bank only slowly, to walk the pumps back toward the middle of their range.

What Is Modernizing It

Three developments reach these loops. Variable-frequency drives have replaced motor-generator sets on recirculation pumps, and twenty-six were in operation in boiling water reactors in the United States by 2014 [23]. On-line monitoring mines the data the control system already records for changes in the relationships between measured parameters, and can extend the calibration period of a sensor that has

not drifted [21]. Operating plants are moving from analog to digital control through individual system upgrades [21], and control systems that adopt office computers, operating systems and network protocols have far less isolation from the outside world and a greater need to be secured [24].

6.8 Limitations

The model is a control-structure teaching model and not a safety analysis code. The couplings, the signs, the gains and the time constants are in the right places and of the right size, and the lessons transfer directly, but there is no hot channel, no critical power ratio, no regional neutronics and no emergency cooling, so the application cannot say anything about fuel damage, about margin to dryout, or about what happens more than a few minutes after a scram. The decay ratio of (6.16) is a calibrated correlation rather than a solved channel, so the onset boundary is meaningful and the limit-cycle amplitude is illustrative, and the drift-flux correlation is least reliable exactly in the low-flow corner the application most wants to demonstrate. The turbine has no bypass valves and no shaft dynamics, which removes the system a real plant relies on to absorb a load rejection. Disturbances enter as commanded steps and scheduled events instead of as the slow wandering a plant does, and the exact repeatability that a frozen step and a frozen scan buy is itself an unreality. Where each individual assumption breaks down is set out in Section 6.3.8.

The pump drive and the feedwater flow have no rate limits. They follow their commands through the first-order lags τ_p and τ_{fw} alone, where a real drive and a real feedwater pump set would also be slew limited. The scenarios step these demands by up to 50 points at once (the runback), so a limit of a few percent per second would stretch each step into a ramp of many seconds and change the fast transients this chapter measures.

The MPC's model leaves out the density-wave oscillator, so the decay ratio and the exclusion region are not in its plan. It also leaves out the noise, the feedwater temperature, the admission capacity, the isolation and relief valves, the trips, the rods' future moves and the regulator's anti-windup.

> **OPEN CHALLENGE** **Oscillations the Core Average Can Hide**
>
> Uprated boiling water reactors may run at up to 120 percent of their original licensed power with core flow as low as 80 percent of rated. Simulations of a turbine trip without scram from there show the core going unstable once both recirculation pumps trip. They also show that stability cannot be judged from the core average power alone: as the oscillation moves from the core-wide mode to higher modes, the average may fall while single channels swing harder. Good data at these conditions are scarce, so the predicted cladding temperatures carry a large uncertainty [26]. Progress would be monitors and models that resolve the regional modes, which the point model of this chapter cannot represent, checked against such data.

6.9 Discussion

Three of the results in Section 6.6 generalize past this plant. The first is that a structural change beat a tuning change by a factor of 4.6 on the same disturbance with identical gains, and that no gain on the single-element loop reaches the three-element result, because the limiting quantity is a right-half-plane zero rather than a margin [3]. The second is that a scheduling law scales whatever gain it is handed. Turning the scheduling of (6.21) on with the challenge's excessive power gain in place multiplied the worst level excursion by a factor of 7.6, through a low-level scram, and turning it on with correct gains bought nothing (67.0 mm against 67.5 mm with the rod controller in service, 102.6 mm against 46.5 mm without it), because the process gain it was written to compensate does not vary the way the derivation assumed. A scheduling rule is a claim about a gain, and the gain is worth measuring before the rule is written. The third is that a computed quantity can carry more diagnostic value than a measured one. The collapsed level of (6.14) is what an operator needs during any feedwater transient and no plant instrument measures it, and in Figure 6.2 it is the only signal that never has the wrong sign.

The engineering consequence is that the level indication and the trips reading it are correct instruments giving misleading information at exactly the moment the information matters. A reactor tripped on low water level 0.835 s after the turbine valves drift shut, while the vessel is gaining 276 kg, is neither an instrument failure nor a controller failure. It follows from a differential-pressure cell measuring the mixture it straddles, and the remedy in service is a collapsed-level calculation on the operator display built on (6.13) and (6.14) from signals a real plant already has.

Methods beyond the implemented architectures fit the same equations. Seventeen states with hard actuator limits, a right-half-plane zero and a stability requirement written on the operating point are a nonlinear model predictive control problem, and posing it that way turns the exclusion region of (6.16) into a constraint instead of a procedure and coordinates the feedwater, recirculation and valve moves as one problem rather than four [27]. The linear predictive controller of Section 6.5.6 is the first step, and it leaves the decay ratio out of its plan. The broader operational context is that reactors are increasingly asked to follow variable renewable generation rather than to hold base load [28], which moves them toward the reduced-flow part of the map where the recirculation gain, the level loop's zero and the decay ratio are all at their least favorable.

6.10 Conclusion

This chapter studied a boiling water reactor as an interactive laboratory for coupled loops with sign-reversing dynamics. The simulator integrates seventeen states, including two-group point kinetics with Doppler and void feedback, a drift-flux boiling channel with a subcooled-boiling correction, jet-pump recirculation against a quadratic loop resistance, a dome pressure on a lumped energy balance, and a vessel inventory split exactly into liquid, sub-surface steam and dome steam so that the indicated and the collapsed level are separate outputs. Measured on the running model, recirculation flow moves 0.659 % of rated power per percent of core flow with 63 % of the change in 3.3 s and no rod motion, a 0.10 MPa pressure change is worth 43.4 pcm, the plant returns 30.6 pcm of negative feedback per percent of power, and 28 °C of lost feedwater heating raises power 9.45 % with nothing on the reactor touched. A +10 % feedwater step drives the indication −40.9 mm for 6.10 s while the inventory gains 193 kg. Three-element control holds a 95–70 % recirculation ramp to −63.0 mm and no time outside the band, against +287.4 mm and 221.6 s for single element at identical gains. On the graded challenge, correcting the six tuning constants and the feedforward gain reduces the worst level excursion from 410 mm to 47 mm, while enabling the gain scheduling with the original gains in place scrams the plant on low level and raises it to 3097 mm. A model predictive controller cuts the pressure step's level dip from 114.9 mm to 17.7 mm, at 2.5 times the integrated pressure error. A fuller model would add a two-node axial power shape so that the regional oscillation mode appears, a hot channel with a critical power ratio, turbine bypass valves and a shaft, and xenon dynamics. The boiling water reactor is an accessible interactive plant on which to learn how loops are paired when the plant is already controlling itself.

KEY POINTS

- The steam bubbles in the core slow the chain reaction, so moving the valves, the feedwater, the pumps or the rods changes the bubbles, and each loop pushes on the other three.
- A +10 % feedwater step drives the indication −40.9 mm while the inventory gains 193 kg.
- Three-element control holds the ramp to −63.0 mm, against +287.4 mm for single element at identical gains.
- A scheduling rule is a claim about a gain, and the gain is worth measuring before the rule is written.
- The exclusion region is entered by reducing flow.

Where else this applies. Chapter 4 carries the same three feedwater structures on a drum, and Volume 6, Chapter 2 makes the same pairing decision before any loop is tuned.

References

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

[2] K. Iinoya and R. J. Altpeter, "Inverse response in process control," *Industrial & Engineering Chemistry*, vol. 54, no. 7, pp. 39–43, 1962.

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

[4] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[5] 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.

[6] C. Lin and L. M. Grossman, "Optimal control of a boiling water reactor in load-following via multilevel methods," *Nuclear Science and Engineering*, vol. 92, no. 4, pp. 531–544, 1986.

[7] J. A. Mowrey, P. R. Boylan, and S. I. Abdel-Khalik, "Modeling of boiling water reactor feedwater control algorithms," *Nuclear Technology*, vol. 124, no. 3, pp. 223–242, 1998.

[8] N. Zuber and J. A. Findlay, "Average volumetric concentration in two-phase flow systems," *Journal of Heat Transfer*, vol. 87, no. 4, pp. 453–468, 1965.

[9] P. Saha and N. Zuber, "Point of net vapor generation and vapor void fraction in subcooled boiling," in *Proceedings of the 5th International Heat Transfer Conference*. Begell House, 1974, pp. 175–179.

[10] W. C. Lipinski, "Control aspects of the experimental boiling water reactor power plant," in *1956 IRE National Convention*. IEEE, 1956, pp. 84–89.

[11] 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.

[12] F. Araya, K. Yoshida, M. Hirano, and Y. Yabushita, "Analysis of a neutron flux oscillation event at LaSalle-2," *Nuclear Technology*, vol. 93, no. 1, pp. 82–91, 1991.

[13] B. R. Upadhyaya and M. Kitamura, "Stability monitoring of boiling water reactors by time series analysis of neutron noise," *Nuclear Science and Engineering*, vol. 77, no. 4, pp. 480–492, 1981.

[14] 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.

[15] W. Wulff, H. S. Cheng, and A. N. Mallen, "Causes of instability at LaSalle and consequences from postulated scram failure," in *International Workshop on Boiling Water Reactor Stability*, Upton, NY, Oct. 1990, brookhaven National Laboratory report BNL-NUREG-45273. [Online]. Available: https://www.osti.gov/biblio/6397094

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

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

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

[19] 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.

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

[21] 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.

[22] 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.

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

[24] 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.

[25] 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.

[26] L.-Y. Cheng, J.-S. Baek, A. Cuadra, A. Aronson, D. Diamond, and P. Yarsky, "BWR anticipated transients without scram leading to instability," in *2013 ANS Winter Meeting and Nuclear Technology Expo*, Washington, DC, Nov. 2013, brookhaven National Laboratory report BNL-101129-2013-CP. [Online]. Available: https://www.osti.gov/biblio/1091775

[27] J. D. Hedengren, R. A. Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[28] A. Ho, B. W. Billings, J. D. Hedengren, and K. M. Powell, "Flexible operation of nuclear hybrid energy systems for load following and water desalination," *Renewable Energy Focus*, vol. 51, p. 100641, 2024.

6.11 Exercises

6.1 Concept check · 1 min

Which assignment of the four handles does the coordinated architecture of this chapter use?

(a) Pressure on the turbine valves, level on feedwater, power on the pumps, rods behind the pumps

(b) Power on the turbine valves, level on feedwater, pressure on the pumps, rods on core power

(c) Pressure on the pumps, level on the turbine valves, power on feedwater, rods on the level

(d) Power on the rods, level on feedwater, pressure on the turbine valves, pumps held fixed

6.2 Concept check · 2 min

With the pressure regulator in service and the level loop out, the feedwater demand steps from 100 to 110 %. The indicated level falls 40.9 mm in the first 3.45 s. What is happening in the vessel?

(a) The transmitter lags the surface by its 1.2 s time constant and still shows the old level

(b) The colder inlet water shortens the boiling length and condenses steam below the surface while the inventory rises

(c) The colder feedwater lowers the dome pressure, so the cell, ranged at the design pressure, reads low

(d) The vessel loses water at first, because the valve opens before the extra feedwater reaches the vessel

6.3 Concept check · 1 min

The model carries the equation $C_e\,\mathrm{d}p/\mathrm{d}t = Q_{\mathrm{th}} + W_{\mathrm{fw}}h_{\mathrm{fw}} - W_{\mathrm{out}}\left(h_f + h_{fg}\right)$. What is it?

(a) The mass balance on the vessel inventory that sets the level

(b) The point-kinetics balance on the neutron density

(c) The vessel energy balance that sets the dome pressure

(d) The drift-flux correlation for the channel void fraction

6.4 Theory and calculation · 6 min

The rod bank has a total worth $W_{\mathrm{rod}} = 0.030$ on the S-curve of Equation (6.6), and one dollar is $\beta = 0.0056$. The bank starts 35 % inserted.

(a) Compute the differential worth at 35 % in pcm per percent of travel and in cents per percent. *Checkpoint:* 47.6 pcm and 8.51 ¢ per percent

(b) Compute the integral worth of a move from 35 to 42 % inserted, in pcm and in cents, and compare it with seven times the answer to (a). *Checkpoint:* 366 pcm, 65.4 ¢, against 333 pcm

(c) The plant returns 30.6 pcm of negative feedback per percent of power. Estimate the settled power change after the move of (b) and compare it with the 11.46 % the chapter measures. *Checkpoint:* about 12.0 % lower

6.5 Theory and calculation · 6 min

At the design point the dome is at $p_0 = 7.17\,\mathrm{MPa}$, the turbine valves are at $y_{T,0} = 0.85$ and pass $W_{s,0} = 1800.2\,\mathrm{kg/s}$, and the energy storage is $C_e = 1.40 \times 10^7\,\mathrm{kJ/MPa}$.

(a) Compute the nozzle coefficient K_t of Equation (6.12). *Checkpoint:* 295.4 kg/s per unit valve per MPa

(b) Compute the saturated steam enthalpy $h_f + h_{fg}$ at p_0 from the fits of Equation (6.1). *Checkpoint:* 2770.7 kJ/kg

(c) Compute the initial rate of change of dome pressure per percent of valve opening, with power and feedwater held, and the initial rate after the 4 % park of scenario 9. Compare the second with the measured peak of 17.7 kPa/s and say which term of the balance makes up the difference. *Checkpoint:* −4.19 kPa/s per percent, −16.8 kPa/s for 4 %

6.6 Theory and calculation · 6 min

The decay ratio of Equation (6.16) has $\mathrm{DR}_0 = 0.25$, $a_{\mathrm{DR}} = 1.20$ and $b_{\mathrm{DR}} = 0.77$.

(a) After the runback of scenario 10 the plant settles at 68.9 % power on 56.2 % flow. Compute the decay ratio. *Checkpoint:* 0.459

(b) Compute the damping ratio ζ of Equation (6.17) at that decay ratio, and the decay ratio at which ζ reaches zero. *Checkpoint:* $\zeta = 0.123$, zero at $\mathrm{DR} = 1$

(c) With both pumps stopped the plant settles at 50.4 % power on 34.6 % flow. Compute the decay ratio there and say which of the two terms of the exponent contributes more. *Checkpoint:* 0.716

6.7 App challenge · 6 min

Load scenario 2, “Shrink: the level goes the wrong way”, as shipped (default seed). Press Run at the default speed of 2 on the Vessel view, pause between 40 and 50 s and export the CSV, whose time column is in seconds.

(a) Find the minimum of `L_ind_m` after the feedwater step at 30 s, its time, and `L_col_m` in the same row. Give the levels in mm to within 0.5 mm. *Checkpoint:* −40.9 mm at 33.4 s, collapsed +11.8 mm

(b) Read `V_steam_sub_m3` in that row and before the step, and express the change as a fraction of the design 80.71 m^3. Give the change to within 0.05 m^3. *Checkpoint:* −1.68 m^3, 2.1 %

(c) The step adds 180.0 kg/s. From the mass balance alone, at $\rho_f = 736.5\,\mathrm{kg/m}^3$ over $A_s = 32\,\mathrm{m}^2$, what would the level have risen by at the time of (a)? *Checkpoint:* 7.64 mm/s, about 26 mm

6.8 App challenge · 8 min

Load scenario 5, “Single-element feedwater through a load change”, as shipped (default seed), and run all 400 s at speed 10. Read the Off level tile and export the CSV, whose time column is in minutes. Repeat with scenario 6, “Three-element feedwater, same load change”.

(a) For each run give the Off level tile, the largest deviation of `L_ind_m` from zero in mm to within 1 mm, and its time in minutes to within 0.02 min. *Checkpoint:* single: 222 s, +287.4 mm at 3.69 min. three: 0 s, −63.0 mm at 2.17 min

(b) Give the largest `L_col_m` of each run in mm and say what the difference between the two runs means for the water put into the vessel. *Checkpoint:* 428.8 mm against 166.4 mm

(c) The level gains are identical. Explain in two sentences why single element loses by a factor of about 4.6. *Checkpoint:* the error it acts on reads backwards for the first tens of seconds

6.9 App challenge · 8 min

Load the challenge, “Challenge: 70 to 100 % on recirculation flow”, at the default seed, and change one setting at a time from the case as given, each before Run. Run each attempt the full 600 s at speed 10 and read the verdict line (worst level, worst pressure, peak decay ratio).

(a) Set Rods behind the pumps to On (Advanced) and leave everything else as given. Give the verdict values and say which target fails. *Checkpoint:* 77 mm, 0.166 MPa, 0.45; the pressure band fails

(b) Reload the challenge, set only Gain scheduling to On, and run. Name the first trip and its approximate time from the marks on the trends. *Checkpoint:* Level 3 low-level scram near 104 s

(c) Explain from Equation (6.21) why switching the scheduling on was harmful here. *Checkpoint:* $4.5 \times 2.03 = 9.15$

6.10 Application area · 3 min

A plant engineer who has run fired drum boilers asks why a boiling water reactor puts its turbine valves on pressure and moves power with the recirculation pumps. Answer with the two measured coefficients the chapter gives.

6.11 Application area · 4 min

After a trip of both recirculation pumps the plant settles on natural circulation at about half power and a third of rated flow. (i) Why is this the dangerous corner of the operating map, and how is it left? (ii) Name two limitations of this chapter’s model that bound what it can say about the event.

6.12 Review the agent · 5 min

An AI assistant reviewing the challenge writes: “Jet-pump head goes as the square of pump drive, so at 46.7 % drive the gain from pump speed to power is less than half its value at 95 %. Switch Gain scheduling on first. It scales the power-loop gain by $f_{p,0}/f_p$ and restores the loop’s performance at low drive.” Find the flaw and say how you would check it on the app before trusting the rule.

CHAPTER 7

Molten Salt Reactor Simulator Where the Pump Sets the Stability Margin

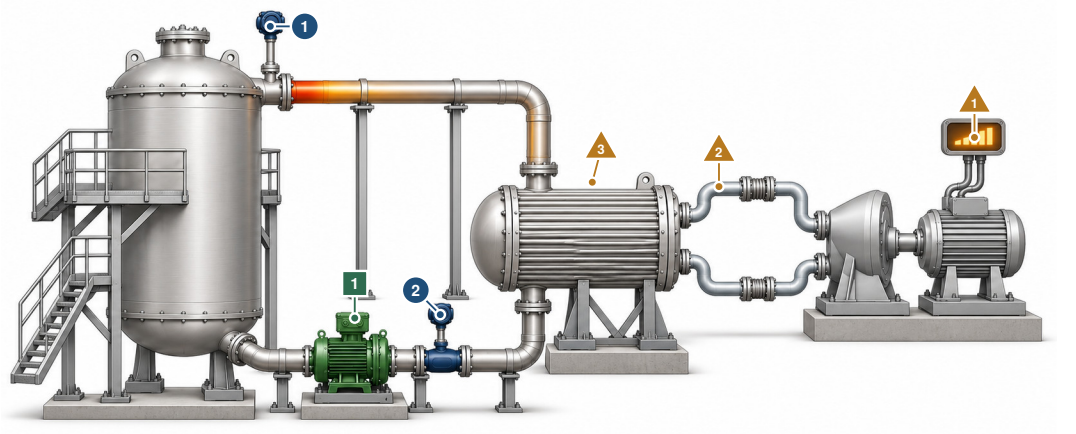

The controller moves 1 fuel salt pump speed

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

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

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

IN PLAIN TERMS A Pump That Also Sets the Safety Margin

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

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

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

THREE WAYS TO READ THIS CHAPTER

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

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

Learning objectives

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

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

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

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

7.1 Introduction

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

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

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

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

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

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

7.2 Background

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

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

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

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

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

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

7.3 Reactor Model

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

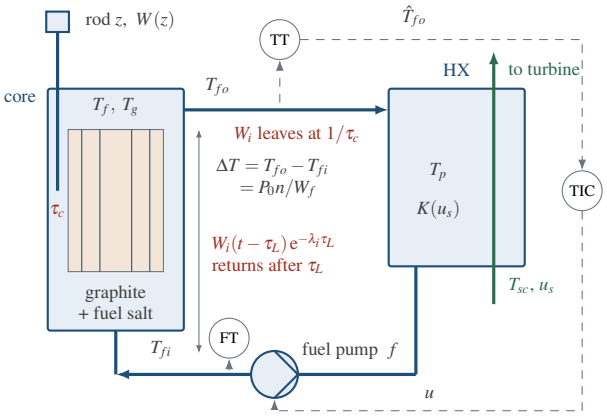

Figure 7.1 The fuel salt loop with the core, the hot leg, the heat exchanger and the fuel pump. Precursor group W_i leaves the core at a rate $1/\tau_c$, decays along the external loop by $e^{-\lambda_i \tau_L}$ and re-enters one loop transit later, and the same salt carries the heat, so the cold leg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

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

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

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

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

7.3.1 Prompt Jump and Circulating Precursor Kinetics

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

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

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

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

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

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

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

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

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

Pause. Try Exercise 7.2 before reading on.

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

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

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

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

DID YOU KNOW? Twenty-Five Seconds Round the Loop

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

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

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

7.3.2 Reactivity and the Criticality Condition

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

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

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

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

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

7.3.3 Thermal Balances of the Core, Loop and Heat Exchanger

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.3.4 Actuators and Instruments

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

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

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

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

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

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

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

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

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

7.3.5 Parameters, Initialization and Numerics

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

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

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

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

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

7.3.6 Assumptions

Each simplification below is stated with the place it fails.

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

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

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

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

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

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

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

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

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

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

7.4 Selection of the Manipulated Variable

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

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

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

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

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

7.4.1 Steady-State and Step Response of Three Candidates

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

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

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

Pause. Try Exercise 7.3 before reading on.

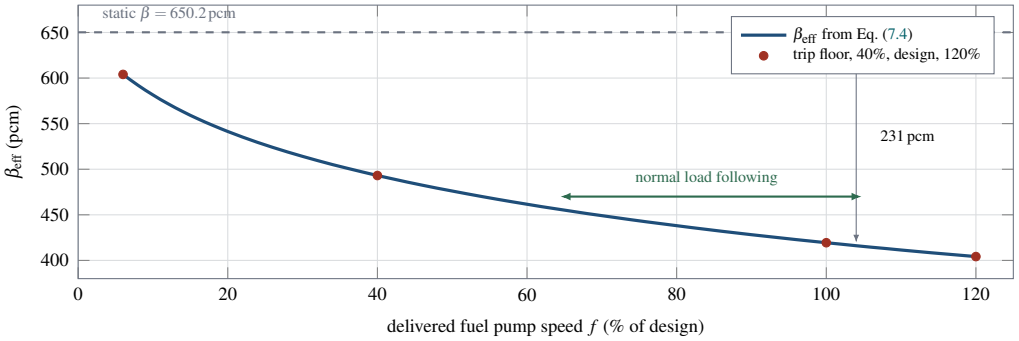

Figure 7.2 The effective delayed neutron fraction as a function of delivered fuel pump speed. The curve is Equation (7.4) evaluated with $\tau_c = 9/f$ and $\tau_L = 16/f$ in seconds. Holding the core outlet at 665 °C while the turbine moves from 110 down to 30 % load runs the pump between 104.8 and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

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

7.4.2 Effect of Pump Speed on the Stability Margin

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

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

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

7.5 Controller Design

7.5.1 Velocity Form PID Algorithm

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

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

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

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

7.5.2 Integrator Count and Steady-State Offset

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

7.5.3 Load Feedforward and Dynamic Compensation

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

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

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

7.5.4 Inner Flow Cascade

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

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

7.5.5 Model Predictive Control

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

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

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

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

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

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

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

Table 7.5 Scenario Sequence and the Mechanism Each Exercises

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

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

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

7.6 Results

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

7.6.1 Self-Regulation Without Control

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

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

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

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

7.6.2 Proportional Offset and the Transport-Limited Gain

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

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

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

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

7.6.3 Feedforward Timing on a Fast Load Ramp

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

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

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

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

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

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

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

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

7.6.4 Cascade Performance and Noise Response

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

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

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

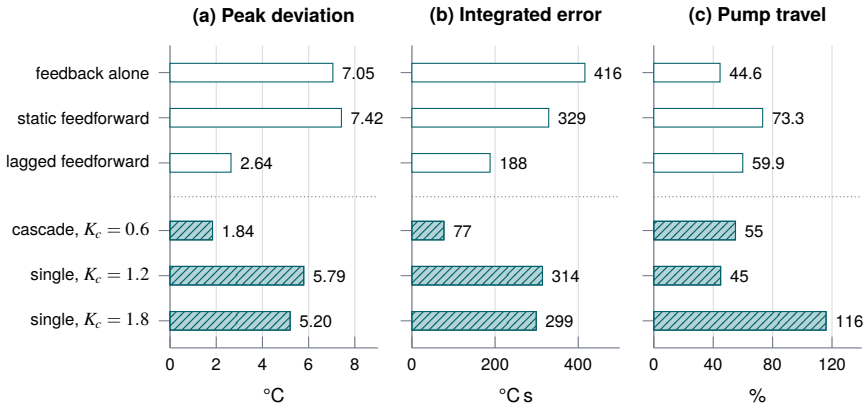

Figure 7.3 Tables 7.7 and 7.8 drawn as bars, every value read at $t = 600$ s. White bars, above the dotted line: the 45 % load reduction at 4 %/s of scenario 7 under three architectures. Hatched bars, below it: the pump delivery cut to 85 % of command of scenario 8, on the cascade and on the single loop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

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

7.6.5 Rod Insertion at Four Pump Speeds

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

7.6.6 Fuel Pump Trip

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

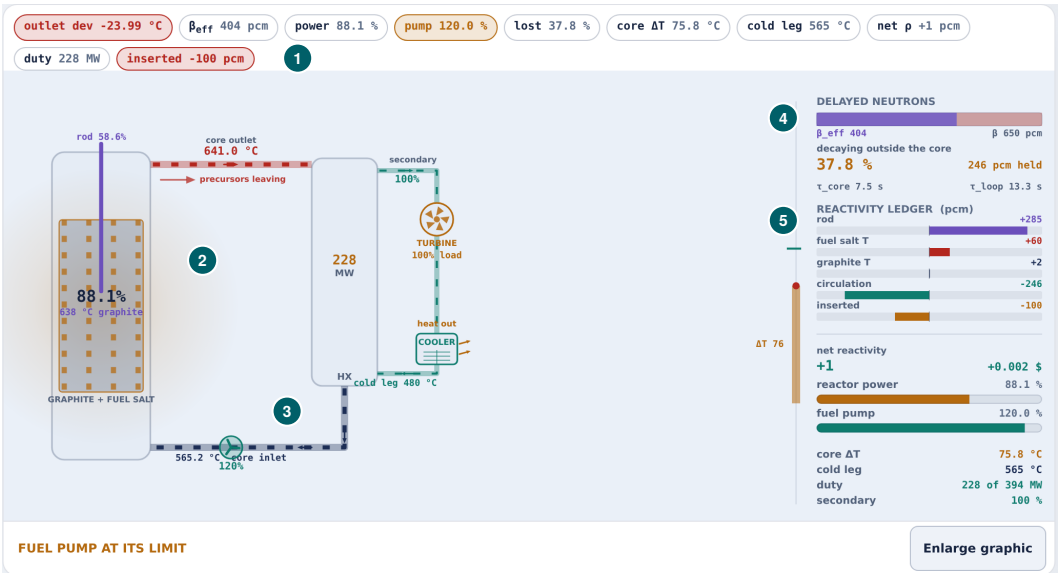

Figure 7.4 Recovery from a −100 pcm rod drop at the top of the pump range, in the plant view of the app. Callouts: (1) the status badges, (2) the core with its power and graphite temperature, (3) the fuel pump on the cold leg, (4) the delayed neutron panel and (5) the reactivity ledger. Scenario 5 runs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

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

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

7.6.7 Heat Exchanger Fouling and a Positive Graphite Coefficient

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

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

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

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

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

7.6.8 Graded Challenge

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

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

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

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

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

7.6.9 Predictive Control and the Margin Constraint

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

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

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

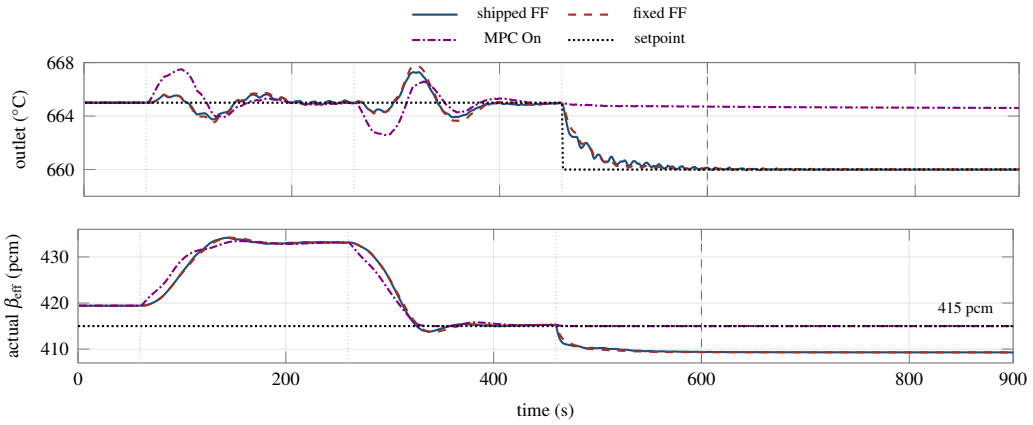

Figure 7.5 Outlet tracking (upper panel) and actual delayed-neutron margin (lower panel) over 900 s. Shipped FF is the solid line, fixed FF the dashed line and MPC On the dash-dotted line. The thick dotted line is the setpoint above and the 415 pcm margin below. Thin dotted event markers locate the received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

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

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

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

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

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

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

7.7 Discussion

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

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

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

7.8 Technology in Practice

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

What Measures It

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

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

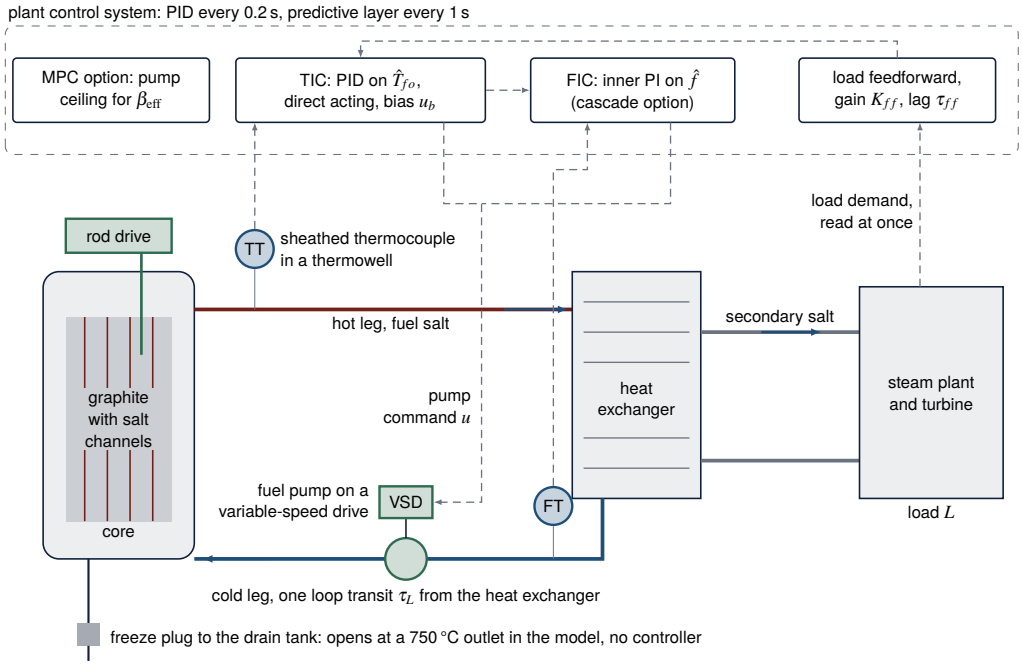

Figure 7.6 The temperature loop of the chapter with plant hardware in place of the symbols of Figure 7.1. Sensors are blue circles and final elements green: the core outlet thermocouple TT in its thermowell on the hot leg, the flow transmitter FT on the cold leg, the fuel pump on its variable-speed drive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

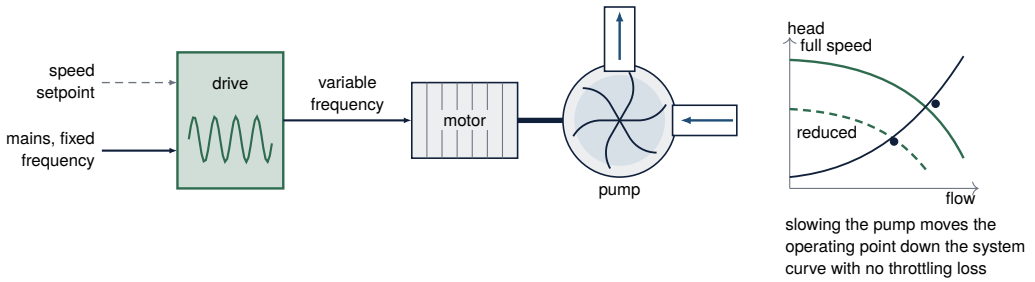

Figure 7.7 A variable-speed drive on a centrifugal pump, from the book's instrument library. Notice that the drive changes the pump curve, not the system curve, so the flow moves along the system curve as the speed changes. In this chapter the same speed also sets the time the fuel salt spends inside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

What Moves It

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

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

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

Where the Controller Runs

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

IN PRACTICE Load Changes Planned Around Thermal Fatigue

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

What Is Modernizing It

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

7.9 Limitations

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

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

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

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

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

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

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

7.10 Conclusion

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

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

KEY POINTS

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.11 Exercises

7.1 Concept check · 2 min

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

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

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

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

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

7.2 Concept check · 2 min

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

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

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

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

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

7.3 Concept check · 2 min

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

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

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

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

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

7.4 Theory and calculation · 5 min

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

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

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

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

7.5 Theory and calculation · 6 min

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

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

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

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

7.6 Theory and calculation · 6 min

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

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

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

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

7.7 App challenge · 8 min

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

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

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

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

7.8 App challenge · 8 min

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

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

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

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

7.9 App challenge · 8 min

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

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

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

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

7.10 Application area · 4 min

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

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

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

7.11 Application area · 4 min

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

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

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

7.12 Review the agent · 5 min

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

CHAPTER 8

Mission-Profile Load Following and Preview Control of an Aircraft Molten Salt Reactor

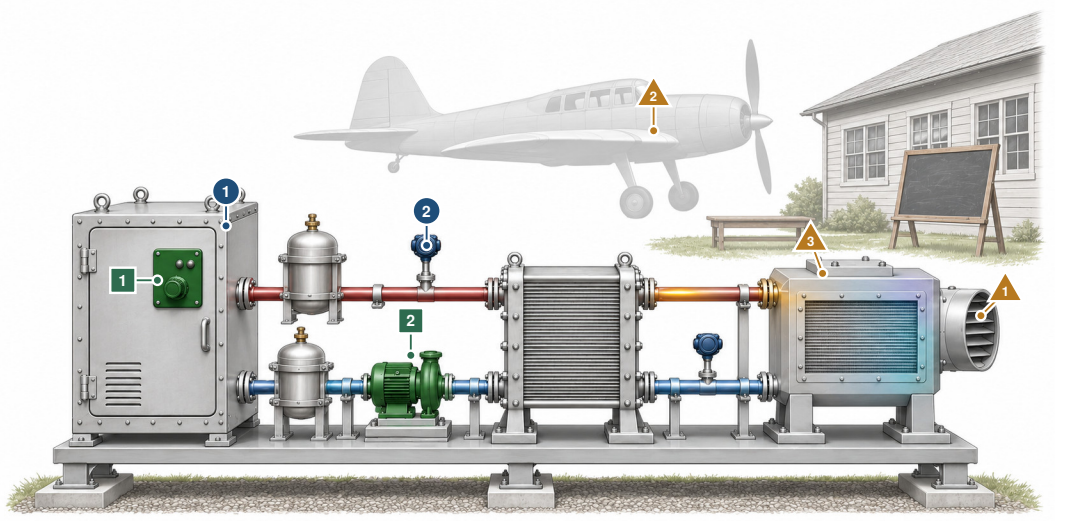

The controller moves 1 control rod at the reactor 2 fuel pump speed

The controller reads 1 reactor power, from a neutron detector 2 core outlet temperature

What disturbs it 1 engine throttle 2 altitude and flight speed 3 fouling of the air heater

Notice that the throttle acts at the far right, on the air heater, while the rod and the pump act at the far left, a whole loop of salt away. Illustration, not to scale.

IN PLAIN TERMS **A Reactor That Must Read the Flight Plan**

The main idea. This reactor heats the air of two jet engines in place of burning fuel. Its fuel is a hot liquid salt that a pump moves round a loop. A control rod and the pump together set the reactor's power and its temperature.

The challenge. At takeoff the engines ask for many times more heat within a minute. The reactor feels the change only after the cooled salt comes all the way round. The rod moves slowly, and the salt must never get hotter than its metal pipes allow.

Why it matters. A flight follows a plan, so a controller that reads the plan ahead can move before the demand does.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	The first page, the fact box of the Introduction, the history boxes of Section 8.2, the try-it and fact boxes of Section 8.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 8.6, Section 8.7, Section 8.8 and the open challenge

The control question A molten salt reactor that heats the air of two turbojets must follow a thermal demand that moves by a factor of fourteen in a takeoff, with a rate-limited rod and a fuel pump that both move power and outlet temperature. What is advance knowledge of the flight profile worth to it, and what does a coordinated predictive controller add on an event the profile does not contain?

Learning objectives

- Compute the demanded duty and the core temperature rise from the model, and explain why a circulating core is critical at positive reactivity.
- Calculate the relative gain of the rod and pump pairing from a measured gain matrix, and explain why it moves from 0.86 at takeoff to 3.58 at idle.
- Distinguish an actuator that has run out of process gain from one that has run out of travel.
- Measure what preview is worth to a feedforward and to a dynamic matrix controller on a scheduled takeoff.
- Explain why two independent loops break the outlet limit in a go-around that the predictive controller survives.

Prerequisites Sections 1.4.2 and 1.4.3 (point kinetics, circulating precursors), 1.6.2 (velocity-form PID), 1.6.3 (predictive control) and 1.6.4 (scenarios and data), and Chapter 7 for a circulating-fuel core with one pump loop.

The app https://apmonitor.com/apps/aircraft-msr/

See also Volume 6, Chapter 2 for pairing and decoupling with the relative gain array. Volume 9, Chapter 5 for feedforward from published orders, Volume 9, Chapter 4 for coordinated two-input two-output predictive control and Volume 9, Chapter 9 for a controlled variable held in a band.

8.1 Introduction

A turbojet raises the temperature of compressed air ahead of the turbine, and nothing in the thermodynamic cycle requires that the temperature rise come from combustion. Between 1946 and 1961 that observation funded the Aircraft Nuclear Propulsion program, whose reactor half produced the Aircraft Reactor Experiment (ARE) at Oak Ridge and whose engine half produced turbojets that ran on reactor heat at the Idaho test site [1, 2]. The program was cancelled for reasons of shielding mass and changing strategic priorities rather than for reasons of control, and the control problem it posed has not been retired. It has become more relevant, because every current proposal for a molten salt reactor serving a process, an isolated grid or a marine plant is a proposal for a reactor that follows a load it does not set [3, 4].

An aircraft is the extreme case of that duty. The thermal demand of the plant studied here moves by a factor of fourteen in forty-five seconds during a takeoff roll and by a factor of nine in twelve seconds on a go-around, against a control rod limited to 0.6 % of travel per second and a fuel salt whose containment alloy sets a hard upper temperature. No grid-connected reactor is asked for anything comparable. What distinguishes the problem from a large load disturbance, however, is not its size. It is that the disturbance is *published*. Throttle, altitude and Mach number twenty seconds in the future are known in the present, because a flight is flown to a plan, while the reactor's own measurement of the same disturbance arrives only after crossing an air heater metal mass, an intermediate salt loop, an exchanger metal mass, a 10 s fuel loop transit and a thermocouple. Figure 8.1 sets the two paths side by side.

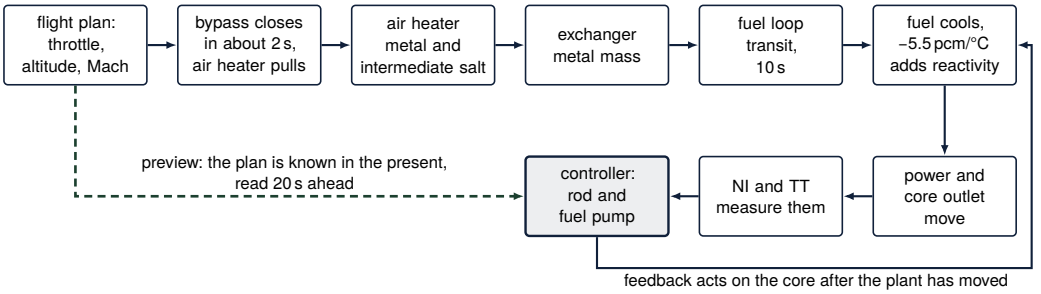

Figure 8.1 Two paths from the same throttle change. Solid arrows: the plant's own path, which crosses the air heater, the intermediate salt loop, the exchanger metal and the fuel loop transit before the fuel temperature coefficient moves the power, and is measured only after that. Dashed arrow: the preview path, by which the flight plan reaches the controller before the throttle moves. Notice that feedback can only start at the end of the solid path, while preview starts at its beginning.

> ✦ **DID YOU KNOW? Five Percent a Minute**
>
> Most current power reactor designs can maneuver between 50 and 100 percent of rated thermal power and ramp as fast as 5 percent of rated power per minute, although these maneuvers are not part of normal daily operation [5]. The takeoff of this chapter moves the demand by a factor of fourteen in forty-five seconds.

Simulation is the natural way to teach that distinction. A block diagram or a transfer function can express a feedforward path, but neither exposes what makes this plant hard, which is that the steady-state gain matrix moves by more than an order of magnitude across one flight and that the pairing chosen at the design point stops being the right pairing before the aircraft lands. The misconception the application is built to correct is that saturation means an undersized actuator. At 6.6 MW the core temperature rise of this plant is about 7 °C at design flow, so the fuel pump cannot move the outlet temperature 23 °C no matter how much travel it has. That is an actuator that has run out of process gain, it looks identical to an undersized valve on a trend, and it is fixed by a different remedy.

The class of problem is broader than nuclear propulsion. Any plant driven by a published trajectory has preview available and most control systems built for steady operation do not use it. Optimal preview tracking has an established theory for known reference models [6], and predictive control in receding-horizon form accepts a setpoint trajectory as naturally as a setpoint [7, 8]. On the plant of this chapter the value of preview can be measured against the same plant, the same actuators and the same limits as a controller that does not use it.

This chapter works with an interactive aircraft molten salt reactor and quantifies what advance knowledge of a flight profile is worth to a reactor that must follow it. Nothing in it is a claim that a nuclear aircraft is practical. Section 8.2 places the problem in its history. Section 8.3 develops a reduced-order model in which the aircraft enters as a scheduled disturbance through two thermodynamic quantities. Section 8.4 states the four control architectures as implemented, among them a two-input two-output dynamic matrix controller whose free response is built from a closed-form steady state rather than from an integrated model, which makes preview control affordable in a browser at about 20 µs of wall-clock time per execution on a desktop processor. Section 8.5 describes the app and its eleven scenarios. Section 8.6 compares the four architectures on one mission under actuator rate limits and a temperature constraint. Section 8.7 names the hardware such loops run on, Sections 8.8 and 8.9 give the limitations and the conclusion, and the exercises follow in Section 8.10.

8.2 Background

The ARE operated at Oak Ridge in November 1954, circulating a molten fluoride fuel salt and rejecting its heat through helium and water rather than through an engine [1, 9], and the History box

below tells how that circulation showed itself.

> ⧖ HISTORY **Watching the Fuel Go Round (1954)**
>
> The Aircraft Reactor Experiment at Oak Ridge pumped its molten fluoride fuel through the reactor and on to two heat exchangers that passed the heat to helium, and the helium gave it to water that was then discharged. At full flow the fuel took about 47 seconds to go round, about 8 of them in the core. The helium blower ran through a magnetic clutch, which gave smooth control at any power where the temperature coefficient was the controlling factor. The circulation showed on the fission chamber recorders as uranium was added from October 30, 1954. Each enriched slug mixed slowly into the carrier salt and made a pip every time it passed through the core, almost exactly 2 minutes apart at minimum pump speed and later 47 seconds apart, which gave the same flow as the flowmeter [10].

Two findings from that experiment shape the model of this chapter. The first is that circulating the fuel carries delayed neutron precursors out of the core before they decay, which modifies the inhour equation without destabilizing the reactor [11]. The second is that the experiment was operated largely by changing the heat removal and allowing the negative temperature coefficient to find the power, which is self-regulation used as an operating procedure rather than as a safety argument [9].

The engine half of the program ran turbojets on reactor heat in its direct-air-cycle Heat Transfer Reactor Experiments [2], as the second History box tells.

> ⧖ HISTORY **Turbojets on Reactor Heat (1956)**
>
> The engine half of the aircraft nuclear propulsion program built a series of Heat Transfer Reactor Experiments, in which compressor air was heated directly in the reactor. In January 1956 the first of them ran a modified turbojet on nuclear power alone. The engine was started on chemical fuel with its air passing through the cold reactor. The control rods were then withdrawn and the reactor brought to power by demanding more neutron flux, and as the reactor heated the air, the chemical fuel valve, sensing the temperature rise, gradually closed. Engine speed was held by the area of the exhaust nozzle. The third experiment ran two modified turbojets and reached full nuclear power in 1959 [12]. The plant of this chapter heats the air through an intermediate salt loop instead.

The arrangement in that work, in which compressed air is heated by the reactor and expanded through the turbine, is the arrangement modeled here, with an intermediate salt loop interposed so that the fuel salt never enters the engine.

Quantitative dynamic modeling of circulating-fuel reactors matured on the Molten Salt Reactor Experiment that followed. The lumped structure used in this chapter, with modified point kinetics, a small number of well-mixed thermal nodes and a pure transport delay around the external loop, is the structure of the MSRE theoretical dynamics analysis [13] and of the stability analysis that preceded it [14], and the operating experience against which such models were compared is documented [15]. The reduction of precursor transport to a single loop transit time with an exponential survival factor remains the standard point-kinetics treatment, and modern work has quantified how much it loses against transport and computational fluid dynamics solutions of the same problem [16, 17] and against Monte Carlo evaluations of the effective delayed fraction [18].

Control of molten salt reactors under load following is an active subject. Core outlet fuel temperature has been treated as the primary controlled variable with fuel circulation as the manipulated variable [3], and load following transients have been examined for natural circulation designs where the flow is not manipulable [4]. What the aircraft case adds to that literature is a second controlled variable that must be held simultaneously, a disturbance an order of magnitude faster, and a disturbance trajectory that is known in advance.

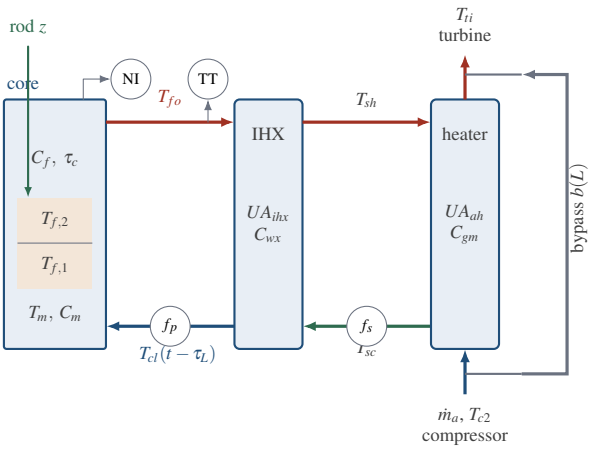

Figure 8.2 Aircraft molten salt reactor propulsion arrangement. Fuel salt circulates through two core nodes at temperatures $T_{f,1}$ and $T_{f,2}$ against a moderator and structure node T_m, leaves at T_{fo}, is cooled in the intermediate heat exchanger, and returns to the core inlet one external loop transit τ_L later. The intermediate salt carries the duty to an air heater, where compressed air at $\dot{m}_a$ and T_{c2} is heated and remixed with the fraction $b(L)$ routed around the heater. The control rod position z and the fuel pump speed f_p are the manipulated variables, and the neutron instrument NI and the outlet thermocouple TT supply the two measurements.

8.3 Aircraft Reactor Model

The plant is a notional molten salt reactor rated at 125 MW thermal driving two turbojets. Fission heats a circulating fluoride fuel salt directly. The salt is pumped through the core, gives its heat to a clean intermediate salt in an intermediate heat exchanger (IHX), and returns. The intermediate salt carries the heat to an air heater in the engine, where it passes to compressed air. Heated air and bypassed air remix ahead of the turbine. Figure 8.2 shows the arrangement and labels every symbol that appears in the equations below.

Table 8.1 classifies the variables the control problem is written in.

8.3.1 Circulating-Fuel Point Kinetics

The fuel is a liquid and it moves, so a delayed neutron precursor born in the core is carried out of it before decaying, as Section 1.4.3 explains and Chapter 7 develops for a ground plant. With a core transit time τ_c and an external loop transit τ_L, the six precursor groups obey

$$\frac{dW_i}{dt} = \beta_i n - \lambda_i W_i - \frac{W_i}{\tau_c} + \frac{W_i(t-\tau_L)\, e^{-\lambda_i \tau_L}}{\tau_c}, \tag{8.1}$$

in which $W_i = \Lambda c_i$ scales the concentration by the neutron generation time, n is power as a fraction of rated, and the last two terms are loss to the loop and delayed return reduced by decay along the way. The prompt neutron mode is removed algebraically by the prompt jump approximation [19] of Equation (1.1), with six groups, which gives

$$n = \frac{\sum_i \lambda_i W_i}{\beta - \rho}. \tag{8.2}$$

The generation time cancels from the model entirely, which is what allows a 0.05 s integration step to carry a thirty-minute flight. The β in the denominator of (8.2) is the *static* delayed fraction,

Table 8.1 Controlled, manipulated and disturbance variables

Symbol	Meaning	Units	Range
Controlled variables			
P	reactor thermal power	% of rated	0 to 150
T_{fo}	core fuel outlet temperature	°C	600 to 900
Manipulated variables			
z	control rod position	% withdrawn	0 to 100
f_p	fuel pump speed	% of design	40 to 125
Disturbance variables			
L	propulsion demand (throttle)	%	10 to 105
h	altitude	ft	0 to 40 000
M	Mach number	—	0 to 0.80
Principal states			
W_i	scaled precursors, $i = 1$ to 6	MW	—
$T_{f,1}, T_{f,2}$	core fuel node means	°C	—
T_m	moderator and structure	°C	—
T_{cl}	fuel cold leg plenum	°C	—
T_{wx}, T_{gm}	exchanger metal masses	°C	—
T_{sh}, T_{sc}	intermediate hot and cold	°C	—
T_{ti}	turbine inlet gas	°C	—
$\dot{m}_a$	engine air mass flow	kg/s	30 to 245
b	air heater bypass fraction	—	0 to 0.92

650 pcm for the six-group thermal uranium-235 data used here, because it comes from production and production does not depend on where the precursor ends up.

Setting dW_i/dt to zero in (8.1) with $W_i(t-\tau_L) = W_i$ gives the effective delayed fraction that the fuel pump controls, which is Equation (1.5) of Chapter 1 with the pump speed written in,

$$\beta_{\text{eff}} = \sum_i \beta_i \frac{\lambda_i \tau_c}{\lambda_i \tau_c + 1 - e^{-\lambda_i \tau_L}}, \qquad \tau_c = \frac{\tau_{c,0}}{f_p}, \quad \tau_L = \frac{\tau_{L,0}}{f_p}. \tag{8.3}$$

At the design flow of 6 s in the core and 10 s outside it, (8.3) gives 390 pcm against a static 650 pcm, so 40.1 % of the delayed neutrons decay outside the core and the rod must hold 260 pcm of permanent excess reactivity to pay for circulation. Applying the same expression to the published ARE transit times of 8.3 s and 47 s gives 378 pcm and a loss of 41.8 %. Because both transit times scale inversely with pump speed, β_{eff} runs from 464 pcm at 40 % flow to 373 pcm at 125 %, and the fuel pump is therefore an actuator that moves the plant's own stability margin every time the temperature loop acts.

8.3.2 Reactivity and Criticality

Reactivity is measured from a design reference state at rated power,

$$\rho = \rho_0 + W(z) - W(z_0) + \alpha_f\left(\bar{T}_f - T_{f,0}\right) + \alpha_m(T_m - T_{m,0}), \tag{8.4}$$

with $\bar{T}_f = (T_{f,1} + T_{f,2})/2$ and $W(z)$ the usual S-curve integral rod worth. A circulating core is critical at *positive* reactivity, because it has to replace the delayed neutrons it is pumping away (Section 1.4.3), so $\rho_0 = \beta - \beta_{\text{eff}}$ at design flow and the steady-state condition is $\rho_{ss} = \beta - \beta_{\text{eff}}$ at whatever flow is in

force. Equation (8.4) is the reactivity ledger of Equation (1.3) without a xenon term, and it pins a weighted core temperature at every power, which is why the mean fuel temperature barely moves across the mission and what a load change actually moves is the core temperature *rise*.

Pause. Try Exercise 8.1 before reading on.

8.3.3 Core and Loop Energy Balances

The core is divided into two fuel nodes in series, which is the coarsest division that separates the mean fuel temperature that criticality pins from the outlet temperature that the containment alloy sees,

$$\begin{aligned}\frac{C_f}{2}\dot{T}_{f,1} &= \tfrac{1}{2}(1-f_M)Q - 2W_f\left(T_{f,1}-T_{fi}\right)\\ &\quad - \tfrac{UA_{fm}}{2}\left(T_{f,1}-T_m\right), \qquad (8.5)\\ \frac{C_f}{2}\dot{T}_{f,2} &= \tfrac{1}{2}(1-f_M)Q - 2W_f\left(T_{f,2}-T_{f,1}^{out}\right)\\ &\quad - \tfrac{UA_{fm}}{2}\left(T_{f,2}-T_m\right), \qquad (8.6)\\ C_m\dot{T}_m &= f_M Q + \tfrac{UA_{fm}}{2}\left(T_{f,1}+T_{f,2}-2T_m\right), \qquad (8.7)\end{aligned}$$

with $Q = P_0 n$ and $W_f = C_f/\tau_c$ the flow capacity rate. Each flowing node carries a linear temperature profile, so its mean is the average of its inlet and outlet and the outlet follows algebraically as $T^{out} = 2T - T^{in}$. The core inlet is the cold leg plenum temperature delayed by one external loop transit, $T_{fi}(t) = T_{cl}(t-\tau_L)$, and the same τ_L that carries the heat carries the precursors in (8.1), because it is the same salt in the same pipe.

8.3.4 Heat Exchange

Both exchangers use the counterflow effectiveness relation rather than a stream exchanging with an isothermal wall. The distinction is not cosmetic on this plant, because the intermediate salt leaves the IHX at 759 °C while the fuel leaves it at 659 °C, and a lumped wall model cannot represent a cold outlet hotter than the hot outlet. With $C_{\min}$ and $C_{\max}$ the smaller and larger capacity rates and $C_r = C_{\min}/C_{\max}$,

$$\varepsilon = \frac{1-e^{-N(1-C_r)}}{1-C_r e^{-N(1-C_r)}}, \quad N = \frac{UA}{C_{\min}}, \quad Q = \varepsilon C_{\min}\Delta T_{in}. \qquad (8.8)$$

Both conductances follow the usual turbulent 0.8 power of mass flow, split as 0.4 per stream, so $UA = UA_0(W_1/W_{1,0})^{0.4}(W_2/W_{2,0})^{0.4}$. The metal mass of each exchanger is carried as a separate capacity that absorbs $UA_m(\bar{T} - T_{metal})$ during a transient and contributes nothing at steady state, which gives the exchanger thermal inertia without also giving it a thermal resistance it does not have.

8.3.5 Aircraft and Engine

The aircraft enters the reactor model through exactly two thermodynamic quantities. International standard atmosphere and ram compression give the total conditions at the compressor face, and an ideal compressor with fixed pressure ratio π_c and isentropic efficiency η_c gives the discharge temperature,

$$T_{t1} = T_{amb}(h)\left(1+0.2M^2\right), \quad T_{c2} = T_{t1}\left[1+\frac{\pi_c^{0.2857}-1}{\eta_c}\right]. \qquad (8.9)$$

Engine air flow follows the corrected flow group and a compressor map written as a power of spool speed,

$$\dot{m}_a = \dot{m}_{max}\, \nu(L)^{2.2}\, \frac{\delta}{\sqrt{\theta}}, \qquad \nu(L) = 0.35 + 0.0065\, L, \tag{8.10}$$

with δ and θ the total pressure and temperature ratios to sea level static. The group $\delta/\sqrt{\theta}$ is 1.04 at sea level and 0.38 at 35 000 ft, while T_{c2} is 248 °C at sea level and 165 °C in the cruise. The two act against each other, so takeoff and high-altitude cruise are physically different disturbances at the same throttle.

Propulsion demand is expressed as the turbine inlet temperature the engine schedules, because thrust at a given air flow is set by how hot the gas is when it reaches the turbine. The demanded reactor duty follows,

$$Q_{dem} = \dot{m}_a c_{p,a} \left(T_{ti}^{SP}(L) - T_{c2} \right), \qquad b = 1 - \left(\tfrac{L}{100} \right)^{\kappa}. \tag{8.11}$$

The second half of (8.11) is the air heater bypass, the fraction of compressor discharge routed around the heater and remixed downstream. It is scheduled open loop on throttle, so the reactor still has to make the heat, and it exists because without a cold air path this plant has no turndown at all. A fixed-geometry air heater with the core held at its design temperature pulls 26 MW at ground idle, 21 % of rated against a demand of 7 %, and the only way to shed the difference would be to cool the salt to 409 °C, which freezes it.

8.3.6 Steady State in Closed Form

Every link in the chain is a fixed thermal resistance at fixed flows, so the whole steady state collapses to two affine relations,

$$\bar{T}_f = T_{c2} + A\, Q, \qquad T_m = T_{c2} + \left(A + \tfrac{f_M}{U A_{fm}} \right) Q, \tag{8.12}$$

with A assembled from the reciprocal conductances of the air heater, the intermediate loop, the IHX and the core. Substituting (8.12) into the criticality condition $\rho_{ss} = \beta - \beta_{\text{eff}}$ leaves one linear equation in Q, so the settled power and every settled temperature at a given rod position, pump speed and flight condition are a closed form and never have to be integrated. That property is what makes both the model feedforward of Section 8.4.4 and the predictive controller's free response of Section 8.4.5 affordable in a browser.

8.3.7 Parameters

Table 8.2 lists the parameters. The column marked *ARE* gives the corresponding published value where one exists, so that a reader can see exactly where the model stops being anchored to an experiment and becomes notional.

8.3.8 Assumptions

The model is a control model. Point kinetics with the prompt jump approximation is silent about anything approaching prompt critical, and the application annunciates $\rho \geq \beta$ rather than integrating through it. The precursors are assumed uniformly mixed in the core, so a precursor about to leave is treated as worth the same as one just born near the flux peak, which transport solutions of the same problem show is an approximation rather than an identity [16, 17]. A pump speed change alters both transit times instantly, so salt already in the loop is treated as having always traveled at the new speed. Reactivity coefficients are constants, there is no burnup, and xenon is not modeled.

On the thermal side there are two fuel nodes, one moderator node and one lumped metal mass per exchanger, so no spatial temperature or flux distribution exists and hot channel factors cannot appear. The whole external transport is applied to the return leg rather than split between a hot leg and a cold leg. The 850 °C limit, the 900 °C trip and the 500 °C freeze alarm are simulation constraints chosen to make the control problem bite rather than the qualification limits of any particular alloy or salt.

Table 8.2 Model parameters, with the published Aircraft Reactor Experiment value where one exists

Symbol	Meaning	ARE	Model	Rationale for the difference
P_0	rated thermal power, MW	2.5	125	two turbojets rather than a helium and water sink; results are reported in percent
$\tau_{c,0}$	core transit at design flow, s	8.3	6.0	compact core sized for an airframe rather than a test loop
$\tau_{L,0}$	external loop transit, s	47	10.0	one exchanger in a nacelle rather than eight in a test cell
β	static delayed fraction, pcm	470	650	six-group thermal uranium-235 data, so β_{eff} can be checked by hand
α_f	fuel temperature coefficient, pcm/°C	−5.5	−5.5	taken directly
α_m	moderator coefficient, pcm/°C	—	−1.0	small and negative by design
f_M	fission energy outside the fuel	0.108	0.108	taken directly
v_{rod}	rod speed limit, %/s	—	0.60	10.5 pcm/s at the design position, the order of the ARE regulating rod
W_{tot}	total rod worth, pcm	—	900	260 pcm is spent permanently on paying for circulation
T_{fo}	core outlet at rated power, °C	828	799	alloy limit set at 850 °C to leave transient margin
T_{fi}	core inlet at rated power, °C	654	659	comparable core temperature rise
C_f	core fuel heat capacity, MW s/°C	—	5.36	gives a 3 s core thermal time constant
C_m	moderator and structure, MW s/°C	—	24	gives a 40 s time constant, the slowest thermal mode
UA_{ihx}	IHX conductance, MW/°C	—	2.53	sized for 125 MW at the design terminal temperatures
UA_{ah}	air heater conductance, MW/°C	—	0.758	sized for 699 °C turbine inlet at takeoff
$\dot{m}_{max}$	engine air, sea level static, kg/s	—	232	low turbine inlet temperature forces a large air flow
π_c	compressor pressure ratio	—	6.0	low on purpose, because a cool discharge leaves room for the reactor
κ	bypass schedule exponent	—	0.67	fits the demanded duty at sea level and is deliberately wrong at altitude

On the aircraft side there is no flight model at all. No lift, no drag, no weight and no equation of motion, so nothing here can say whether the aircraft would climb at the rate the profile asserts. Thrust appears only as the turbine inlet temperature the engine did or did not receive. There is no shielding and no mass accounting, which means the reason the historical program was cancelled is entirely absent from the model.

8.4 Control Architectures

8.4.1 Pairing

Both actuators move both controlled variables. Withdrawing the rod raises the core temperature level, which raises the duty a fixed heat sink pulls, so the rod raises power and outlet temperature together. Speeding the fuel pump narrows the core temperature rise and cools the outlet, and raises the power slightly because a faster flow lowers the thermal resistance A in (8.12). The steady-state gains measured on the application at the sea-level takeoff condition are

$$\begin{bmatrix} \Delta P \\ \Delta T_{fo} \end{bmatrix} = \begin{bmatrix} +0.549 & +0.020 \\ +3.03 & -0.69 \end{bmatrix} \begin{bmatrix} \Delta z \\ \Delta f_p \end{bmatrix}, \tag{8.13}$$

in units of percent of rated power and degrees Celsius per percent of actuator. The relative gain array [20] of the pairing rod to power and pump to temperature is therefore 0.86, which is close to unity and says that the two loops barely interact at the design point. Volume 6, Chapter 2 treats a pairing whose relative gain of one still hides a one-way interaction.

That is not a property of the plant. Repeating the same four measurements at landing idle gives $\partial P/\partial z = 0.023$, a factor of twenty-four weaker, and $\partial T_{fo}/\partial f_p = -0.15$, a factor of four and a half weaker, while $\partial T_{fo}/\partial z$ is between 2.3 and 3.1 °C per percent at every steady condition on the mission except the top of the descent, where the rod sits at 73 % on the flat of its worth curve and the figure is 1.5. The relative gain is 3.58 there. The pairing chosen at the design point is strongly interacting by the time the aircraft is on the ground, and at low power the rod is the better temperature handle while the pump is barely a handle at all.

The intermediate salt circulator is instructive as a negative result. With the rod parked, moving it from 40 % to 130 % of design at cruise changes the settled reactor power by 0.2 points and the core outlet by 0.1 °C, because the air stream through the heater (0.04 MW/°C) is then the smaller capacity rate by a factor of more than seven even at 40 % circulation, and a circulator that only sets the spread of the intermediate loop leaves the duty where it was. The same move at the sea-level design point changes the power by 19.7 points, because at full throttle the heater air stream (0.27 MW/°C) and the intermediate stream at 40 % flow (0.31 MW/°C) are nearly balanced. It is a pump, it costs mass, it is a handle only where the rod is already strong, and it is not a manipulated variable.

8.4.2 Temperature Program

No fixed core outlet setpoint exists that both ends of the mission can reach. With the power loop closed, the fuel pump can place the outlet anywhere between 779 and 950 °C at the top of the takeoff climb and only between 740 and 752 °C at taxi idle, and those bands do not overlap. The setpoint is therefore a program that slides with the demanded power,

$$T_{fo}^{SP} = T_{prog} - k_T\,(100 - P_{SP})\,, \tag{8.14}$$

with $T_{prog} = 795\,^{\circ}\mathrm{C}$ and $k_T = 0.55\,^{\circ}\mathrm{C}$ per percent of power as shipped. Setting k_T to zero restores a fixed setpoint and the pump then spends most of the flight on a stop, which is the subject of scenario 6 in Table 8.3.

8.4.3 Decentralized Feedback

Both loops use the shared velocity-form PID of Section 1.6.2, with the derivative taken on the measurement, executed every $\Delta t_c = 0.25\,\mathrm{s}$ against a 0.05 s integration step. Equation (8.15) writes it in this chapter's symbols, with s_k for the running sum S_k of Equation (1.14) and c_k for the output v_k before limiting of Equation (1.15),

$$\begin{aligned} s_k &= s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\ c_k &= u_{bias,k} + s_k + I_{k-1} \\ u_k &= \mathrm{slew}[\mathrm{clamp}(c_k, u_{\min}, u_{\max})\,, u_{k-1}, v\,\Delta t_c] \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - c_k}{\tau_t}\right]\Delta t_c \end{aligned} \tag{8.15}$$

with $P_k = K_c e_k$, $D_k = -K_c \tau_D\, a\,(\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f)/\Delta t_c$ and $e_k = a(\mathrm{SP}_k - \mathrm{PV}_k)$, where a is the action flag and PV^f is the measurement passed through a first-order filter of 1 s. The output u_k is c_k clipped to the actuator range and then limited to a move of $v\,\Delta t_c$ per scan, and the integral alone absorbs the difference by back-calculation with a tracking time $\tau_t = \tau_I$. Setting τ_I to the top of its range (600 s) removes integral action and with it the back-calculation.

The power loop has a positive process gain, so $a = +1$ and the loop is reverse acting in the convention of Section 1.3. The temperature loop has a negative process gain, because a faster pump cools the outlet, so $a = -1$ and the loop is direct acting. The power loop's slew rate v is the rod drive rate of 0.6 %/s, so a demand the drive cannot execute is never counted as executed, and the temperature loop has no slew limit. The power loop's bias is the rod position at the start of the run, or the feedforward z_{ff} of Section 8.4.4 when that architecture is selected, and the temperature loop's

bias is the 100 % design pump speed. In manual operation and under the predictive controller both loops track the applied actuator positions, so every transfer is bumpless.

8.4.4 Preview Feedforward

The feedforward is solved from the steady state of the plant as declared, with the parameters of Table 8.2 at their defaults, rather than identified from a step test, because the operating point moves by a factor of fourteen in power across one flight and no single identified model covers it. Given the flight condition at $t + t_{lead}$, the demanded duty from (8.11) and the closed form (8.12), one bisection on the monotone rod worth curve returns the rod position that makes the plant critical at that duty. The feedforward is

$$z_{ff}(t) = \frac{1}{\tau_{ff} s + 1} z_{ss}(t + t_{lead}) + K_{ff} \left[P_{SP}(t + t_{lead}) - P_{SP}(t) \right]. \tag{8.16}$$

Two properties of (8.16) are worth stating because both were measured rather than assumed. The first is that the steady-state term contributes almost nothing to a takeoff. The rod position the plant needs is 54.1 % at taxi and 53.5 % at the top of the takeoff climb, so the destination barely moves and the entire job is transient. What a takeoff needs is the rod above its steady position for the length of the ramp, while the fuel, the moderator and two exchanger metal masses are charged to their new temperatures, and the size of that lead is available only from the schedule. That is the second term. Volume 9, Chapter 5 uses published orders the same way, released one travel time ahead at each gate.

The second property is that the steady-state term has to be lagged, and the lag is not a refinement. At low power the rod position this plant needs goes *up* as the throttle comes *down*, because the bypass over-bypasses and a smaller air stream needs a hotter core to drive even a small duty through it. Closing the throttle therefore asks the rod to move one way eventually and the opposite way immediately, and an unlagged steady-state feedforward gets the immediate half exactly backwards. Removing the 45 s lag takes the integrated power error over the takeoff window from 448 to 926 % s.

The feedforward enters the power loop as its bias $u_{bias,k}$ in (8.15), and the fuel pump is left on feedback. A pump feedforward solves for the flow that would hold the outlet on its program at the demanded duty, which during a descent is a steady state the plant will not reach for another minute. Placing that flow in the temperature loop's bias at the same 20 s preview, a variant tested and not shipped, drives the pump from 40 % to 110 % during the descent while the plant is still at the old power. Over the full flight it quadruples the pump travel (481 % against 117 %) and raises the outlet error from 8475 to 9270 °C s, while the power error moves only from 2817 to 2728 % s.

8.4.5 Dynamic Matrix Control

The predictive architecture is a two-input two-output dynamic matrix controller [7], not the quadratic-program controller that most apps offer (Section 1.6.3). It works over a 120 s horizon in twenty-four blocks of 5 s, with four moves per input. The objective is

$$J = \sum_{k=1}^{N_p} \left[w_P e_{P,k}^2 + w_T e_{T,k}^2 \right] + \sum_{j=1}^{N_m} \left[\lambda_z \Delta z_j^2 + \lambda_f \Delta f_{p,j}^2 \right], \tag{8.17}$$

and the unconstrained solution is the single matrix-vector product

$$\Delta u = \left(A^T Q A + \Lambda \right)^{-1} A^T Q e = K e, \tag{8.18}$$

after which the first move of each input is clamped into its actuator box and into the rod rate limit. The gain K is 8×48, eight moves against twenty-four predictions of each output, and is computed when the predictive controller takes over by Gauss-Jordan elimination on an 8×8 system. It is rebuilt

in place when the demand or the power has moved more than eight points of rated, which occurs 22 times over the full flight and ten times during the takeoff alone. Move suppression is normalized against the diagonal of $A^T QA$, so a weight of unity means a move costs as much as the error it removes, without which a value reasonable at takeoff freezes the controller at idle.

Two departures from the textbook form are load-bearing. The first is the free response. Textbook DMC shifts a stored prediction and corrects it with a step bias, which assumes everything unmodeled is a step. On this plant it is a ramp, because the throttle is moving, the bypass valve is moving with it and the reactor's own temperature feedback is dragging the power after them at two points per second, and a step-disturbance prediction starts a takeoff behind and never recovers. The free response here is instead built from the closed form (8.12) evaluated at the previewed flight condition with the actuators where they are,

$$\hat{y}_k = y + (y_k^{ss} + \text{bias} - y)\left(1 - e^{-kT_s/\tau_{free}}\right), \tag{8.19}$$

with $T_s = 5\,\text{s}$ the execution period and τ_{free} equal to 26 s for power and 40 s for the outlet. Each execution therefore costs twenty-five algebraic solves and no differential equations. The measured bias is carried as a level, which is what makes the controller offset free when the air heater fouls. The step responses and the free response use the parameters of Table 8.2 at their defaults, fixed when the application loads, and the prompt jump in the rod-to-power response is scaled by the measured power, so a changed parameter or the fouling reaches the controller only as model mismatch through that bias.

The second departure is that the outlet temperature is a zone of $\pm 12\,^{\circ}\text{C}$ rather than a setpoint, which is standard industrial practice for a variable that is wanted in a range [8] and the way Volume 9, Chapter 9 treats a comfort band. It matters here because the fuel pump saturates for about a third of the flight, so a target the pump cannot reach is a constant term in (8.17) that the optimizer can serve only with the rod, which moves the outlet by 2.3 to 3.1 °C per percent across the mission. At the shipped weights the optimizer barely does so. Closing the zone to zero moves the full-flight power error from 941 to 868 % s, an improvement rather than a cost, while the thrust deficit rises from 10 350 to 11 543 °C s.

The whole execution, free response and matrix-vector product together, is timed by the application whenever the predictive controller takes over and written to the exported file. It costs about 20 µs of wall-clock time on a desktop processor, once every 5 s of plant time.

8.5 Interactive Application

The app runs on the framework of Section 1.6. Twenty states are integrated with fixed-step fourth-order Runge-Kutta at 0.05 s, with the controller executing at 0.25 s and the predictive controller at 5 s. The external loop delay is a ring buffer of 8192 samples pushed every 0.05 s and interpolated linearly, which holds 409 s of history and covers the natural convection floor where τ_L reaches 200 s. Simulation speed is selectable from real time to fifty times, and the thirty-minute flight is therefore about half a minute of wall clock at the fastest setting.

Figure 8.3 shows the interface. Three tabs present the same simulation from three points of view. The mission tab draws the aircraft on its own scheduled altitude trace with the engine beneath it as compressor, air heater and turbine with the bypass duct looping around the heater. The plant tab draws the energy path end to end with a temperature ladder of every terminal temperature and the 850 °C limit across it. The neutronics tab draws the static delayed fraction as one bar split at β_{eff}, both transit times, a plot of β_{eff} against pump speed with the operating point marked, and the reactivity ledger as signed bars in pcm.

Eleven scenarios are provided, listed in Table 8.3 with what each establishes and what completes it. Each pins a flight phase, an architecture, a tuning and a disturbance schedule, and carries the briefing, Learn card and completion rule of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt. A finished run exports as a comma



Figure 8.3 The application in scenario 1 at run time 158 s, mission time 218 s, in the initial climb at 2700 ft and Mach 0.29, with both actuators parked, cropped to the mission tab. Power has overshot to 149.4 % of rated against a demand of 93.5 % and the high flux alarm is annunciated (5), while the core outlet is at 818 °C, 32 °C below the alloy limit. The status badges (1) lead with the power deviation. The mission tab shows the aircraft on the scheduled altitude profile (2), the phase name and flight condition, and the engine drawn as compressor, air heater and turbine with the bypass duct (3). The panel leads with the turbine inlet temperature (718 °C against a schedule of 699 °C) (4), which is the quantity that says whether the aircraft received the thrust it asked for.

separated file carrying the flight condition, both loops, every loop temperature, the reactivity ledger, both transit times, the accumulated metrics, every disturbance setting and the parameters the run used, so a single file reproduces the experiment without the application.

8.6 Results

Every number in this section was measured on the application at the tuning stated, with each case selected, its settings changed and then run. Integrated absolute error accumulates from the start of the run and never stops rising, so each figure is read at the end of the run its scenario defines and two figures are comparable only over the same window. Two windows are used throughout, a 300 s takeoff window that opens 30 s before the throttle moves (the case of scenarios 7 and 10), and the full 1800 s flight of scenario 8. Integrated errors are the scorecard's own, accumulated from the 4 Hz log, while peak power errors are taken against the continuous demand at the 0.05 s integration step. Peak deviations, peak temperatures and actuator travel do not have the window problem.

8.6.1 Inherent Load Following and Its Limits

With the rod and the fuel pump parked and nothing in automatic, the reactor follows the takeoff on its own and then overshoots it, the self-regulation that Chapter 2 studies in a solid-fuel core. Power reaches 149.5 % of rated against a peak demand of 100.1 % about fifty seconds after the throttle reaches full, 56.0 points above the demand at that moment. The largest gap between power and demand over the flight is 90.9 points the other way, at the end of the ramp, when the power has barely left its taxi value. The high flux alarm at 118 % annunciates on the way past.

Table 8.3 Scenarios and Completion Requirements

#	Scenario, *kind*	What it establishes	Required to complete
1	Nobody driving *demonstration*	inherent following, and its overshoot	Fly the mission in manual with the rod and the fuel pump untouched to $t = 4$ min, through the takeoff.
2	Hold and measure *experiment*	the gain matrix at two operating points	At the held sea-level, full-throttle condition in manual, step the rod up 2 % (from 54.2 to about 56.2 %) after $t = 5$ s and hold it 4 min, then step the pump up 10 % and hold it 4 min.
3	Takeoff by hand *experiment*	the rod rate limit	In manual, start moving the rod between $t = 1$ s and the end of the throttle ramp at $t = 75$ s, fly to $t = 2$ min with the outlet below 850 °C, then set the rod to the steady position the model gives for the demand.
4	Power loop *experiment*	one loop, and the temperature it ignores	Fly the whole mission to landing in automatic on the power loop alone, with the fuel pump left at 100 %.
5	Two loops *experiment*	the second loop and what it costs	Fly the mission in automatic on two loops until a flight phase pins the fuel pump on a stop, then on to landing.
6	Actuator authority *demonstration*	saturation without an undersized actuator	At ground idle on two loops with the program slope $k_T = 0$, run 5 min with the pump on a stop, then 5 min with $k_T \geq 0.5$ °C per percent. Then set the throttle to 100 %, Reset so the run starts lined out there, and run 5 min without a trip.
7	Preview feedforward *experiment*	lead, lag and anticipation gain	Four takeoff runs on preview feedforward to $t = 5$ min, each setting held throughout: as shipped, $t_{lead} = 0$, $t_{lead} = 45$ s, and $K_{ff} = 0$ at $t_{lead} = 20$ s.
8	Whole flight compared *experiment*	four architectures, no single winner	Four whole missions to landing in automatic, one on each architecture: power loop alone, two loops, preview feedforward and MPC.
9	Go-around *experiment*	the constraint, crossed in a transient	Three runs through the go-around at $t = 100$ s, each held on its architecture for 70 s between $t = 90$ and 165 s: two loops, preview feedforward and MPC.
10	Coordinated MPC *experiment*	the value of the preview switch	Three takeoff runs on MPC to $t = 5$ min: preview on with the ± 12 °C zone (IAE$_P$ below 1500 % s), preview off (above 3000 % s), and preview on with the zone at zero (below 1500 % s).
11	Graded challenge *challenge*	degradations a feedforward cannot see	Fly the whole 1800 s mission with power within ± 1.2 % of demand from $t = 1188$ s, the outlet never above 850 °C and total rod travel under 150 %, with the model parameters and disturbances as given.[a]

[a] With the shipped two-loop tuning ($K_{c,P} = 0.35$, $\tau_{I,P} = 150$ s, $K_{c,T} = 0.5$, $\tau_{I,T} = 180$ s) the worst power error over the judged window is 5.07 %, more than four times the band, although the outlet peaks at only 801 °C and rod travel is 112 %. With the suggested MPC architecture it is 0.92 %, inside the band with 0.28 points to spare, with a peak outlet of 818 °C and 99 % of rod travel.

The mechanism is a delay acting on a stiff feedback. The bypass valve closes in about two seconds and the air heater begins pulling at once, but the cooled salt must cross the air heater metal, the intermediate loop, the exchanger metal and a 10 s fuel loop transit before the core is affected. By the time the fuel temperature has fallen far enough for the −5.5 pcm/°C coefficient to supply the reactivity the ramp needed, the ramp is over and the reactivity is still present. The cost to the aircraft is 48 °C of turbine inlet temperature below schedule at the worst point, integrating to 4383 °C s by landing.

Table 8.4 Measured mission performance. IAE_P in percent-seconds, IAE_T in °C-seconds, thrust deficit is the integral of turbine inlet temperature below schedule in °C-seconds, and peak $|e_P|$ is the largest gap between power and the continuous demand at the 0.05 s integration step

	Takeoff window, 300 s				Full flight, 1800 s			
Architecture	IAE_P	peak $\lvert e_P\rvert$	peak T_{fo}	deficit	IAE_P	IAE_T	peak T_{fo}	deficit
Nobody driving	6235	90.9	819	1781	9990	7723	819	4376
Power loop only	687	10.93	817	64	1312	26525	817	5974
Two loops	641	10.41	810	65	1487	10667	812	7822
Preview feedforward	448	14.8	808	57	2817	8475	809	6791
Dynamic matrix control	729	15.6	818	24	941	20144	818	10350
DMC, preview disabled	6358	99.9	782	6231	10276	32259	786	31385

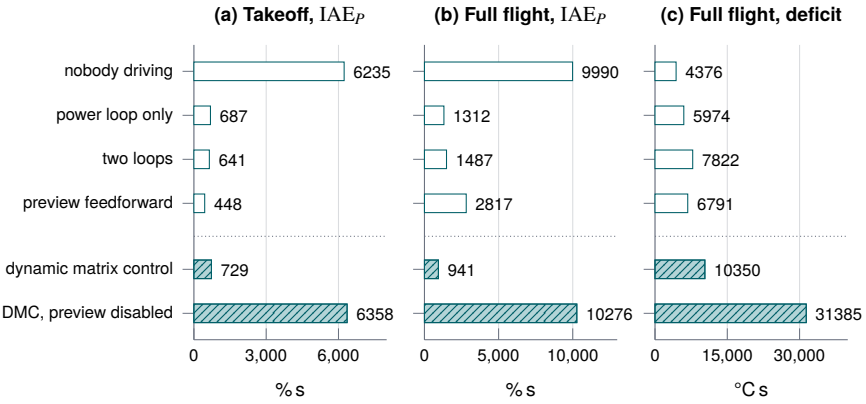

Figure 8.4 Three columns of Table 8.4 drawn as bars. White bars, above the dotted line: the plant with nobody driving and the three architectures without the predictive controller. Hatched bars, below it: the dynamic matrix controller with its preview and with the preview disabled. Notice that the same predictive controller with its preview disabled lands beside the plant with nobody driving in panels (a) and (b), and that the shortest bar belongs to a different row in each panel.

> **TRY IT IN THE APP** **1 · Fly the Mission With Nobody Driving**
>
> Load scenario 1 and press Run, leaving the control rod and the fuel pump alone. Watch reactor power against the demand trace through the takeoff. The power lags at first, then climbs on its own past the demand to 149.5 % of rated, and the high flux alarm at 118 % annunciates on the way past.

8.6.2 Architecture Comparison

Table 8.4 reports the four architectures on both windows at the shipped tuning of $K_{c,P} = 0.9$, $\tau_{I,P} = 45\,\mathrm{s}$, $K_{c,T} = 1.1$ and $\tau_{I,T} = 60\,\mathrm{s}$, with the feedforward at $t_{lead} = 20\,\mathrm{s}$, $K_{ff} = 0.20$ and $\tau_{ff} = 45\,\mathrm{s}$, and the predictive controller at $w_P = 4.0$, $w_T = 0.04$, $\lambda_z = 0.15$, $\lambda_f = 0.30$ and a $\pm 12\,^{\circ}\mathrm{C}$ zone. With the power loop alone the fuel pump stays parked at 100 %.

The takeoff power peaks are read against the engine's step-level demand at each physics scan. Readings against the continuous demand trace give 11.1 and 10.6 points for the first two controlled rows. Figure 8.4 draws three of the columns as bars.

No architecture wins every column, and the rows that appear to lose are the informative ones. The power loop alone removes most of the tracking error, taking IAE_P over the flight from 9990 to

1312 % s and the worst gap from 90.9 points to 10.93, and controls no temperature at all, integrating 26 525 °C s of deviation from the program. Adding the second loop cuts that to 10 667 °C s and costs tracking (1487 % s), because the pump is now moving and every pump move changes the thermal resistance of the whole chain.

Preview feedforward wins the takeoff and loses the flight. Over the 300 s window it integrates 448 % s against 641 for the two loops alone, and over the full flight it integrates 2817, which is worse than either loop architecture. The explanation is in (8.16). The anticipation term acts only while the demand is changing, so it is free during the quiet phases, while the lagged steady-state term still leads the plant during the slow throttle reductions of the descent even at a 45 s lag.

The predictive controller wins the flight by a wide margin, 941 % s against 1487 for the two loops, because it coordinates both actuators against a setpoint trajectory it can see. Its temperature figure of 20 144 °C s is a direct consequence of the $\pm 12\,^{\circ}\mathrm{C}$ zone rather than a failure, and closing the zone to zero returns the temperature figure to 10 358 °C s and the power figure to 868 % s at the cost of the thrust deficit (11 543 against 10 350 °C s).

8.6.3 The Value of Preview

The last row of Table 8.4 is the central experiment. It is the same controller with the same model, the same horizon, the same weights and the same actuators, and the only change is that the setpoint trajectory over the horizon is replaced by the present setpoint held constant, which is what ordinary model predictive control does. Integrated power error over the takeoff goes from 729 to 6358 % s, almost a factor of nine, and the thrust deficit from 24 to 6231 °C s. Over the full flight the tracking error goes from 941 to 10 276, which is slightly worse than the 9990 of the open-loop case.

> ✦ **DID YOU KNOW? Preview Off, Worse Than Nobody Driving**
>
> From the simulator: over the full flight, the dynamic matrix controller with its preview disabled integrates 10 276 % s of power error, slightly more than the 9990 % s of the plant with nobody driving. With the preview it integrates 941 % s (Table 8.4).

The feedforward architecture isolates the same effect in a simpler controller. Setting t_{lead} to zero, which converts (8.16) into an ordinary steady-state feedforward acting on the present flight condition, takes the takeoff error from 448 to 681 % s, which is worse than having no feedforward at all. Setting the anticipation gain K_{ff} to zero and keeping the preview gives 657. Neither half improves on the 641 % s of the two loops alone, and together they are worth 193 % s against having neither.

Preview also buys the thrust deficit directly, and at a price. Extending t_{lead} to 45 s at $K_{ff} = 0.2$ takes the deficit from 57 to 11 °C s with the worst shortfall down from 6 °C to 3, while the integrated power error rises to 930 % s. That trade is physical rather than a tuning artifact. During a ramp the reactor must supply the demanded duty and the energy going into its own thermal capacities at the same time, so tracking the demand exactly and satisfying the engine exactly are different objectives.

8.6.4 Go-Around

The go-around is the transient the mission is built around and the only one that is not in the profile. The aircraft is on approach at about a quarter throttle with the demand at 9 % of rated power and a core that has been cooling for ten minutes, and at $t = 100\,\mathrm{s}$ the throttle goes to 95 % in twelve seconds. Table 8.5 reports it.

The two independent loops exceed the containment limit. Power passes the demand by 144.3 points at 152 s, and the core outlet crosses 850 °C at 147 s, peaks at 859 °C and spends 8.3 s above the limit. Preview feedforward does not help (142.9 points, 859 °C), because the event is not in the profile and there is nothing to read ahead.

🕮 **Pause.** Try Exercise 8.3 before reading on.

Table 8.5 Go-around from the approach, 300 s window. Over and under are the largest excess of power above the demand and shortfall below it, in points of rated power

Architecture	over	under	peak T_{fo}	$t > 850\,^{\circ}\mathrm{C}$	deficit
Nobody driving	62.6	72.4	792	0 s	2174
Power loop only	14.5	63.0	793	0 s	1354
Two loops	144.3	63.5	859	8.3 s	2228
Preview feedforward	142.9	61.8	859	8.3 s	2150
Dynamic matrix control	51.6	65.8	815	0 s	1417

The overshoot comes from the two loops acting together rather than from the plant or from the power loop. With nobody driving, the power lags the demand by 72.4 points while the throttle opens and passes it by only 42.9 points at 182 s, once the salt cooled by the opened heat sink has crossed the loop and the fuel temperature coefficient has delivered its reactivity. Its larger excess in Table 8.5 comes later, when the throttle returns to the approach schedule. The power loop alone, with the pump parked at 100 %, passes the demand by 7.9 points during the go-around and keeps the outlet at 793 °C. Closing the temperature loop as well takes the pump from 61 % to its 40 % stop as the throttle opens and then to 125 % by 146 s as the outlet climbs, and because the pump also sets β_{eff}, every one of those moves is a reactivity change. The power stalls near 54 % between 130 and 136 s, the power loop keeps withdrawing the rod at its rate limit to 70.3 %, and that stored rod reactivity arrives together with the plant's own delayed insertion.

Of the three architectures that control the outlet, the predictive controller is the only one that stays under the limit. Its worst power error is a 65.8-point shortfall while the throttle is still opening, it passes the demand by 51.6 points, and the outlet peaks at 815 °C. It does so by moving the fuel pump and the rod together against a prediction that already contains the plant's own passive response, so it does not add its own on top. The tell-tale is the fuel pump. The two-loop system reaches 125 % only one second before the outlet crosses the limit, while the predictive controller puts the pump on its 125 % stop at the execution that coincides with the go-around command, before the throttle has moved.

> **TRY IT IN THE APP 9 · Go-Around From the Approach**
>
> Load scenario 9 and run it on Two loops through the go-around at 100 s, watching the core outlet on the Plant tab. It peaks at 859 °C and spends 8.3 s above the 850 °C limit. Pause, press Reset, choose MPC before pressing Run and repeat: the outlet now peaks at 815 °C.

8.6.5 Degradation and the Graded Challenge

The graded scenario runs the full flight from a deliberately sluggish tuning with 0.8 °C of thermocouple noise, 0.5 % of flux channel noise, and three undeclared events: a 20 % loss of air heater conductance at 620 s, a 6 °C thermocouple bias at 900 s and a fuel pump delivery gain of 0.85 at 1180 s. The target is reactor power within ± 1.2 % of demand through the last third of the flight (from $t = 1188$ s) with the outlet below 850 °C and rod travel under 150 %. The sluggish tuning as given misses by more than a factor of four, at 5.07 %.

The predictive controller holds 0.92 % with a peak outlet of 818 °C for 99 % of rod travel and 162 % of pump travel. The two loops at the tuning of Table 8.4 reach 2.68 % for 258 % of rod travel, and tightening them to $K_{c,P} = 1.6$ with $\tau_{I,P} = 30$ s buys 1.55 % for 397 % of rod travel. Preview feedforward at the same tuning is the worst of the four at 8.93 %.

The feedforward result is the one worth stating carefully. A feedforward is an open-loop correction computed from a model, so a plant degradation does not disturb it, it invalidates it, and the feedback

loop underneath then has to clean up a correction that is confidently wrong. The feedforward's model keeps the declared air heater conductance, so the fouling at 620 s leaves it unchanged. The predictive controller survives the same two degradations because its model bias in (8.19) is measured every execution, so the degradation appears in the bias rather than in the controlled variable.

8.6.6 Saturation Without an Undersized Actuator

Holding the aircraft at ground idle (15 % throttle) for five minutes with the program slope k_T set to zero leaves the outlet 23.1 °C below a fixed 795 °C setpoint with the fuel pump pinned on its 40 % stop, and integral action never moves it. The reachable band at that condition is 763 to 773 °C. Restoring the slope slides the setpoint to 743 °C, which is outside the band in the other direction, and the pump pins on its 125 % stop with the outlet 19.5 °C high after a further five minutes. At 100 % throttle, with the run started lined out there by a Reset, both settings work and the pump sits at 101.5 and 103.1 %.

The pump has not run out of travel. At 6.6 MW the core temperature rise is 7.4 °C at design flow, and the pump's whole range only moves it between 6.0 and 18.6 °C, which cannot move an outlet temperature twenty-three degrees. The actuator that can move it there is the rod, worth 2.3 °C per percent at that condition against the pump's 0.15, which is the relative gain of 3.58 written out.

Pause. Try Exercise 8.2 before reading on.

8.7 Technology in Practice

The two measurements of Figure 8.2 are, in a plant, a neutron detector outside the reactor vessel and a thermocouple in the core outlet line, and the two manipulated variables are a regulating rod on its drive and a fuel pump on a speed drive. Figure 8.5 draws the loops with that hardware in place of the schematic's symbols.

What Measures It

Neutron flux is measured by fission chambers, ion chambers and self-powered detectors, and detectors outside the vessel give the flux as a representation of reactor power [21]. In a molten salt reactor the flux reading serves as the fast indication of power for feedback control, while the longer-term power comes from heat balances and from samples of the fuel salt [22]. The core outlet temperature is read by thermocouples in thermowells inserted into the flow [22], and every thermowell adds to the first-order lag of its sensor through its mass and specific heat [23]. The chapter carries both instruments as a first-order lag with colored noise and a constant bias, and the graded scenario adds thermocouple noise, flux channel noise and a thermocouple bias.

TECHNOLOGY Sensor: An Ion Chamber Outside the Vessel

Power-range neutron detectors are typically uncompensated ion chambers, since the neutron signal is large against the gamma [21]. In a molten salt reactor they would sit outside the vessel, should answer a change of power within about a millisecond, and be calibrated against a measurement of thermal power. No high-sensitivity neutron detector that works above 600 °C is commercially available. Heat lowers the resistance of a chamber's ceramic insulators and causes spurious discharges whose pulses cannot be told from neutron pulses [22], so a detector run too hot reports counts that are not neutrons. The Aircraft Reactor Experiment cooled its special fission chambers in the reflector with helium [10]. The chapter's NI is a lag with colored noise and a bias.

What Moves It

The power loop writes the rod drive at no more than 0.6 %/s of travel, and the temperature loop writes the fuel pump between 40 and 125 % of design speed with no slew limit. The rod worth follows an S-curve, and at the top of the descent the rod sits on the flat of that curve, where it moves the outlet temperature least.

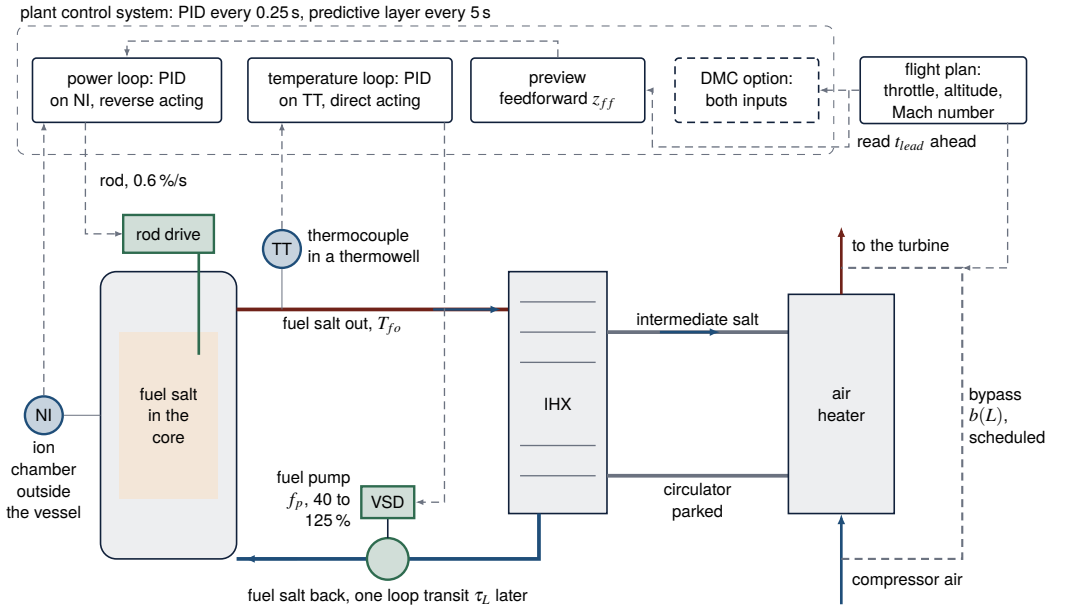

Figure 8.5 The two loops of the chapter with plant hardware in place of the symbols of Figure 8.2. Sensors are blue circles and final elements green: the neutron instrument NI, an ion chamber outside the vessel, the core outlet thermocouple TT in its thermowell, the regulating rod on its rate-limited drive and the fuel pump on a variable-speed drive (VSD). The power loop writes the rod, with the preview feedforward in its bias, and the temperature loop writes the pump. The predictive option replaces both loops and reads the flight plan over its horizon. Notice that the flight plan reaches the controllers directly, while the plant feels the same throttle only through the air heater, the intermediate salt, the exchanger and the fuel loop.

TECHNOLOGY Final Element: A Regulating Rod on a Slow Drive

The Aircraft Reactor Experiment had one stainless-steel regulating rod. It moved 12 inches about the center of the reactor, was worth 0.4 percent in reactivity over that stroke and could be moved at 0.011 percent per second. Its three larger rods of boron carbide and iron each took 13 minutes to withdraw [10]. In power reactors the rod position is indicated twice, by a detector on the drive housing and by counting the signals sent to the drive, and after a dropped or stuck rod the two disagree. Power supply failures and drift have dropped or slipped rods [24]. The chapter's rod moves at 0.6 %/s, 10.5 pcm/s at the design position.

Where the Controller Runs

Control systems of this kind run loops with cycle times from milliseconds to minutes, and a supervisory layer sends setpoints to the field controllers below it [25]. The chapter's two PID blocks execute every 0.25 s over a 0.05 s integration step: a reverse-acting block on the neutron instrument with the rod drive's slew limit on its output and the preview feedforward in its bias, and a direct-acting block on the outlet thermocouple whose setpoint follows the temperature program. The dynamic matrix controller is the supervisory layer. It executes every 5 s, writes both actuators at once and reads the flight plan over its 120 s horizon.

> **IN PRACTICE** **Measuring the Plant While It Runs**
>
> For the reactivity control of a molten salt reactor, the variables to watch are the position of the control element and the response of the neutron flux to each change. Small perturbations may be impressed on the power to evaluate the reactor's dynamics online, as the pseudo-random binary movements of the regulating rod did at the Molten Salt Reactor Experiment. In the United States, general design criterion 26 of 10 CFR Part 50, Appendix A, requires two independent means of bringing and holding the reactor subcritical [22]. Scenario 2 of this chapter measures the gain matrix in the same spirit, by stepping the rod and then the pump at a held condition.

What Is Modernizing It

Three developments reach these loops. Neutron noise, cross-correlated between several detectors outside the vessel, is expected to show how the fuel salt is distributed in the core, for example after a flow blockage or with vessel vibration [22]. Operating plants are moving from analog to digital instrumentation and control through individual system upgrades [21]. And control systems that adopt office computers, operating systems and network protocols have far less isolation from the outside world and a greater need to be secured [25], which matters to a controller that reads a flight plan from outside the reactor.

8.8 Limitations

The reduced-order structure used here is the structure that the Aircraft Reactor Experiment and Molten Salt Reactor Experiment literature supports for circulating-fuel *neutronics* and thermal *time scales* [13, 14]. The thermal system of this aircraft plant is notional. It has not been validated against any experiment, and published work on reduced-order models of this class cautions that transient core temperature predictions from them are less reliable than their power predictions. The temperatures in this application should be read as a control-relevant model rather than as a thermal-hydraulic prediction, and nothing here is evidence about whether a nuclear aircraft could be built.

The predictive controller is unconstrained and then clipped. Move and position limits are imposed by clamping the computed move into the actuator box rather than by solving a constrained quadratic program [7], and the fuel temperature limit enters as a weighted zone rather than as a hard constraint. The go-around result in Table 8.5 is therefore correctly stated as an observation that the outlet stayed below 850 °C and not as a guarantee that it could not exceed it.

The air heater bypass is a schedule rather than a governor, which is a deliberate choice. A real engine control system would close a loop around turbine inlet temperature and hide the reactor behind it, and there would then be no reactor control lesson. The consequence is that the schedule is wrong in ways a real design would not be, particularly at altitude, and the power loop is cleaning that up as well as rejecting the disturbance.

Instruments are a first-order lag plus colored noise plus a constant bias. A thermocouple in a salt thermowell has a nonlinear calibration, drifts, and fails in ways that do not resemble a constant offset. Finally, the fixed step and the deterministic pseudorandom noise make every run reproduce exactly on any machine, which a real plant does not, and that repeatability should be treated as a teaching convenience.

> **OPEN CHALLENGE** **A Neutron Detector for the Hot Core**
>
> Salt reactor cores run near 750 °C, and no high-sensitivity neutron flux detector that works above 600 °C is commercially available. Inside the core, long exposure to heat lowers the resistance of ceramic insulators, causes spurious discharges, outgasses the chamber body into its fill gas and weakens the electrode supports. Detectors outside the vessel avoid this, but they will not give a detailed picture of the flux inside the core [22]. Progress would be a detector that survives years at core temperature, so that power could be read where it is made, not only through the vessel wall.

8.9 Conclusion

This chapter studied a reduced-order model and an interactive app for a notional molten salt reactor supplying heat to aircraft turbojets. The reactor model retains circulating-fuel point kinetics with delayed neutron precursor transport, thermal reactivity feedback and lumped energy balances in the form the Aircraft Reactor Experiment literature supports, and the aircraft enters as a scheduled disturbance through the corrected air flow group and the compressor discharge temperature.

Three results follow from the mission comparison. The plant follows a takeoff without being commanded, lagging it by 90.9 points of rated power at the end of the ramp and then overshooting to 149.5 % of rated, so inherent load following reduces the size of the control problem without removing it. The conventional pairing of control reactivity to power and fuel circulation to core outlet temperature has a relative gain of 0.86 at takeoff and 3.58 at ground idle, so a pairing chosen at the design point is strongly interacting before the aircraft lands, and the fuel pump saturates at low power because its process gain has gone rather than because it is too small. Advance knowledge of the flight profile is worth more than any retuning of the feedback loops, taking the integrated power error over a takeoff from 6358 to 729 % s for a predictive controller that is otherwise identical.

The go-around is the case that separates coordination from feedforward. Two independent loops that hold every steady phase of the flight put the fuel outlet 9 °C above its limit for eight seconds, feedforward does not help because the event is not scheduled, and a coordinated predictive controller peaks at 815 °C. The eleven scenarios of the app, each with a completion rule, and its data export carry these studies of multivariable pairing, feedforward timing and preview control on a plant whose disturbance is a published plan, and the exercises of Section 8.10 build on them.

KEY POINTS

- A reactor whose fuel is a pumped liquid salt follows a takeoff on its own, but it lags the demand and then overshoots to 149.5 % of rated.
- The rod-to-power and pump-to-temperature pairing has a relative gain of 0.86 at takeoff and 3.58 at ground idle.
- At low power the fuel pump saturates because its process gain has gone, not because it is too small.
- Preview of the flight profile takes the takeoff power error of the same predictive controller from 6358 to 729 % s.

Where else this applies. Volume 9, Chapter 5 also acts on published orders ahead of time.

References

[1] E. S. Bettis, R. W. Schroeder, G. A. Cristy, H. W. Savage, R. G. Affel, and L. F. Hemphill, “The Aircraft Reactor Experiment–Design and construction,” *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 804–825, 1957.

[2] S. Comassar, “General electric direct-air-cycle aircraft nuclear propulsion program: Aircraft nuclear propulsion application studies (comprehensive technical report),” General Electric Company, Cincinnati, OH, Tech. Rep. APEX-911, 1962.

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

[4] J. P. Carter, R. Christensen, and S. Yoon, “Numerical analysis of dynamic load following response in a natural circulation molten salt pebble bed reactor,” *Nuclear Engineering and Design*, vol. 394, p. 111818, 2022.

[5] 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.

[6] H. A. Pak and R. Shieh, “On MIMO optimal preview tracking control for known trajectory models,” *Optimal Control Applications and Methods*, vol. 12, no. 2, pp. 119–130, 1991.

[7] C. E. Garcia and A. M. Morshedi, “Quadratic programming solution of dynamic matrix control (QDMC),” *Chemical Engineering Communications*, vol. 46, no. 1-3, pp. 73–87, 1986.

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

[9] E. S. Bettis, W. B. Cottrell, E. R. Mann, J. L. Meem, and G. D. Whitman, "The Aircraft Reactor Experiment–Operation," *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 841–853, 1957.

[10] 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.

[11] W. K. Ergen, A. D. Callihan, C. B. Mills, and D. Scott, "The Aircraft Reactor Experiment–Physics," *Nuclear Science and Engineering*, vol. 2, no. 6, pp. 826–840, 1957.

[12] G. Thornton and A. J. Rothstein, "General electric direct-air-cycle aircraft nuclear propulsion program: Program summary and references (comprehensive technical report)," General Electric Company, Flight Propulsion Laboratory Department, Cincinnati, OH, Tech. Rep. APEX-901, 1962.

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

[14] S. J. Ball and T. W. Kerlin, "Stability analysis of the Molten-Salt Reactor Experiment," Oak Ridge National Laboratory, Tech. Rep. ORNL-TM-1070, 1965.

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

[16] J. Shi and M. Fratoni, "GeN-Foam model and benchmark of delayed neutron precursor drift in the Molten Salt Reactor Experiment," *EPJ Web of Conferences*, vol. 247, p. 06040, 2021.

[17] M. Elhareef and Z. Wu, "Benchmarking a point reactor kinetics method with delayed neutron precursors transport using a multiphysics coupled framework," *Annals of Nuclear Energy*, vol. 217, p. 111366, 2025.

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

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

[20] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[21] 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.

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

[23] 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.

[24] 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.

[25] 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.

8.10 Exercises

8.1 Concept check · 2 min

At steady power and design flow the control rod of the aircraft reactor holds about 260 pcm of positive reactivity, and Equation (8.4) sets $\rho_0 = \beta - \beta_{\text{eff}}$. Why is a circulating-fuel core critical at positive reactivity?

(a) The negative fuel coefficient has to be paid for at rated power, so the rod holds the power defect

(b) The moderator coefficient is small, so the rod must supply what the moderator does not

(c) The prompt jump approximation removes the prompt neutrons, so the rod must replace them

(d) Precursors that decay in the external loop are lost to the chain, so prompt multiplication makes up the share

8.2 Concept check · 1 min

In scenario 6 at ground idle with the program slope $k_T = 0$, the core outlet settles 23.1 °C below its setpoint with the fuel pump on its 40 % stop, and the integral never moves it. What is the right diagnosis?

(a) The pump has run out of process gain, because the core rise it can redistribute is only a few degrees

(b) The pump is undersized, and a pump with a wider speed range would reach the setpoint

(c) The integral has wound up, and back-calculation is switched off on the temperature loop

(d) The thermocouple lag is too long, so the loop cannot see the outlet move in time

8.3 Concept check · 2 min

In the go-around of scenario 9, preview feedforward leaves the core outlet at 859 °C, no better than the two loops alone. Why does the preview not help here?

(a) The feedforward lag of 45 s is too slow to follow a 12 s throttle change

(b) The go-around is not in the scheduled profile, so the preview has nothing ahead to read

(c) The rod rate limit stops the feedforward term from reaching the rod during the ramp

(d) The feedforward acts only on the rod, so the pump overshoot that drives the outlet is unchanged

8.4 Theory and calculation · 6 min

Equation (8.13) gives the steady-state gains at the sea-level takeoff condition. The app's steady-state model gives, at landing idle, $\partial P/\partial z = 0.02326$, $\partial P/\partial f_p = -0.001116$, $\partial T_{fo}/\partial z = 2.287$ and $\partial T_{fo}/\partial f_p = -0.1523$, in the same units.

(a) Compute the relative gain λ_{11} of the pairing rod to power and pump to temperature at takeoff. *Checkpoint:* 0.862, within 0.005.

(b) Compute λ_{11} at landing idle. *Checkpoint:* 3.58, within 0.02.

(c) With the temperature loop closed and holding the outlet exactly, the steady gain of the rod on power is K_{11}/λ_{11}. Compute it at both conditions and compare it with the open-loop gain. *Checkpoint:* 0.637 at takeoff and 0.0065 at idle, % power per % rod.

8.5 Theory and calculation · 5 min

The core fuel salt has a heat capacity $C_f = 5.357\,\mathrm{MW\,s/°C}$ and a core transit time of 6.0 s at design flow, and both scale as in Equation (8.3). Take the duty as fixed.

(a) Compute the flow capacity rate W_f at design flow and the core temperature rise Q/W_f at the rated 125 MW and at 6.6 MW. *Checkpoint:* 0.893 MW/°C, 140 °C and 7.39 °C.

(b) Compute the rise at 6.6 MW with the pump on its 40 % and 125 % stops. *Checkpoint:* 18.5 °C and 5.91 °C, each within 0.1 °C.

(c) Each core node carries a linear profile, so the mean fuel temperature is very nearly the average of inlet and outlet and $T_{fo} = \bar{T}_f + \Delta T/2$. With the mean held fixed, how far can the whole pump range move the outlet at 6.6 MW? *Checkpoint:* 6.3 °C, within 0.1 °C.

8.6 Theory and calculation · 7 min

At the sea-level takeoff condition the Mach number is 0.25 and the ambient temperature is 288.15 K. The compressor has $\pi_c = 6.0$ and $\eta_c = 0.85$, the air has $c_{p,a} = 1.15\,\mathrm{kJ/(kg\,K)}$, and the engine schedules a turbine inlet temperature of 700 °C at full throttle. Use Equations (8.9) to (8.11) and Table 8.2.

(a) Compute the compressor discharge temperature T_{c2}. *Checkpoint:* 248.0 °C, within 0.5 °C.

(b) At sea level $\delta/\sqrt{\theta} = (1+0.2M^2)^3$. Compute it, the engine air flow at full throttle and the capacity rate $\dot{m}_a c_{p,a}$ of the air through the heater with the bypass closed. *Checkpoint:* 1.038, 240.8 kg/s and 0.277 MW/°C.

(c) Compute the demanded duty Q_{dem} and compare it with the rated power. *Checkpoint:* 125.2 MW, within 0.3 MW.

8.7 App challenge · 7 min

Load scenario 1, Fly the mission with nobody driving (MANUAL), and change no setting. Run at 20× to $t = 4$ min at the default seed, so the run takes 12 seconds, and export the CSV file. Its time column is in minutes. The throttle reaches full at $t = 105$ s.

(a) Find the peak of `power_pct` and the time at which it occurs, in seconds. *Checkpoint:* 149.5 %, within 0.1 %, at 157 s, within 1 s.

(b) Find the largest shortfall of `power_pct` below `power_sp_pct` and when it occurs. *Checkpoint:* 90.8 points, within 0.2, at 104.5 s, within 1 s.

(c) Find the largest value of `tti_sp_C` minus `tti_C`, and explain with the chapter's mechanism why the power first lags and then overshoots. *Checkpoint:* 48 °C, within 0.5 °C, at about 113 s.

8.8 App challenge · 8 min

Load scenario 10, Coordinated MPC with the profile in view. Each run goes to $t = 5$ min at 20× (15 seconds) at the default seed. Read the power loop's IAE on the scorecard or in the last row of column `power_iae`, and the thrust deficit in the last row of column `thrust_deficit_Cs`.

(a) Run the case as shipped, with MPC setpoint preview on Preview on. *Checkpoint:* IAE 729 % s and deficit 24 °C s, each within 2.

(b) Pause, Reset, set MPC setpoint preview (Advanced) to Preview off before Run, and run again. Give the ratio of the two power IAE values. *Checkpoint:* IAE 6358 % s, within 5, deficit 6231 °C s, within 5, ratio 8.7.

(c) Reload the scenario, set MPC temperature zone (Advanced) to 0 with the preview on, and run a third time. Does closing the zone cost the power tracking on this window? *Checkpoint:* IAE 638 % s, within 2.

8.9 App challenge · 8 min

Load scenario 9, Go-around from the approach. The go-around is commanded at $t = 100$ s. Run each case to $t = 160$ s at 10× (16 seconds) at the default seed and export the CSV file, whose time column is in minutes.

(a) Run on Two loops as shipped and read the Peak fuel margin tile. *Checkpoint:* −9 °C, an outlet peak of 859 °C.

(b) Pause, Reset, choose MPC before Run, and run again. Read the same tile. *Checkpoint:* 35 °C, an outlet peak of 815 °C.

(c) In each file find the first row at which `pump_cmd_pct` reads 125 %, in seconds, and explain what the difference shows. *Checkpoint:* 146.5 s on two loops and 100.25 s on MPC, each within 0.5 s.

8.10 Application area · 3 min

Without the air heater bypass, a fixed-geometry heater with the core near its design temperature pulls 26 MW at ground idle.

(a) The demand there is 7 % of the rated 125 MW. By what factor does the heater's pull exceed it? *Checkpoint:* About 3, 21 % against 7 %.

(b) The only other way to shed the surplus is to cool the salt to 409 °C. Why is that not available, and what does the bypass cost the power loop instead?

8.11 Application area · 4 min

A safety reviewer wants to cite Table 8.5 as evidence that the dynamic matrix controller keeps the fuel outlet below 850 °C in a go-around.

(a) Give two reasons from the chapter why the table does not support that claim.

(b) Name one change to the controller that would make the limit a constraint it enforces, and say what would still not be proven.

8.12 Review the agent · 5 min

An AI assistant writes: “The relative gain of rod to power and pump to temperature is 0.86 at takeoff, so the pairing is nearly decoupled. Tune the two loops one at a time at the sea-level design point. Because the relative gain is a property of the plant, the same pairing and tuning will behave the same way for the whole mission.” Find the flaw and say how you would check it in the app.

CHAPTER 9

Multi-Module Steam Header Simulator for Load Allocation and Plantwide Coordination

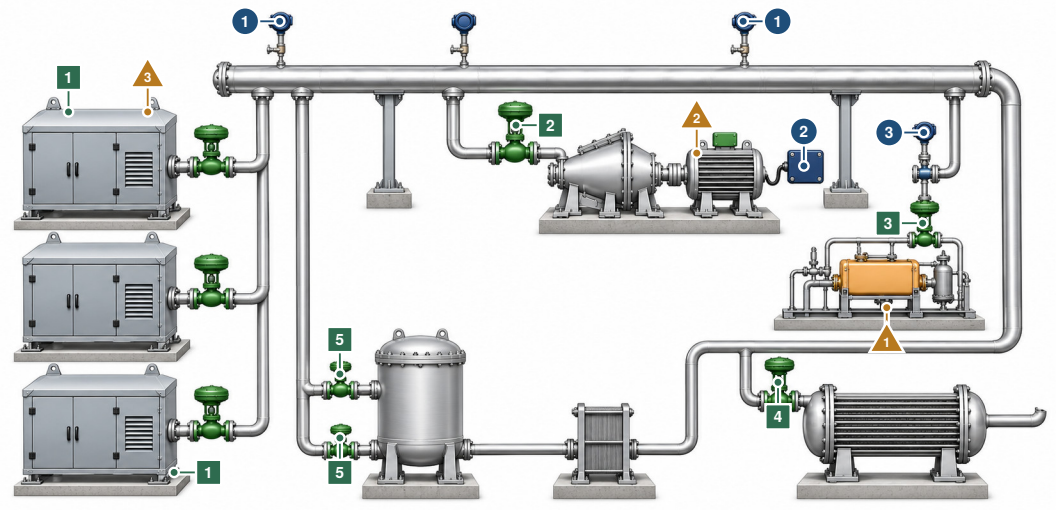

The controller moves 1 reactor power of each module 2 turbine valve 3 process steam valve 4 bypass valve to the condenser 5 storage charge and discharge valves

The controller reads 1 header pressure, three transmitters 2 generator output 3 process steam flow

What disturbs it 1 process steam request 2 grid load 3 a module trip

Notice that every module and every user meets in the one header pressure, the only number they all see. Illustration, not to scale.

IN PLAIN TERMS **Three Reactors Sharing One Steam Pipe**

The main idea. Three small reactors send steam into one shared pipe that feeds a turbine and a factory. The pressure in the pipe shows whether the reactors make as much steam as the users take.

The challenge. Give each reactor its own controller for that pressure, and the reactors quietly drift apart while the pressure looks perfect. One controller must answer the pressure, and a separate rule must share the work so that no reactor runs nearer its heat limit than the others.

Why it matters. A pipe at the right pressure can hide a reactor pushed to its limit.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of Section 9.1, the history box of Section 9.2, the try-it and fact boxes of Section 9.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 9.6, Section 9.7, Section 9.8 and the open challenge

The control question Three small reactors feed one header that holds 7.55 s of steam. Who should answer a header pressure error, and by what rule should the load be divided when each module's ceiling moves with the header?

Learning objectives

- Compute the time scales of the header, the modules and their ramp limits.
- Explain why independent integral loops on one measurement walk apart.
- Derive margin-proportional allocation and compare it with an equal split.
- Compare master-slave with coordinated control and price the bypass rest.
- Decide which product a runback should shed after a module trip.

Prerequisites Sections 1.6.2, 1.3, 1.4.5 and 1.6.4, and Section 1.6.3 for the predictive controller. Chapter 4 and Chapter 5 follow the steam of one reactor.

The app https://apmonitor.com/apps/smr-steam-header/

See also Volume 6, Chapter 7 and Volume 6, Chapter 8 for plantwide control structure. Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging a fast, costly actuator, and Volume 4, Chapter 12 for allocation among actuators at different prices.

9.1 Introduction

Three reactors that feed one pipe share one number. Header pressure rises when the modules make more steam than the users take and falls when they take more, so it is the plant's statement of whether supply matches demand, and it is also the variable no single piece of equipment owns. Every module pushes it up, every user pulls it down, and the pressure reaches back into all of them at once, because it sets the saturation temperature each steam generator boils against and the flow each user's valve passes. A block diagram of one module and one pressure loop is a correct picture of a plant that does not exist. The plant that exists has three producers with different capability, four consumers with different value, and an inventory of steam between them that lasts 7.55 s.

That configuration produces a failure that a transfer function cannot show. Put an identical, honestly tuned proportional-integral controller on each module, all three holding the same header pressure setpoint, and apply no disturbance at all. The three module powers walk apart, reaching 70.5 MW of spread at 600 s and 91.7 MW at 1200 s, while header pressure stays between 3.385 and 3.401 MPa inside a tolerance band of ±0.08 MPa. Nothing on the header alarms and nothing oscillates. The only signs are per-module ones, a margin-spread score that turns red at 135 s with the modules 20.5 MW apart and a module coolant temperature warning at 271 s, by which time they are 48 MW apart. The misconception this chapter corrects is the belief that a controlled variable held on setpoint is evidence that the plant is under control, when what is under control is one measurement and the megawatts behind it are free to redistribute until two of the three modules sit against their own thermal limits.

The class of problem is general. Several nearly identical producers discharge into one small-inventory buffer, consumers draw from it, and the buffer pressure or level is the only variable they all see. Parallel boilers on a common steam header are the oldest instance, and the plantwide question of which loop owns a shared variable and by what rule the rest of the plant is divided is the subject of a settled literature [1–3]. Parallel compressors on a common discharge manifold pose the same allocation problem with efficiency rather than thermal margin as the currency to be divided [4], and parallel chillers on a common chilled-water loop add a sequencing decision on top of the split [5]. Synchronous generators and grid-forming inverters on a common bus divide load through droop for exactly the reason a header master exists, because more than one integral controller on one frequency or one pressure leaves the division of output indeterminate [6]. In every case the regulation of the shared variable is the easy half. Allocation is the decision that matters, and it is made against constraints that belong to the individual producers and never appear in the shared measurement.

Multi-module reactor plants make the problem urgent. Header pressure control for a plant with several reactors and several loads has been posed and solved as a decoupling problem [7], and

unit-specific load-following across a multi-unit small modular reactor plant has been treated as an optimization over the individual units rather than as a single regulation loop [8]. That work addresses the designer choosing a control structure and presumes that the structure is already accepted as where the answer lies. A student shown a block diagram with one pressure controller has no way to discover why a second one is forbidden. This chapter works with an interactive three-module steam header plant, set out in enough detail to rebuild it from the text, and measures what five control architectures and three allocation rules do to the same plant under the same upsets. Its central idea is the separation of the header problem into a fast regulation component, answered by a bypass valve in seconds, and a slow allocation component, answered by three rate-limited reactors against a capability ceiling that moves as they load. The measured runs show that the two components have different owners and different failure modes, on a plant where a well-regulated variable and a badly run plant coexist and can be watched together. The reactor model and the allocation arithmetic are standard.

> ✦ **DID YOU KNOW? Twelve Reactors and One Refueling Tool**
>
> Small modular reactors are usually planned with several units on one site. One design, of 77 MW electrical modules, can house up to twelve units on a site. Because the site has a single disassembly tool, only one unit can be refueled at a time, a job of ten days [9].

Section 9.2 traces load division on a common header from parallel boilers to multi-module reactor plants. Section 9.3 collects the eighteen-state model, its instrument chain and the constants that reproduce the simulator. Section 9.4 measures the time scales, the coupling between the modules and the capability ceiling that moves with the header, and derives the allocation rule that equalizes margin. Section 9.5 describes the common loop algorithm, the five architectures, the allocator and its spill pass, the two-stage runback and the model predictive controller. Section 9.6 reports the twelve scenarios, from independent control fighting itself to the graded challenge and the predictive controller against coordinated allocation. Section 9.7 sets the plant in its hardware, Section 9.8 states the limitations, Section 9.9 draws out what generalizes, and Section 9.10 summarizes the measured results.

9.2 Background

Dividing one load among several boilers on a common header is one of the oldest problems in industrial process control, and the earliest arrangement was mechanical. A boiler master reduced header pressure to a single air signal that reached every firing rate controller in parallel through a manual bias station on each one. The bias station is the ancestor of every allocation rule that followed, because it decided how the common demand was divided while the master decided only how large the demand was. Only one integral device existed in that scheme, which is the reason it worked and the reason a single pressure master feeding a participation rule is still the recommended arrangement [3]. Where a second was added, one header controller per boiler, the division of firing rate became indeterminate and the plant drifted for reasons no single loop explained.

Dynamic models of the boiler side matured well after the control practice, and the drum boiler acquired a compact and widely reused nonlinear model whose central result is that shrink and swell couple pressure to level and tie the two loops together [10], the inverse response treated in Section 1.4.4 and Chapter 4. The essential feature of that model here is the drum itself, because a recirculating boiler holds a large water and steam inventory and can absorb a load change for a minute or more before its firing rate has to answer. Remove that inventory and the timing of the whole plant changes. A once-through steam generator boils to superheat over its length and holds almost nothing, so feedwater and steam flow are tied together and the unit cannot buffer anything, which is why feedwater control across several once-through generators is treated as a coupled multi-unit problem [11]. Helical-coil geometry is the standard choice inside an integral reactor vessel, and its

thermal-hydraulic analysis shows the moving boundaries between economizer, boiling and superheat regions that a reduced-order model has to lump [12].

Small modular reactors bring the parallel-boiler problem back with the inventory removed. Factory fabrication, staged commissioning and independent dispatch of the individual modules are the economic arguments for the concept, and the same arguments have been made and revised across several deployment cycles [13]. Integral pressurized water reactor designs place a helical-coil once-through steam generator inside a natural-circulation vessel and bring several such modules into one turbine hall, coupled to more than one product through a common steam supply [14]. That coupling is now studied as a system in its own right, with dynamic models built for plants selling electricity and hydrogen from the same steam [15] and for plants that add thermal storage as a further participant on the steam balance [16]. The header in all of those arrangements belongs to none of the products, so multi-reactor, multi-load header pressure control has been posed directly as a decoupling design [7] and multi-unit load following as an allocation across units with individual constraints [8]. Both answers match the practice parallel boilers arrived at, one master and one rule, and both add the constraint the older plants did not have, which is that each producer's capability moves with the shared variable.

> ⧖ HISTORY **Steam Bypass on Demand (2009)**
>
> From 2009 a surplus of baseload generation in one Canadian province led its system operator to require nuclear plants to cut their output when generation ran high and demand ran low. Selected units kept their reactors at full thermal power and cut generator output by up to 300 MW of electricity for short periods, sending the excess steam to the condenser through the steam bypass system. That system had been designed for rare loss of load events and not for frequent use, so its capability had to be evaluated first. Because the plants had several units each, the operating staff could choose which unit to run at reduced power, sparing units with equipment problems. A review published in 2018 counted more than 1500 such operations without an event at one eight-unit station [17].

9.3 Process Model

The simulated plant is three integral pressurized water reactors of 250 MW rated thermal power, each a single vessel standing in a pool with no primary coolant pumps and a helical-coil once-through steam generator wound around the riser inside the vessel. Each module discharges through a main steam isolation valve into one header, which supplies a turbine generator, an industrial process-steam customer, a bypass to the condenser, a thermal storage loop and a relief valve, as shown in Figure 9.1 with every symbol the equations below use. The three modules are deliberately not identical, because a plant one year after commissioning has modules at different burnup with steam generators in different condition, and that difference is what an allocator has to divide.

Table 9.1 classifies what the control system interacts with. The plant is expressed in the units of steam-plant operation, pressures in MPa, temperatures in °C, flows in kg/s, thermal and electrical power in MW, enthalpies in kJ/kg, storage in MWh thermal and time in seconds.

🕮 **Pause.** Try Exercise 9.1 before reading on.

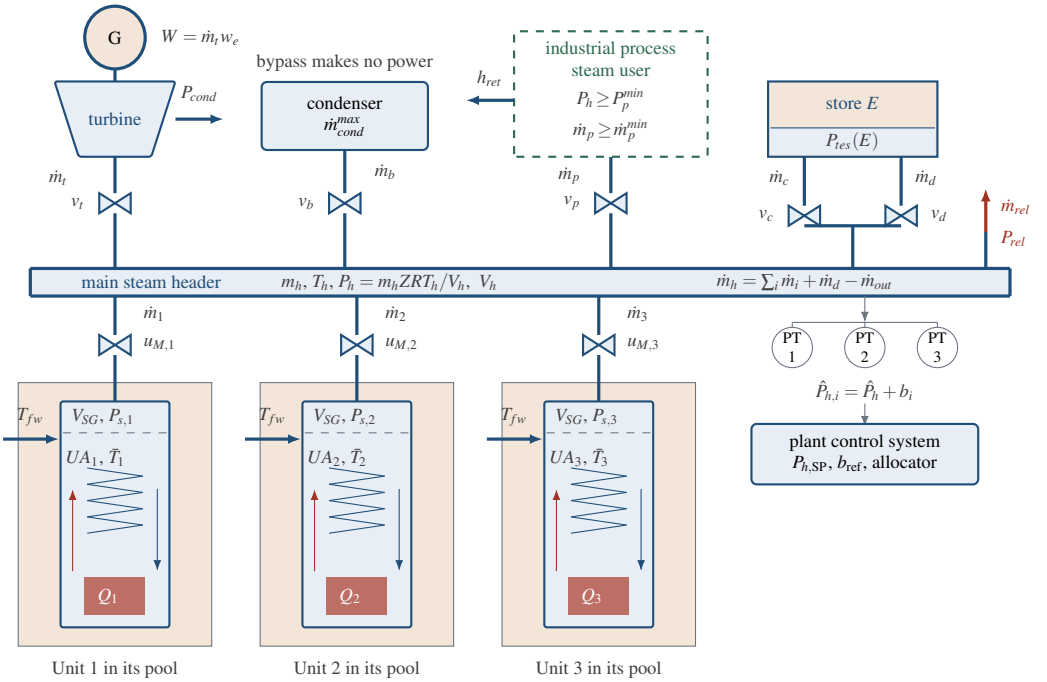

Figure 9.1 Three integral pressurized water reactors on one main steam header, with every symbol used in the equations. Core power Q_i heats a primary loop of average temperature $\bar{T}_i$, which passes $UA_i(\bar{T}_i - T_{sat}(P_{s,i}))$ into a once-through steam generator of steam-side volume V_{SG} at pressure $P_{s,i}$. Flow $\dot{m}_i$ reaches the header through an isolation valve of fractional opening $u_{M,i}$. The header holds mass m_h at temperature T_h in volume V_h, and its pressure P_h sets both the saturation temperature in every steam generator and the flow through every user's valve. Three pressure transmitters read the same header with different calibration biases b_i.

9.3.1 Steam Properties

Saturation temperature, saturated vapor enthalpy and saturated liquid enthalpy are quadratic fits over 1.5 to 6 MPa, the band this header visits, as shown in Equation (9.1) with P in MPa and $x = \ln P$.

$$\begin{aligned} T_{sat} &= 180.30 + 41.81\,x + 6.348\,x^2 \\ h_g &= 2767.5 + 21.03\,P - 3.030\,P^2 \\ h_f &= 768.6 + 170.5\,x + 43.5\,x^2 \\ h_f(T) &= 4.19\,T + 3.4 \times 10^{-4}\,T^2 \\ h(P,T) &= h_g(P) + c_p \max(0,\, T - T_{sat}(P)) \end{aligned} \tag{9.1}$$

where c_p is the superheated heat capacity in kJ/kg/K. The fits reproduce saturation temperature to better than 0.15 °C and saturated vapor enthalpy to about 2 kJ/kg, and every lookup clamps its argument to 0.5 to 8 MPa so that an excursion outside the fitted range returns the edge value rather than an extrapolated one. The equation of state is Equation (9.2), and it serves both the header and each steam generator plenum.

$$P = \frac{mZR(T + 273.15)}{1000\,V} \tag{9.2}$$

where Z is a compressibility factor, $R = 0.4615\,\mathrm{kJ/kg/K}$ is the gas constant of steam, m is mass in kg and V is volume in m^3.

Table 9.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts, and the Range of a CV Is That of Its Setpoint.

Symbol	Role	Nominal	Range
P_h	CV, measured	3.4012 MPa	2.90–3.75 MPa
W	CV, measured	132.7 MW	20–150 MW
$\dot{m}_p$	CV, measured	55.7 kg/s	0–95 kg/s
b (bypass)	CV, mid-ranged	3.0 %	0–20 %
$Q_{i,dem}$	MV	190.0/176.7/163.4 MW	50–260 MW
v_t, v_p	MV	70.6, 54.7 %	0–100 %
v_b	MV	3.0 %	0–60 %
v_c, v_d	MV	0 %	0–100 %
grid load	DV, measured	133 MW	40–150 MW
process request	DV, measured	56 kg/s	0–95 kg/s
f, T_{fw}	DV, measured	60 Hz, 149 °C	±0.6, 100–200
$m_{f,i}$ fouling	DV, unmeasured	1.00/0.93/0.86	0.55–1.05
b_i spread	DV, unmeasured	0.015 MPa	0–0.05 MPa
trip i	DV, event	none	any module
Q_i	state	190.0/176.7/163.4 MW	0–300 MW
$\bar{T}_i$	state	302.9/302.1/301.3 °C	120–400 °C
$P_{s,i}$	state	3.802/3.752/3.704 MPa	0.6–8 MPa
m_h, T_h	state	1724 kg, 296.8 °C	$\geq$ 20, 150–430
E	state	112.5 MWh	0–250 MWh
x_k, five valves	state	valve positions	0–100 %
W_{lag}	state	132.7 MW	≥ 0

9.3.2 Module Side

Each module carries three states. Core thermal power follows a rate-limited demand through one first-order lag, as shown in Equation (9.3).

$$\tau_Q \frac{\mathrm{d}Q_i}{\mathrm{d}t} = Q_{i,dem} - Q_i, \qquad \left|\frac{\mathrm{d}Q_{i,dem}}{\mathrm{d}t}\right| \leq \frac{r_i}{100}\frac{Q_{rate}}{60} \tag{9.3}$$

where τ_Q is the closed-loop core power time constant in seconds, r_i is the module's ramp rate in percent of rating per minute and Q_{rate} is the rated thermal power in MW. Reactivity control, rod worth, xenon and the moderator and Doppler feedbacks are all folded into τ_Q and the rate limit. That rate limit is the constraint the whole application is organized around, because the three modules together can move only 31.3 MW of core power per minute and no controller gain changes the number.

The primary loop is the module's integrator, as shown in Equation (9.4).

$$C_{p1}\frac{\mathrm{d}\bar{T}_i}{\mathrm{d}t} = Q_i - Q_{SG,i}, \qquad Q_{SG,i} = UA_i\,(\bar{T}_i - T_{sat}(P_{s,i})) \tag{9.4}$$

where C_{p1} is the heat capacity of coolant, vessel and internals in MJ/°C, $\bar{T}_i$ is the average primary coolant temperature and $UA_i = UA_{SG}\,m_{f,i}$ is the steam generator conductance in MW/°C with $m_{f,i}$ the per-module fouling multiplier. Both sides carry units of MW. The heat sink in (9.4) is the saturation temperature of the module's own steam, which follows the header, so raising header pressure reduces every module's heat removal at the same instant. That single term is why no disturbance on this plant is local, and the ratio $C_{p1}/UA_1 = 180/3.43 = 52.5$ s is the slowest time constant the control system can reach.

The steam side of a once-through generator holds almost nothing, and its pressure is the third state, as shown in Equation (9.5).

$$\begin{aligned} \dot{m}_{gen,i} &= \frac{1000\,Q_{SG,i}}{\max\big(200,\ h(P_{s,i},T_{st,i}) - h_f(T_{fw})\big)} \\ \frac{\mathrm{d}P_{s,i}}{\mathrm{d}t} &= \frac{(\dot{m}_{gen,i} - \dot{m}_i - \dot{m}_{sv,i})\,ZR(T_{st,i}+273.15)}{1000\,V_{SG}} \end{aligned} \tag{9.5}$$

where $\dot{m}_{gen,i}$ is the steam made in kg/s, V_{SG} is the whole steam-side volume of coil, plenum and lead pipe in m^3, $\dot{m}_i$ is the flow to the header and $\dot{m}_{sv,i}$ is safety valve flow. Feedwater flow is set equal to $\dot{m}_{gen,i}$, so there is no feedwater control and no drum level. Outlet superheat is a power law referenced to a design duty rather than an area calculation, as shown in Equation (9.6).

$$\begin{aligned} T_{st,i} &= T_{sat}(P_{s,i}) + \mathrm{clamp}(\theta_i,\,6,\,140) \\ \theta_i &= \Delta T_{SH}\left(\frac{Q_{SH}}{\max(Q_{SG,i},5)}\right)^{0.55} \end{aligned} \tag{9.6}$$

where ΔT_{SH} is the superheat at the reference duty Q_{SH} in MW and θ_i is the unclamped superheat. Superheat falls as duty rises because feedwater flow tracks steam flow against a fixed heat transfer area, and the clamp keeps the correlation inside its fitted band.

Flow from a module to the header is driven by the pressure difference between them, as shown in Equation (9.7).

$$\dot{m}_i = u_{M,i}\,K_{line}\sqrt{10^6\,(P_{s,i}-P_h)\,\rho_i}, \qquad \rho_i = \frac{1000\,P_{s,i}}{ZR\,T^K_{st,i}} \tag{9.7}$$

where $u_{M,i}$ is the fractional isolation valve opening, K_{line} is the line flow coefficient, ρ_i is steam density in $\mathrm{kg/m}^3$ and $T^K_{st,i}$ is the outlet temperature in kelvin. The square root makes the split between modules a hydraulic outcome before it is a control decision, and at the design point each module sits 0.30 to 0.40 MPa above the header. Main steam safety valves are shut in every normal condition and open linearly above their set pressure, as shown in Equation (9.8).

$$\dot{m}_{sv,i} = K_{sv}\max(0,\ P_{s,i} - P_{sv}) \tag{9.8}$$

They are the only outlet an isolated steam generator has for the decay heat still arriving in front of it, which is why an isolated module's pressure settles just above P_{sv} rather than running away.

A scram collapses core power to decay heat and shuts the isolation valve, as shown in Equation (9.9).

$$\begin{aligned} Q_{i,dem} &= 0.065\,Q_{i,0}\left(1 + \frac{t - t_{trip}}{10}\right)^{-0.2} \\ u_{M,i} &= \max\left(0,\ 1 - \frac{t - t_{trip}}{\tau_{MSIV}}\right) \end{aligned} \tag{9.9}$$

where $Q_{i,0}$ is the power the core was making at the trip and τ_{MSIV} is the closure time in seconds. The decay expression is a single-term fit to the Way and Wigner shape of Equation (1.9), giving 6.5 % at once, 4 % after ninety seconds and 2.6 % after a quarter of an hour, and τ_Q is replaced by the faster τ_{scram} while the trip flag is set. A module also trips itself if $\bar{T}_i$ exceeds $\bar{T}_{trip}$.

9.3.3 Header and Users

The header is two states, a mass and a temperature, and its pressure follows from (9.2). The balances are Equations (9.10) and (9.11).

$$\frac{\mathrm{d}m_h}{\mathrm{d}t} = \sum_i \dot{m}_i + \dot{m}_d - \dot{m}_{out} \tag{9.10}$$

$$\begin{aligned} m_h(c_p - ZR)\frac{\mathrm{d}T_h}{\mathrm{d}t} = &\sum_i \dot{m}_i(h_i - u_h) + \dot{m}_d(h_d - u_h) \\ &- \dot{m}_{out}(h_h - u_h) - \dot{Q}_{loss} \end{aligned} \tag{9.11}$$

where h_i is the enthalpy of the steam arriving from module i, h_d is the enthalpy of storage discharge steam, $h_h = h(P_h, T_h)$ is the header enthalpy, $u_h = h_h - ZR(T_h + 273.15)$ is the header specific internal energy in kJ/kg, $\dot{Q}_{loss} = UA_h(T_h - T_{amb})$ is the pipe heat loss in kW with UA_h in kW/°C, and $\dot{m}_{out}$ is the total draw of Equation (9.14). Both sides of (9.10) are in kg/s and both sides of (9.11) are in kW.

Every user is a valve on the same pressure, in the critical-flow form of Equation (9.12).

$$\dot{m}_k = \frac{x_k}{100} K_k \frac{P_h}{\sqrt{T_h + 273.15}}, \qquad k \in \{t, p, b, c\} \tag{9.12}$$

where x_k is the valve position in percent and K_k is the coefficient of the turbine governor, the process valve, the bypass or the storage charge valve. Equation (9.12) is the second coupling term in the model, because it makes every user's draw proportional to header pressure at a fixed valve position. The users therefore self-regulate a little, and the industrial customer's flow falls exactly when its pressure requirement is already in trouble. Storage discharge is driven the other way, as shown in Equation (9.13).

$$\begin{aligned} \dot{m}_d &= \frac{x_d}{100} K_d \max(0,\ P_{tes} - P_h) \\ P_{tes} &= P_{tes}^{lo} + \left(P_{tes}^{hi} - P_{tes}^{lo}\right)\frac{E}{E_{max}} \end{aligned} \tag{9.13}$$

where E is the stored energy in MWh, so the delivery pressure falls with the charge and a nearly empty tank cannot push steam into a healthy header at all. The relief valve is the linear form of (9.8) about P_{rel}, and the total draw is Equation (9.14).

$$\dot{m}_{out} = \dot{m}_t + \dot{m}_p + \dot{m}_b + \dot{m}_c + \dot{m}_{rel} \tag{9.14}$$

Generator output is a polytropic expansion to the condenser through one combined efficiency, followed by a first-order output lag, as shown in Equation (9.15).

$$\begin{aligned} w_e &= \eta_t\, c_p(T_h + 273.15)\left[1 - \left(\frac{P_{cond}}{P_h}\right)^{\kappa}\right] \\ \tau_{gen}\frac{\mathrm{d}W_{lag}}{\mathrm{d}t} &= \frac{\dot{m}_t w_e}{1000} - W_{lag} \end{aligned} \tag{9.15}$$

where w_e is the specific work in kJ/kg, η_t absorbs the isentropic, wetness, mechanical and generator losses, and $\kappa = (\gamma - 1)/\gamma$. At the design point (9.15) returns 783 kJ/kg, so 169.5 kg/s through the governor valve makes 132.7 MW. Process heat and condenser loading are Equation (9.16).

$$\dot{Q}_p = \frac{\dot{m}_p\,(h_h - h_{ret})}{1000}, \qquad L_{cond} = \frac{0.9\,\dot{m}_t + 1.35\,\dot{m}_b}{\dot{m}_{cond}^{max}} \tag{9.16}$$

where h_{ret} is the condensate return enthalpy. The coefficients in (9.16) state that a kilogram of bypass steam has done no work on the way and loads the condenser about a third harder than a kilogram of

turbine exhaust, so the bypass is not a free place to put steam. Storage charge and discharge move the energy state through two fixed efficiencies, as shown in Equation (9.17).

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \frac{1}{3600}\left[\frac{\dot{m}_c\,(h_h - h_f(T_{fw}))\,\eta_c}{1000} - \frac{\dot{m}_d\,(h_d - h_f(T_{fw}))}{1000\,\eta_d}\right] \tag{9.17}$$

where η_c and η_d are the recovery efficiencies and 3600 converts MW to MWh/s. Storage steam is barely superheated at $\Delta T_{SH,d}$ above saturation, because a kettle boiler makes saturated steam, so a hard discharge cools the header. Each of the five valves follows its command through a first-order stroking lag, as shown in Equation (9.18).

$$\tau_{V,k}\,\frac{\mathrm{d}x_k}{\mathrm{d}t} = u_k - x_k \tag{9.18}$$

9.3.4 Instrument Chain

The controller reads three instruments, each a first-order filter of the true value written with the exact zero-order-hold update of Equation (9.19).

$$\hat{y}_k = \alpha\,\hat{y}_{k-1} + (1-\alpha)\,y_k, \qquad \alpha = \mathrm{e}^{-\Delta t/\tau} \tag{9.19}$$

with $\tau = \tau_{PT}$ for header pressure, τ_{WT} for generator output and τ_{FT} for process steam flow. Optional colored noise is added to the header pressure before the filter, produced by passing a Gaussian sequence through a first-order filter and rescaling so that the stationary standard deviation equals the specified value. The three module controllers do not read the same number, as shown in Equation (9.20).

$$\hat{P}_{h,i} = \hat{P}_h + b_i, \qquad b = [\,\sigma_b,\ 0,\ -\sigma_b\,] \tag{9.20}$$

where σ_b is the transmitter calibration spread. Its default of 0.015 MPa is four tenths of one percent of a 3.4 MPa reading, which is a good calibration and not a fault, and Section 9.6 measures what it does.

9.3.5 Parameters, Initialization and Numerics

The constants are listed in Tables 9.2 and 9.3, of which four are hidden and never displayed. Every other constant is exposed as an adjustable parameter, and every disturbance in Table 9.1 as a slider or a discrete event.

Every scenario opens at a genuine steady state, and finding it is not a marching problem, because raising one module's demand raises its steam generator pressure, which raises the header, which reduces every other module's flow and every user's draw at the same time. The equilibrium is solved as the two residuals of Equation (9.21) in the two header states, with each module's steam pressure found by bisection inside.

$$\begin{aligned} r_1(P_h,T_h) &= \sum_i \dot{m}_i + \dot{m}_d - \dot{m}_{out} \\ r_2(P_h,T_h) &= \sum_i \dot{m}_i(h_i - u_h) + \dot{m}_d(h_d - u_h) \\ &\quad - \dot{m}_{out}(h_h - u_h) - \dot{Q}_{loss} \end{aligned} \tag{9.21}$$

The outer solve is a damped Newton iteration on a numerical 2×2 Jacobian, relaxed to 55 % for the first four passes and converged to $|r_1| < 10^{-10}$ and $|r_2| < 10^{-7}$. The inner bisection drives $\dot{m}_{gen,i} - \dot{m}_i - \dot{m}_{sv,i}$ to zero over 70 halvings, which is monotone in $P_{s,i}$ and therefore safe, and with the isolation valve shut the root lands just above P_{sv}, where an isolated module belongs. The same solver supplies the manual aid that sets the three demands to those that put the header on setpoint, using a secant iteration on total demand rather than a fixed gain because the header gain moves with the

Table 9.2 Module and Header Parameters. Hidden Constants Are Never Displayed.

Symbol	Meaning	Value	Unit
Q_{rate}	module rated power	250	MW
Q_i^{lic}	licensed power, U1/U2/U3	250/240/232	MW
r_i	ramp rate, U1/U2/U3	5.0/3.0/4.5	%/min
UA_{SG}	clean SG conductance	3.43	MW/°C
C_{p1}	primary heat capacity	180	MJ/°C
τ_Q	core power lag	14	s
τ_{scram}, τ_{MSIV}	scram, valve closure	1.6, 6	s
$\bar{T}_{max}$, $\bar{T}_{trip}$	coolant limit, trip	315, 324	°C
V_{SG}	SG steam-side volume	30	m^3
K_{line}	steam line coefficient	0.032473	—
P_{sv}	safety valve set pressure	4.60	MPa
ΔT_{SH}	superheat at reference duty	57	K
V_h	header volume	120	m^3
P_h^0, P_{rel}	design, relief pressure	3.40, 3.85	MPa
UA_h, T_{amb}	heat loss, ambient	2.2, 25	—
Z, c_p	compressibility, heat cap.	0.90, 2.65	—
δ_P	pressure tolerance band	0.08	MPa
$\tau_{PT}, \tau_{FT}, \tau_{WT}$	transmitter responses	0.4, 1.2, 0.8	s
hidden			
K_{sv}	safety valve capacity	150	kg/s/MPa
Q_{SH}	reference duty for (9.6)	176.5	MW
K_{rel}	relief valve capacity	400	kg/s/MPa
f_0	nominal frequency	60	Hz

operating point. It is 6.8 kPa per MW at the design point and 13.6 kPa per MW on a plant at half load held at the same pressure.

The engine of Section 1.6.1 integrates the model by fourth-order Runge-Kutta at a fixed 0.05 s step, chosen because the fastest state is a steam generator pressure with a time constant near 1 s. The control system executes every 0.25 s, a realistic distributed control system scan, the historian logs four samples per simulated second, and the noise seed is 20260906. Every state is clipped after each step to the range given in Table 9.1.

9.3.6 Assumptions

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

No neutronics. Equation (9.3) is a first-order lag on a rate-limited demand, so reactivity control, rod worth, xenon and the moderator and Doppler feedbacks are all inside one 14 s time constant. That is adequate for a load-sharing study and useless for anything about reactivity, so the application cannot show a reactivity-limited maneuver or a xenon-constrained restart. *One lumped steam generator conductance.* A real helical-coil once-through unit has economizer, boiling and superheat sections whose boundaries move with load [12], and its response to a feedwater step is genuinely different from the single lag of (9.4). *Feedwater flow is not independent.* It equals the steam made, so there is no feedwater control, no drum level because there is no drum, and no feedwater pump as a constraint, and coupled feedwater control across several once-through generators is itself a substantial problem [11]. *Natural circulation is asserted rather than modeled.* There is no primary flow state and therefore no flow-limited power, which is a real constraint on a natural-circulation module.

Steam is an ideal gas. Equation (9.2) carries $Z = 0.90$ and a constant superheated heat capacity, and outside roughly 1.2 to 6.5 MPa the property fits are extrapolation and every lookup clamps. *The header is a single lumped volume*, with no pressure drop, no distinct conditions at each take-off and

Table 9.3 User, Storage, Valve, Runback and Economic Parameters. Where a Row Carries Several Symbols the Values Are Listed in the Same Order and the Units Are Those Named in the Meaning Column or Given in the Text.

Symbol	Meaning	Value	Unit
K_t, K_p, K_b	governor, process, bypass valve	1685.31, 714.98, 702.21	—
η_t, κ	turbine efficiency, exponent	0.683, 0.23	—
P_{cond}, τ_{gen}	condenser pressure, output lag	0.007, 1.5	—
$\dot{m}_{cond}^{max}$	condenser steam capacity	200	kg/s
h_{ret}	condensate return enthalpy	440	kJ/kg
P_p^{min}, $\dot{m}_p^{min}$, t_{grace}	industrial user limits	2.90, 34, 45	—
E_{max}, $E(0)/E_{max}$	storage capacity, initial charge	250, 45	—
P_{tes}^{lo}, P_{tes}^{hi}	delivery pressure empty, full	3.05, 5.40	MPa
K_c^{chg}, K_d	charge, discharge coefficient	280.89, 42.5	—
η_c, η_d	charge, discharge recovery	0.96, 0.94	—
$\Delta T_{SH,d}$	superheat of storage steam	10	K
$\tau_{V,t}$, $\tau_{V,b}$, $\tau_{V,p}$, $\tau_{V,c}$	valve stroking times	0.6, 2.0, 4.0, 5.0	s
R, W_{rated}	governor droop, rating	5 %, 150	—
δ_{fast}, δ_{trim}	runback pick-up, trim band	0.10, 0.02	MPa
K_{shed}, $K_{shed,p}$	runback droop gains	220, 260	—
K_{trim}, $K_{trim,p}$	runback trim rates	0.60, 7	—
c_E, c_Q	electricity, steam price	42, 28	$/MWh
c_{int}, c_{trip}	interruption penalty, loss cost	9000, 40 000	—

no condensate, and two users at opposite ends of a real header do not see the same pressure, which would make the multi-controller failure of Section 9.6 worse rather than better. *The turbine has no shaft dynamics*, so a grid load rejection here is a steam-balance event rather than the mechanical one it would really be. *Storage is a state of charge*, with no thermocline, no temperature limits and no self-discharge, so the application overstates how neatly a store can be dispatched. *Two constraints stop a module, and the economics are four prices.* Only the average coolant temperature limit and the licensed power appear, where a real protection system carries a dozen more, and the prices exist to make the trade between shedding electricity and shedding a customer concrete rather than to value a contract.

9.4 Coupling and Capability

9.4.1 A Fast Buffer Between Slow Plants, and Six Actuators

Three time scales decide everything the control system can do, and all three were measured on the shipped model at the design point. The header holds 1724 kg of steam and passes 228.2 kg/s, which is 7.55 s of residence, and the steam side of a module holds about 470 kg against an average 76.1 kg/s, which is 6.18 s, so the inventory between a reactor core and a customer's reboiler is under

fourteen seconds of throughput. Against that, the primary loop time constant of (9.4) is 52.5 s and the aggregate ramp capability of (9.3) is 31.3 MW/min. The coupling is an order of magnitude faster than anything thermal that could answer it, and that ratio is the whole control problem. A recirculating boiler with a drum would hold thirty times the module inventory and buffer the header for a minute, which is why the parallel-boiler practice of Section 9.2 could tolerate a slow master [10].

Header pressure has six actuators, and every architecture that follows is a decision about which of them answers a given error. Stepping each actuator from the design point in manual mode, 10 s into a run, and running to a new steady state gives the steady gain and the time to 63 % of the final pressure change, counted from the step. A station moved before Run only sets the steady state the plant starts from. Unit 1 from 190 to 210 MW moves the header by $+0.1367$ MPa with $t_{63} = 195$ s, and all three demands together by 30 MW move it by $+0.2052$ MPa. The governor valve from 70.6 to 78.0 % moves it by -0.2578 MPa in 32 s, and the bypass from 3 to 13 % by -0.1501 MPa in 36 s.

Two readings of those numbers matter. The header gain is 6.84 kPa per MW of module duty from the single-module step and 6.84 kPa per MW from the three-module step, so it is a property of the header rather than of which module moved. The two valve times are about 4 s apart, because a pressure change on this header settles over about half a minute whatever valve moved, so what separates the actuators is the speed at which the flow itself changes. A bypass valve strokes in 2 s and passes 1.00 kg/s per percent, so closing a bypass resting at 3 % puts 3 kg/s back into the header inside two seconds. A reactor demand takes 96 s to move 20 MW at 12.5 MW/min before the 14 s power lag and the 52.5 s primary loop have even started, and the measured t_{63} of 195 s is set by those three in series. Everything fast on this plant is a valve and everything free is a reactor.

The single-module step also measures the coupling directly. Raising Unit 1 from 190 to 210 MW raises its own steam generator pressure by 0.204 MPa and, through the header, raises Unit 2 by 0.127 MPa and Unit 3 by 0.128 MPa. The flows of the two untouched modules dip while the header rises and then return to within 0.01 kg/s of where they started, because in steady state a module's steam production is its own heat balance and nothing else. Their margin does not return. The saturation temperature underneath each has risen with its steam pressure, so (9.4) settles at a higher $\bar{T}_i$ for the same duty, and Unit 2 loses 1.95 °C of margin while Unit 3 loses 1.99 °C without either being touched. Loading one module spends thermal margin belonging to all three.

9.4.2 The Ceiling Is Not the Rating, and It Moves

What actually stops a module is the average coolant temperature limit, not the licensed power, as shown in Equation (9.22).

$$Q_i^{max} = \min\Big(Q_i^{lic},\ UA_i\,(\bar{T}_{max} - T_{sat}(P_{s,i}))\Big) \tag{9.22}$$

The physical ceiling in Equation (9.22) is distinct from what the controller can estimate. The allocator and the demand clamps use

$$\begin{aligned}\widehat{Q}_i^{max} = \max\big(0, \min\big(Q_i^{lic},\\ \widehat{Q}_{SG,i} + UA_{i,0}(\bar{T}_{max} - \bar{T}_{i,meas})\big)\big).\end{aligned} \tag{9.23}$$

The primary heat balance estimates $\widehat{Q}_{SG,i}$ from the commanded core response and the change in measured average coolant temperature. Conductance, primary heat capacity and core response times are held at their declared design values. Hidden fouling and steam generator pressure never enter the allocator or clamp. The same estimate limits all five structures. With changed conductance, the estimate can differ from the physical ceiling during a transient and approaches it as the measured temperature approaches its limit. In a frozen-instrument regression, a hidden fouling and parameter change moved the former clamp by 57.31 MW and moves the corrected clamp by zero.

At the design point the three ceilings are 232, 218 and 204 MW against licensed powers of 250, 240 and 232 MW, and the plant has 123 MW of thermal reserve. Every ceiling falls as the header

rises, and it falls faster for a fouled module because a smaller UA_i converts the same duty into a larger temperature rise. The consequence is a step test that stops short of its own request. Asking Unit 1, licensed for 250 MW, to make 230 MW produces a demand ramp at 12.5 MW/min that overshoots to 222.6 MW at 240 s with $\bar{T}_1 = 311.4\,^{\circ}\mathrm{C}$ and then settles back to 217.3 MW with $\bar{T}_1$ resting exactly on the 315.00 °C limit, reading 217.5 MW at the 480 s instant the scenario removes the step. Loading it raised its own steam generator pressure from 3.802 to 4.081 MPa and the header from 3.401 to 3.588 MPa, and every one of those kilopascals raised the saturation temperature the primary has to work above. The module ends 12.7 MW short of the request and 32.7 MW short of its license, held there by a constraint that appears in no slider.

9.4.3 Why Allocation by Available Duty Equalizes Margin

The allocation rule follows from (9.4) and (9.22) directly. In steady state $Q_i = UA_i(\bar{T}_i - T_{sat})$, so the margin of module i is Equation (9.24).

$$\bar{T}_{max} - \bar{T}_i = \bar{T}_{max} - T_{sat} - \frac{Q_i}{UA_i} = \frac{Q_i^{max} - Q_i}{UA_i} \tag{9.24}$$

where Q_i^{max} takes the conductance branch of (9.22). Dividing a total demand D in the ratio of the ceilings, $Q_i = D\,Q_i^{max} / \sum_j Q_j^{max}$, and substituting into (9.24) gives Equation (9.25).

$$\bar{T}_{max} - \bar{T}_i = (\bar{T}_{max} - T_{sat}) \left(1 - \frac{D}{\sum_j Q_j^{max}} \right) \tag{9.25}$$

The right side of (9.25) carries no index i, so the rule leaves every module the same number of degrees from its own limit, which is the only sense in which a load share on this plant can be called fair. The equality is exact only where every module boils against the same saturation temperature, and because the three generators sit at slightly different pressures the measured margins under coordinated control at the design demand are 12.8, 12.9 and 13.0 °C, a spread of 0.25 °C. Evaluating the algebra on Unit 1 gives $(232/3.43)(1 - 530/654) = 12.8\,^{\circ}\mathrm{C}$, which matches the measurement. The rule also has a failure mode built into (9.22), because a module whose licensed power binds no longer has a ceiling proportional to its conductance and the equality is lost.

9.5 Control Architectures

9.5.1 The Common Loop Algorithm

Every loop is the shared velocity-form proportional-integral controller of Section 1.6.2, executed every $\Delta t_c = 0.25\,\mathrm{s}$ and written in this chapter's symbols as Equations (9.26) to (9.29).

$$e_k = a\left(y_{\mathrm{SP},k} - \hat{y}_k\right), \qquad P_k = K_c e_k \tag{9.26}$$

$$s_k = s_{k-1} + P_k - P_{k-1} \tag{9.27}$$

$$c_k = u_{bias} + s_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k, u_{min}, u_{max}) \tag{9.28}$$

$$I_k = I_{k-1} + \frac{K_c}{\tau_I} e_k \Delta t_c + \frac{\Delta t_c}{T_t} (u_k - c_k) \tag{9.29}$$

where $a = \pm 1$ is the action, s_k is the running sum S_k of Equation (1.14), accumulated without limit, c_k is the demanded output v_k of Equation (1.15), u_k is the output applied, and T_t is the anti-windup tracking time, which is τ_I on every loop here. No loop in this application has a derivative time and none is exposed, so the 0.3 s derivative filter configured on the two user loops is inert. Anti-windup is the back-calculation of Section 1.6.2, and s_k is never truncated [18]. The reactor ramp limit of (9.3) is applied after the master and valve position controller outputs and lies outside this mechanism, so those integrals keep integrating while the reactors ramp toward their targets, bounded only by the sums of the module floors and ceilings. Manual tracking and the copy of each automatic output to its

Table 9.4 The Eight Implemented Loop Types, Their Pairings, Directions and Default Tunings. Action +1 Is Reverse Acting and −1 Is Direct Acting. Sums Run Over the Modules Still in Service.

CV → MV	Limits	K_c	τ_I (s)	a
$P_h \to D$ (master)	$\sum Q_i^{min}$–$\sum Q_i^{max}$	320	240	+1
$P_h \to Q_{i,dem}$ (per module)	50–Q_i^{max}	950	120	+1
$P_h \to v_b$ (bypass)	0–60 %	260	25	−1
$b \to D$ (position)	$\sum Q_i^{min}$–$\sum Q_i^{max}$	3.0	240	+1
$W \to v_t$	0–100 %	0.35	20	+1
$\dot{m}_p \to v_p$	0–100 %	0.50	15	+1
$P_h \to v_d$ (discharge)	0–100 %	700	40	+1
$P_h \to v_c$ (charge)	0–100 %	600	90	−1

manual station work as Section 1.6.2 describes, every loop not in service in the selected architecture is tracked the same way, and so bumpless transfer and architecture switching both fall out of the tracking. The direction of each loop follows the sign of its process gain and appears as the ISA action flag in Table 9.4, in the convention of Section 1.3.

The two pressure-to-demand loops in Table 9.4 are reverse acting because more reactor power raises the header, and the bypass loop is direct acting because opening the bypass lowers it. The per-module loops and the master work against the same process gain. Section 9.4 measures 6.84 kPa per MW from the single-module step and 6.84 kPa per MW from the three-module step, because the header answers the total duty and not the identity of the module that supplied it, so a per-module controller needs the master gain and no more. The default per-module gain of 950 MW/MPa against a master gain of 320 MW/MPa is three times that. Whatever value is chosen, three such controllers act on the same error at the same instant, so the loop gain the plant sees is three times the loop gain any one of them was tuned for. That multiplication, and not the numerical default, is the structural fault.

9.5.2 Five Architectures

The architecture selector chooses among five answers a plant can give to the question of who moves when header pressure moves.

Manual. The operator sets three module demands and five valve positions, and two manual aids compute the demands that put the header on setpoint and the valve positions that deliver the user setpoints at the present pressure.

Independent. Each module runs its own controller from Table 9.4 on its own transmitter reading $\hat{P}_{h,i}$ of (9.20). Three integral controllers share one measurement, so the loop gain is three times what any of them was tuned for, and because the three transmitters disagree the three integrators chase three targets that cannot all be reached.

Master-slave. One pressure master produces a single total thermal demand D and a participation rule divides it, so only one integrator sees the pressure error and the conflict stops. The reactors remain inside the pressure loop, so the speed of the loop is the speed of the reactors.

Coordinated. Three layers act on different time scales. The bypass controller takes the pressure error directly and its valve strokes in 2 s, a valve position controller then moves the total demand D to walk the bypass back toward its resting position b_{ref} over about four minutes, and the allocator divides D by available duty. Where thermal storage is in service the discharge controller is offered the setpoint 0.04 MPa low and the charge controller 0.04 MPa high, so the store acts only outside a dead band and never fights the master. The valve position controller watches the bypass controller's unsaturated demand c_k of (9.28) rather than the clamped valve, because a bypass clamped shut at zero would otherwise report zero no matter how much steam was missing and the slow layer would never learn it needs to make more. Volume 4, Chapter 10 builds the same kind of hierarchy for a

power block, with a turbine bypass that acts in seconds and discards what it removes, mid-ranged against slower actuators that cost nothing.

Pause. Try Exercise 9.3 before reading on.

MPC. One optimization takes the place of the three coordinated layers, setting every second the ramp of each module and the bypass command over the next two minutes, with the coordinated loops held in reserve as its fallback. Section 9.5.5 writes it out.

9.5.3 The Allocator and Its Spill Pass

Given a total demand D and a participation shape s_i, the allocator computes the per-module target of Equation (9.30).

$$Q_i^{tgt} = \mathrm{clamp}\left(\frac{D\, s_i}{\sum_{j \in \mathscr{F}} s_j},\; Q_i^{min},\; Q_i^{max} \right), \qquad i \in \mathscr{F} \tag{9.30}$$

where $\mathscr{F}$ is the set of modules that have not tripped, $Q_i^{min} = \min(0.20\, Q_{rate},\, Q_i^{max})$ and Q_i^{max} is (9.22). The clamping discards demand, so the calculation runs for up to four passes, each of which sums the spill $\sum_i (Q_i^{raw} - Q_i^{tgt})$, removes the modules now against a limit from $\mathscr{F}$, and redistributes the remainder over those that still have room. A share a module cannot take is not a share, it is a shortfall some other module has to be told about, and an allocator that does not spill is the second most common way plantwide load sharing fails after not having one at all. Three shapes are implemented. Equal gives $s_i = 1$. Margin-proportional gives $s_i = Q_i^{max}$, which is the rule of (9.25). Base and swing holds the base-loaded modules at a fixed demand and gives the first module still running the whole of the variation, which puts the entire ramp-rate burden on one reactor. The target is applied through the rate limit of (9.3), so the allocator decides the destination and the ramp rate decides the path. Volume 4, Chapter 12 divides one demand among three actuators with different prices and limits, the same allocation problem on a vehicle.

9.5.4 The Two-Stage Runback

Neither the architecture nor the allocator creates steam, so a plant that has lost a module needs a way to stop selling something. The runback is two proportional stages on the same measurement, both one-way, as shown in Equations (9.31) and (9.32).

$$\begin{aligned} d_{fast} &= \max\left(0,\; P_{h,\mathrm{SP}} - \delta_{fast} - \hat{P}_h\right), \\ d_{slow} &= \max\left(0,\; P_{h,\mathrm{SP}} - \delta_{trim} - \hat{P}_h\right) \\ \Delta W &= \mathrm{clamp}\left(-K_{shed}\, d_{fast} + \Sigma_{trim},\; -(W_{rated} - 20),\, 0\right) \end{aligned} \tag{9.31}$$

$$\Sigma_{trim} \mathrel{-}= K_{trim}\, d_{slow}\, \Delta t_c \tag{9.32}$$

where ΔW is a bias added to the generator load setpoint. The fast stage is proportional and instantaneous, so it arrests a falling header and returns the load as the pressure recovers, and it is the only thing on the plant that can answer a module trip inside the industrial customer's 45 s ride-through. The slow trim accumulates while the header is outside a narrow dead band and never restores load, because after a trip the surviving reactors reach their ceilings with the pressure still short of setpoint and no master action closes a gap where the steam does not exist. A runback is a latch, and recovering the megawatts is an operator decision made on the load setpoint. The process variant is identical with $K_{shed,p}$ and $K_{trim,p}$ acting on the process flow setpoint, floored a little above $\dot{m}_p^{min}$ so that the runback cannot itself shut the customer down. Chapter 5 meets the opposite event, a load that disappears, with a reactor runback among its last layers.

9.5.5 Model Predictive Control

The fifth architecture is a model predictive controller (MPC), a receding-horizon quadratic program [19, 20] in the arrangement of Section 1.6.3, posed every $T_s = 1$ s, each fourth control scan,

over $N = 120$ steps grouped into twelve move blocks that lengthen from a single step near the present to 75 steps at the far end. It decides four inputs, the ramp rate v_i in MW/s of each running module's demand and the bypass command u_b, and between decisions each demand advances by $v_i \Delta t_c$ per scan. Because the decision is a rate and not a level, the limit of (9.3) becomes the exact hard bound $|v_i| \le \bar{v}_i = r_i Q_{rate}/6000$ on every step of the plan whatever the block lengths, and u_b is held hard to 0–60 %.

The controller's model is the balance of Section 9.3 written out again with the declared values of Tables 9.2 and 9.3, the declared fouling multipliers 1.00, 0.93 and 0.86 and 149 °C feedwater, never the run's, so a fouled generator or an edited parameter is model mismatch. Its 17 states are the three demands, core powers, coolant temperatures and steam generator pressures of (9.3) to (9.7), the header mass and temperature of (9.10) and (9.11), the bypass position, the header transmitter filter of (9.19), and one users' draw w_u in kg/s standing for the turbine, process and storage flows and held constant over the horizon. Each decision linearizes it by central differences about the present estimate and discretizes it with an exact hold and its affine term. The turbine and process loops, the run's fouling and feedwater temperature, transmitter noise and bias, the runback, the trips and the condenser are left out. Stepped open loop from the design point for 130 s, the model overstates the header's response by about a third, 0.075 MPa against the plant's 0.056 MPa for +10 MW on Unit 1 and −0.190 against −0.151 MPa for +5 % of bypass, because the users' valves pass more steam as the header rises until their own loops close them back through (9.12).

Each plan minimizes Equation (9.33).

$$\begin{aligned} \mathscr{J} = \sum_{k=1}^{N} \Big[& q_P \big(\hat{P}_{h,k} - P_{h,\mathrm{SP}}\big)^2 + q_s \sum_i \Big(Q_{i,k} - \sigma_i \sum_j Q_{j,k} \Big)^2 \\ & + \sum_i \frac{\lambda\, v_{i,k}^2}{\bar{v}_i^2} + q_b \big(u_{b,k} - b_{ref}\big)^2 \Big] + \sum_{\ell} \big(w_\varepsilon \varepsilon_\ell^2 + w_\varepsilon \varepsilon_\ell \big) \\ & + \sum_j \Big[\sum_i \frac{\lambda\, \Delta v_{i,j}^2}{\bar{v}_i^2} + s_b\, \Delta u_{b,j}^2 \Big] \end{aligned} \tag{9.33}$$

Here $\hat{P}_{h,k}$ is the predicted transmitter reading, Δ is a change between blocks, and σ_i divides the running total in the ratio of each module's declared available duty $UA_i(\bar{T}_{max} - T_{sat}(\hat{P}_{s,i}))$ capped at its licensed power, the rule of (9.25) evaluated on declared values. The weights are $q_P = 10^4$ per MPa^2, $q_s = 0.01$ per MW^2, $q_b = s_b = 0.02$ per %^{2} and the ramp price $\lambda = 3$, where $\lambda = 1$ would price a full-rate ramp held for a second like 0.01 MPa of header error held for a second. One slack ε_ℓ per softened output, judged every second step and priced at $w_\varepsilon = 10^4$ per °C or MW both quadratically and linearly, softens two limits on each running module, the predicted coolant temperature $\bar{T}_{i,k} \le \bar{T}_{max}$ computed with the declared conductance and the demand between $0.2\,Q_{rate}$ and Q_i^{lic}. A temperature estimate already above $\bar{T}_{max}$ is brought back along $\bar{T}_{max} + (\hat{\bar{T}}_{i,0} - \bar{T}_{max})\,\mathrm{e}^{-kT_s/60\,\mathrm{s}}$, where $\hat{\bar{T}}_{i,0}$ is the present estimate. Each plan is a dense problem in 48 decisions solved by a dual active-set method [21], in a mean 3.80 ms in Chrome on scenario 11, 0.38 % of one core at real time.

A steady-state Kalman filter on the same model, linearized at the declared design point at the 0.25 s scan, updates the estimate every scan under every architecture from the filtered header transmitter alone, the only instrument the pressure loops read, with a noise standard deviation of 0.002 MPa and w_u a random walk of variance $\rho_K (0.05\,\mathrm{kg/s})^2$ per scan, $\rho_K = 1$, which removes steady offset [22]. Neither the MPC nor Coordinated has a load feedforward. One tuning pass on scenario 11 over nine values of λ from 0.01 to 100 and six of ρ_K from 1 to 10^5 admitted only settings that rejected no plan, lifted no relief, kept the customer, tripped no module beyond the scripted one and moved neither the modules nor the bypass more than Coordinated did on the same day, and $\lambda = 3$ with $\rho_K = 1$ was the only one admitted.

Table 9.5 Scenarios and What Each Exercises

#	Scenario	Mechanism	Kind
1	Three reactors and one pipe	header gain, coupling, actuator speed	Experiment
2	One module, one step	rate limit, moving ceiling	Experiment
3	Independent pressure control	three integrators, one measurement	Demonstration
4	Master–slave load sharing	one integrator, reactor-limited speed	Experiment
5	Coordinated allocation	mid-ranging, the resting position	Experiment
6	A large turbine load ramp	fast actuator in the loop	Experiment
7	The chemical plant starts another train	interaction on a common pressure	Experiment
8	The store absorbs the ramp	a fourth participant, and its limits	Experiment
9	Equal megawatts is not equal margin	allocation against a constraint	Experiment
10	Unit 1 scrams	N−1 dispatch and the runback	Demonstration
11	Three reactors and one plan	one optimization against three layers	Experiment
12	Ride through the trip	the dispatch decision, one event	Challenge

The coordinated loops track the applied commands every scan with no preload, so a plan that the solver does not return optimal, finite and inside its bounds is discarded, and Coordinated then holds the demands and the bypass, taking them over with no bump and keeping them until the solver again returns a valid plan. The scram, decay heat and isolation of (9.9), the high-temperature trip, the relief and the runback read no architecture and act above the plan, which after a scram carries the tripped module's decay heat and shut valve as known inputs, pins its ramp at zero and goes on planning the survivors. Every applied demand also passes through the clamp of (9.30) at the ceiling of (9.22), the same clamp the other architectures apply, and that clamp is the estimate of Section 9.4.2, built from the measured coolant temperatures, which follow a fouled module where the plan's declared model does not. The simulated plant, its instruments and the controller's model are thus three separate things.

9.6 Results

Twelve scenarios are arranged as one laboratory sequence, listed in Table 9.5 with the mechanism each exercises. Each fixes the starting mode, the architecture, the split rule, the tuning constants, the disturbances and any timed events, and every figure below was measured by driving the shipped model headlessly from those definitions at the 0.05 s physics step. Each case is credited by the completion rule of Section 1.6.4. The challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

9.6.1 Independent Control With No Disturbance

Scenario 3 holds the plant at the design point with every module running its own pressure controller and nothing disturbing it. The three module powers separate immediately and do not come back, as shown in Figure 9.2. At 600 s the modules are making 132, 202 and 198 MW, a spread of 70.5 MW, and at 1200 s they are at 120, 212 and 198 MW, a spread of 91.7 MW, while header pressure over the whole run stays between 3.385 and 3.401 MPa, a band of 0.0162 MPa inside a tolerance of ±0.08 MPa. The scorecard sees the divergence only through the margin spread, which passes 8 °C at 135 s.

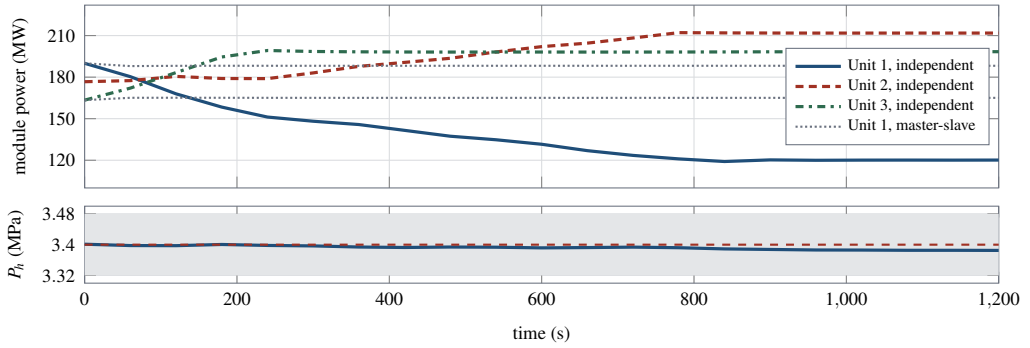

Figure 9.2 Independent per-module pressure control diverging with no disturbance applied, scenario 3. Upper panel: Unit 1 solid, Unit 2 dashed and Unit 3 dash-dotted under independent control, and Units 1 and 3 under master-slave as gray dotted lines. The three module powers separate to 70.5 MW at 600 s and 91.7 MW at 1200 s and stop only when Units 2 and 3 reach their thermal ceilings, while master-slave on the same plant holds a constant 23.2 MW spread. The lower panel shows header pressure over the same run (solid) against its setpoint (dashed) and the shaded ±0.08 MPa tolerance band.

> **TRY IT IN THE APP** **3 · Independent Pressure Control, Fighting Itself**
>
> Load scenario 3 and press Run, at 20× or 60× speed. Nothing disturbs the plant. Watch the three module powers on the Modules tab walk apart, 70.5 MW at ten minutes and 91.7 MW at twenty. Then look at the header pressure trend, which never leaves its band.

The mechanism is (9.20). Unit 1 reads the header 0.015 MPa high, so its integrator winds down, Unit 3 reads it 0.015 MPa low, so its integrator winds up, and Unit 2, whose transmitter happens to be correct, holds the balance. The spread peaks at 93.5 MW at 820 s and the powers then stay within about 1 MW of their final values rather than walking on forever, for a reason worth noticing. Header pressure settles at 3.385 MPa, which is exactly the value Unit 1's transmitter reports as 3.400 MPa, so Unit 1's error reaches zero and its integrator stops, while Units 3 and 2 keep winding up until they reach the Q_i^{max} limits of (9.22), at 196 and 758 s, and stop at 198.4 and 211.8 MW. The final state is a plant with two modules pinned against their thermal ceilings, one at 120 MW, and a controlled variable that never left its band.

Pause. Try Exercise 9.2 before reading on.

Master-slave on the identical scenario holds a spread of 23.2 MW and a header band of 0.0018 MPa, and that 23.2 MW is not a defect but what the margin-proportional rule of (9.25) intends, because the three modules genuinely have different capability. Removing the transmitter spread does not repair the independent architecture, it changes the symptom. With $\sigma_b = 0$ the same three controllers drive header pressure between 2.693 and 3.917 MPa, a band of 1.225 MPa, because three integral controllers acting on one error deliver three times the loop gain any of them was tuned for. No tuning repairs either version, and that is the point. One controlled variable gets one integrator, and the split is a separate decision made by a separate mechanism [1, 3]. Volume 6, Chapter 7 makes the same point on a reactor and column, where the control structure and not the tuning decides how the plant answers a production change.

9.6.2 One Integrator Is Not Fast, It Is Only Consistent

Scenario 4 puts the master in charge and asks the grid for 8 MW more electricity at 90 s, which is about 24 MW of core power that does not exist yet, with the reactors inside the pressure loop. The header falls to 3.2771 MPa, 0.123 MPa below setpoint, and re-enters a ±0.02 MPa band for good

Table 9.6 Measured Header Excursions on Three Upsets. Scenario 6 Steps the Generator Down 23 MW at $t = 60$ s and Raises It 35 MW in Steps of About 9 MW a Minute Apart From $t = 480$ s. Scenario 7 Steps Process Steam 56 to 68 kg/s and Back at 600 s. Scenario 8 Repeats Scenario 6 With the Bypass Resting at 2 % and the Store In and Out of Service. Ranges Are Over the Whole Run After $t = 30$ s, to 900, 780 and 900 s. MPC Rows Use Scenario 11's Settings.

Case	P_h min (MPa)	P_h max (MPa)	bypass peak (%)	relief (s)	customer
6, coordinated	3.166	3.449	24.9	0	kept
6, master-slave	2.647	3.874	3.0	147	lost
6, MPC	2.926	3.455	27.5	0	kept
7, coordinated	3.308	3.416	11.2	0	kept
7, master-slave	3.241	3.567	3.0	0	kept
7, MPC	3.291	3.416	15.0	0	kept
8, store in	3.336	3.448	24.1	0	kept
8, store out	3.124	3.448	24.1	0	kept
8, MPC, store in	3.324	3.455	27.3	0	kept

196 s after the step. Nothing is wrong with the tuning, because the plant cannot move faster than 31.3 MW/min of aggregate ramp rate. Selecting the coordinated architecture and repeating the step gives a nadir of 3.3391 MPa and a recovery in 110 s, an improvement of 0.062 MPa on the deviation and 87 s on the recovery, entirely because a 2 s valve is now in the loop ahead of the reactors.

Scenario 4 also shows why loops designed separately on a shared variable interact. Tuning the generator load controller of Table 9.4 hard, at $K_c = 0.6$ and $\tau_I = 5$ s, and repeating the same 8 MW step drives header pressure between 2.888 and 3.990 MPa over a 900 s run, lifts the relief for 198 s of it and loses the customer. The mechanism is (9.12), because a sagging header reduces the flow the governor valve passes at a fixed position, so a stiff megawatt controller answers a sagging header by opening the valve further, which takes more steam, which sags the header further.

9.6.3 Mid-Ranging and What the Resting Position Costs

Scenario 5 runs the coordinated architecture undisturbed. In steady state the bypass sits on its 3.0 % resting position passing exactly 3.00 kg/s, header pressure is 3.4000 MPa, and the three module margins are 12.8, 12.9 and 13.0 °C. Three percent of bypass is 7.0 MW thermal thrown to the condenser continuously, or 2.35 MW of electricity not sold, which is 99 $/h at the default price, and what it buys is 3 kg/s of instant upward response as well as unlimited downward response. Setting it to zero makes the fast path one-sided, because a shut valve can still catch a rising header in two seconds and can do nothing about a falling one. A setpoint step of +0.05 MPa, which leaves the header low, settles inside 0.01 MPa in 26 s with the bypass resting at 3 % and, at a zero resting position, in 61 s once the bypass has closed to 0.5 % at 349 s and in 64 s once it is at 0.3 % after ten minutes, while a −0.05 MPa step settles in 5.35 s either way. Every mid-ranging scheme makes that trade and the resting position prices it [3].

9.6.4 Load and Process Upsets, With and Without a Fast Actuator

Table 9.6 reports scenarios 6 to 8, each run twice on the same disturbances with the architecture or the store as the only change, and once more under the MPC of Section 9.5.5.

The first pair of rows in Table 9.6 is the clearest single result in the application. Coordinated control holds the header between 3.166 and 3.449 MPa through both halves of the ramp. The step down opens the bypass to 24.9 % within seconds, and the valve position controller then spends four minutes walking the reactors down so the bypass can close again, at a cost of 4.35 MWh to the condenser over the 900 s run, and the ramp up is answered almost entirely by the reactors because nine megawatts a minute of electricity is about 27 MW/min of core power against 31.3 MW/min

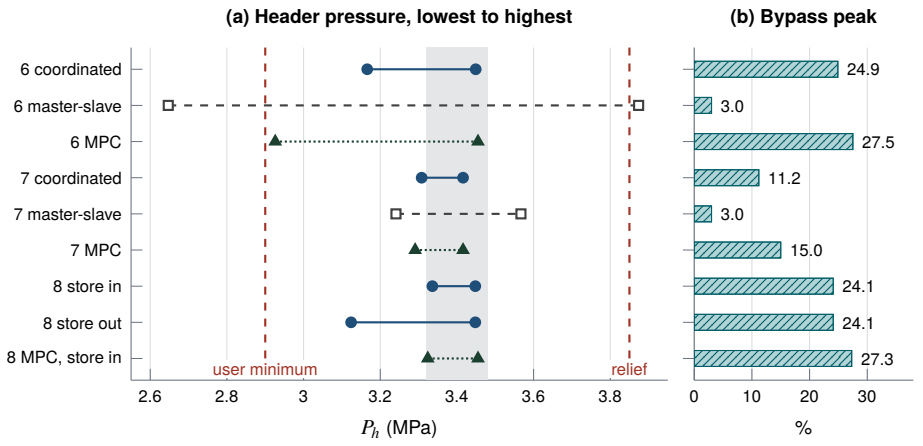

Figure 9.3 Table 9.6 drawn as ranges and bars: scenarios 6 to 8 under coordinated control (solid, circles), master-slave (dashed, squares) and the MPC (dotted, triangles). (a) Header pressure from its lowest to its highest value over each run, against the shaded ±0.08 MPa band, the industrial user's 2.90 MPa minimum and the 3.85 MPa relief pressure (dashed). (b) Bypass peak. Notice that in scenarios 6 and 7 the master-slave run, whose bypass stays at 3.0 %, has the widest range, and that scenario 6 under master-slave crosses both dashed lines: the relief lifts for 147 s and the customer is lost.

available. Master-slave on the identical events fails in both directions. The step down has nothing fast to absorb it, so the header reaches 3.874 MPa and the relief lifts for 147 s, and the ramp up then drives it to 2.647 MPa and the customer is lost. Same reactors, same rate limits, same disturbance, and the only difference is whether a fast actuator is in the loop. Figure 9.3 draws the table as ranges and bars.

Scenario 7 is milder and shows the same ordering, a nadir of 3.308 MPa under coordinated control against 3.241 MPa under master-slave and an overshoot on the return of 3.416 MPa against 3.567 MPa. The process flow controller is part of the mechanism rather than a bystander, because it has to open its valve further than it planned to hold the flow it was asked for while the header is falling, which pulls the header down again.

The store in scenario 8 improves the nadir from 3.124 to 3.336 MPa and reduces the energy thrown to the condenser from 3.97 to 3.37 MWh, from a tank that ends the run at 44.4 % charge against the 45.0 % it started with. Two limits appear at once. The delivery pressure of (9.13) falls with the charge, from 5.40 MPa full to 3.05 MPa empty, so a nearly flat tank cannot push steam into a healthy header, and storage steam is barely superheated, so a hard discharge cools the header through (9.11). Both are real properties of a sensible-heat store, and both are why it is a fourth participant rather than an answer [16].

9.6.5 The Allocation Rule Matters as Much as the Architecture

Scenario 9 fouls Unit 3's steam generator from 0.86 to 0.70 at 120 s and then brings the evening peak. Header pressure dips to 3.303 MPa at 157 s on the fouling step and is back inside its band by 194 s, because the total goes on being delivered, and the only lasting indication is a module temperature warning at 245 s. With the equal split the run ends at 900 s with the three modules at 198, 198 and 166 MW carrying margins of 8.5, 4.9 and 0.0 °C, a spread of 8.43 °C, so Unit 3 is pinned on its 315 °C limit while Unit 1 sits with eight and a half degrees spare (9.2, 4.9 and 0.0 °C once Unit 1 has finished unloading at 1200 s). Switching to the margin-proportional rule redistributes the same total to 208.9, 196.6 and 155.4 MW with margins of 5.41, 5.66 and 4.59 °C, a spread of 1.07 °C. The plant sells the same product from nearly the same steam, 242.2 against 243.7 kg/s, and has simply stopped choosing which module to break. The architecture is identical in both runs, so the 7.36 °C reduction

in margin spread belongs entirely to the splitting rule, which is the result unit-specific allocation across a multi-unit plant is built to obtain [8].

The MPC of Section 9.5.5 still believes Unit 3 is at 0.86, so at its worst instant its model allows Unit 3 203.4 MW where the true ceiling of (9.22) is 165.5 MW, and only the ceiling clamp every architecture shares holds the module. Unit 3 ends 0.10 °C past its limit, against 0.09 °C under Coordinated with the equal split, and the clamp has a cost. The plan keeps asking for a ramp the clamp refuses, the header IAE is 20.04 against 18.04 and 15.16 MPa s for the equal and margin splits, and the header settles 0.004 MPa under setpoint, 3.396 MPa at 2400 s. In a diagnostic copy without the clamp the MPC runs Unit 3 5.12 °C past its limit at 900 s. The comparison is therefore not between a controller that knows the fouling and one that does not, because the clamp and Coordinated's allocator both estimate the ceiling from the measured coolant temperatures, which follow the fouling.

9.6.6 A Module Trip Is a Dispatch Decision

Scenario 10 scrams Unit 1 at 90 s. Core power collapses toward decay heat with the 1.6 s scram time constant, 39 MW three seconds after the trip, and the isolation valve is shut in six through (9.9), so 82.0 kg/s of the header's 228.2 kg/s stops arriving into a pipe holding 7.55 s of steam. The two survivors have ceilings of 218 and 204 MW against a plant that was using 530 MW, so the arithmetic does not close and no architecture makes it close. Three runback choices were measured over 1500 s with the coordinated architecture and the margin split throughout. With no runback the nadir is 1.936 MPa, the header spends 474 s below P_p^{min} and the customer is lost. With the turbine armed the nadir is 2.943 MPa, no time at all is spent below P_p^{min}, the customer is kept and 41.7 MW is shed. With the process user armed the nadir is 2.082 MPa, 473 s are spent below P_p^{min}, the customer is lost anyway and 18.7 kg/s is shed. All three runs end with header pressure back on 3.400 MPa.

The no-runback case deserves attention for how it ends rather than for how deep it goes. The customer's 56 kg/s of demand leaves with the customer, the two surviving reactors then have enough capability for what remains, and header pressure recovers neatly to 3.400 MPa. The plant balanced itself by losing the customer, the header tile is green again by 900 s, and an operator reading the pressure trend at that instant sees a plant that handled a reactor trip. The rest of the scorecard does not agree, with the industrial user, worst margin, reserve and time off band all red.

> **TRY IT IN THE APP** **10 · Unit 1 Scrams**
>
> Load scenario 10 and press Run with the runback on None. Unit 1 trips at 90 s, the header falls to 1.936 MPa and the industrial customer is lost, after which the pressure recovers to 3.400 MPa. Reset, arm the runback on the turbine and run again: the lowest pressure is 2.943 MPa, the customer stays, and 41.7 MW of generation is shed.

Arming the runback on the turbine changes the outcome without changing anything about the reactors. The fast stage of (9.32) takes 79 MW off the generator setpoint within the fifteen seconds after the trip that would otherwise not have existed, the nadir is 2.943 MPa, and the starvation timer never starts. The fast stage then hands load back as the header recovers until the latched trim holds 41.7 MW, and the plant finishes on 3.400 MPa with the two survivors carrying 407 MW. The trade is 41.7 MW of electricity not sold against a customer whose interruption costs 40 000 $ plus 9000 $/h. Arming it on the process user cannot work, for reasons of arithmetic rather than tuning, because the runback may take the customer down only to a little above its 34 kg/s shutdown flow, which is 18.7 kg/s of relief against a shortfall of about 47 kg/s. Figure 9.4 shows the no-runback case 540 s after the trip.

Under the MPC, which after the scram plans the two survivors with Unit 1's ramp held at zero, the armed turbine runback picks up at 92.30 s against 92.05 s. The plan shuts the bypass from 3.02 % in its first decision, at 90.25 s, where Coordinated has it shut at 91.75 s, so the header reaches the pick-up threshold a scan later. The nadir is 2.941 against 2.943 MPa, 42.2 against 41.7 MW is shed,

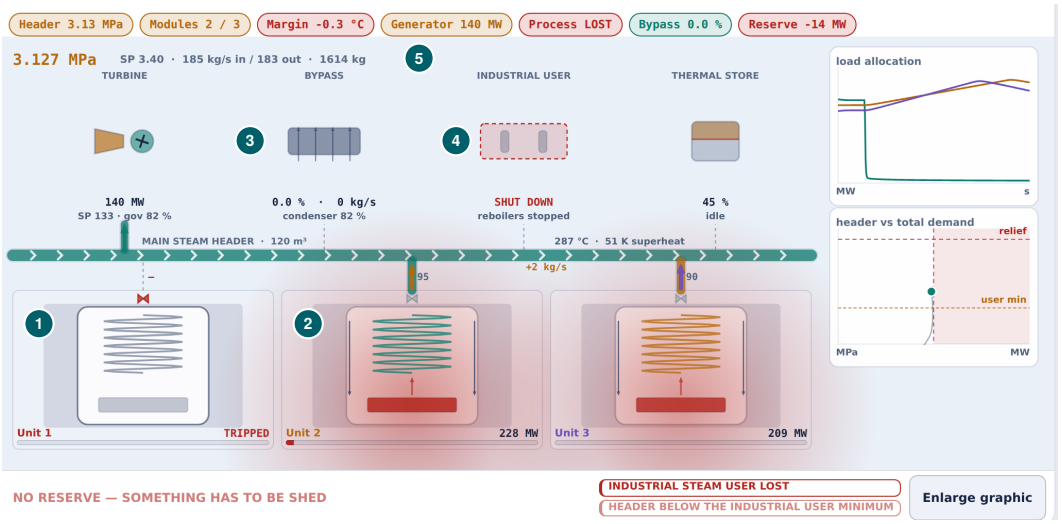

Figure 9.4 Scenario 10 at $t = 10.50$ min, 540 s after the Unit 1 scram, with the runback set to None and the coordinated architecture in service, in the plant view of the app. Header pressure reads 3.127 MPa (5) and is climbing back toward its 3.40 MPa setpoint for the wrong reason, because the industrial user (4) was lost and its 56 kg/s of demand went with it. Unit 1 (1) is tripped and pale, Units 2 and 3 (2) carry 228 and 209 MW, above ceilings of 222 and 207 MW, the worst module margin is –0.3 °C, thermal reserve is –14 MW, and the bypass (3) is hard shut.

and the IAE is 78.02 against 77.42 MPa s. The high-temperature trip, its limit lowered to 302.5 °C below Unit 1's coolant temperature at 10 s of scenario 5, fires at 10.05 s under both.

9.6.7 The Graded Challenge

Scenario 12 loads the plant weakly tuned, on master-slave with an equal split, a bypass hard shut at zero percent and Unit 3 already fouled to 0.78, then scrams Unit 1 at 120 s and asks the customer for 62 kg/s at 450 s, with a target of ±0.08 MPa of 3.40 MPa, no lost customer and no relief lift. Run exactly as loaded, the header reaches 2.091 MPa, the customer is lost 56 s after the scram, and the worst residual in the final third of the run, from $t = 594$ s, is 0.725 MPa, about nine times the band. One change is both necessary and sufficient for this target. A runback armed on the turbine, with everything else as loaded, keeps the customer, lifts no relief and leaves a worst residual of 0.042 MPa, while every run without it loses the customer, including the full suggested set with the runback switched off. The other choices add margin rather than feasibility. Coordinated allocation and the margin-proportional split with the runback, on the weak tuning and a shut bypass, give 0.033 MPa, and the suggested settings add a 6 % bypass resting position, storage on automatic and stronger tuning. They give a nadir of 3.087 MPa, no relief lift, 31.4 MW of generation shed, and a worst residual in the final third of 0.046 MPa against the 0.08 MPa tolerance, inside the band. The split rule cannot matter once the trip has happened, because both survivors then sit on their temperature limits whatever the rule, with a worst margin of 0.0 °C at 900 s in all three runs, so the target rewards the dispatch decision rather than the structure. The MPC confirms it, and the challenge itself is unchanged. With only the architecture switched to MPC the runback stays off, the customer is lost 54.4 s after the scram and the worst residual is 0.736 MPa, while the MPC with the suggested settings meets the target at 0.038 MPa, with Unit 3 ending 0.02 °C past its limit on the shared clamp where the suggested coordinated run ends it on the limit.

Table 9.7 Scenario 11 and the Other Cases, Coordinated Against the MPC at Scenario 11's Settings With the Same Seed, Events and Noise. Header IAE in MPa s From the 4 Hz Historian, to 900 s Except Scenarios 4 and 7, to 600 and 780 s, Unless Stated

Case and measure	Coordinated	MPC
Scenario 11, 0–900 s		
IAE, whole run	72.56	71.61
IAE, load drop, 0–300 s	3.36	1.04
IAE, process step, 300–540 s	1.22	0.35
IAE, trip, 540–900 s	67.98	70.22
At the rate limits (module-s)	1262	383
Module (MW), bypass (%) travel	237, 60.2	126, 58.3
To the condenser (MWh)	3.15	4.72
Worst coolant margin (°C)	1.31	4.08
Plans, rejected, mean solve (ms)	—	900, 0, 3.80
IAE, and to the condenser (MWh)		
4, +8 MW load step	5.58, 0.84	7.95, 1.02
6, evening ramp	47.15, 4.35	111.73, 5.82
7, process step	10.82, 2.09	12.76, 2.97
8, store in service	19.98, 3.37	19.03, 5.65
8, store in service, to 2400 s	22.63	33.18
9, equal split	18.04, 2.32	20.04, 1.84
9, margin split	15.16, 2.21	20.04, 1.84
10, no runback	504.80, 0.87	513.19, 0.79
10, turbine runback	77.42, 1.20	78.02, 1.26
Scenario 12, suggested settings		
Worst residual (MPa), Unit 3 margin	0.046, 0.00 °C	0.038, –0.02 °C

9.6.8 Model Predictive Control Against Coordinated Allocation

Scenario 11 runs one day on the tuned loops with the margin split, the store out of service and the turbine runback armed. The grid takes 23 MW off the generator at 60 s, the process user steps from 56 to 68 kg/s at 300 s, and Unit 1 scrams at 540 s. Table 9.7 sets Coordinated against the MPC on it and on the earlier cases, with Coordinated on each case's own tuning, the MPC at scenario 11's settings throughout and no preview given to either structure, and Figure 9.5 shows the day.

Over the whole run the header IAE is a tie, 71.61 against 72.56 MPa s, 1.3 % apart, and the difference lies in how the steam is found. On the load drop Coordinated opens the bypass to 24.9 % at 125.5 s, then walks the reactors down behind it, 93 MW of module travel by 300 s with the bypass back at 8.0 %. The MPC opens it to 27.6 % at 148.3 s and still holds 20.1 % at 300 s, having moved the modules 40 MW. Over the day the modules spend 383 against 1262 module-seconds ramping at their rate limits and travel 126 against 237 MW, and the price is 4.72 against 3.15 MWh to the condenser, because a ramp is priced in (9.33) and an open bypass only through q_b. Both structures bottom out at 2.97 MPa after the trip. The worst coolant margin of 4.08 against 1.31 °C comes from loading the survivors less, 389.6 against 405.6 MW of core power at 900 s, and not from the plan's temperature limit, which carries no slack on any of the 900 plans. It holds for scenario 11 only, where every steam generator is in its declared condition. The bypass rest weight was lowered once before tuning, from 0.444 to 0.02, and it decides this ranking, since at 0.444 the IAE is 84.24 MPa s with 2.82 MWh to the condenser.

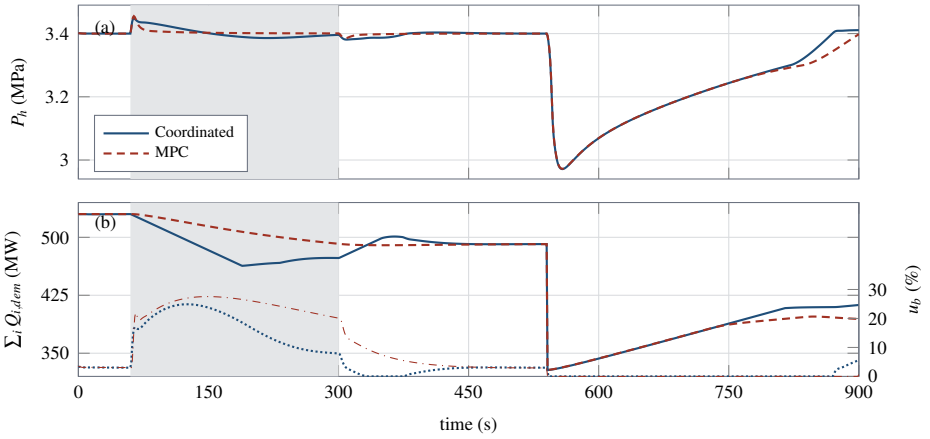

Figure 9.5 Scenario 11 under Coordinated (blue, solid) and the MPC (red, dashed), the decisive load drop shaded. (a) Header pressure. (b) Total module demand (left axis, Coordinated solid and MPC dashed) and bypass command (right axis, Coordinated dotted and MPC dash-dotted), with the bypass held open longer and the reactors moved less by the MPC.

> **+ DID YOU KNOW? An Optimizer That Only Ties**
>
> From the simulator: over the day of scenario 11 the model predictive controller and coordinated allocation hold the header equally well, 71.61 against 72.56 MPa s of integrated error. The plan keeps the modules at their ramp limits for 383 module-seconds instead of 1262, and pays with 4.72 instead of 3.15 MWh sent to the condenser (Table 9.7).

The other cases show where pricing the ramp costs pressure. The MPC spends less time at the rate limits in every case but scenario 10 without the runback, 1059 against 1047 module-seconds, by factors of 1.4 to 9.0, and moves the modules less in all of them. It wins the downward halves, 1.11 against 3.59 MPa s over the first 480 s of scenario 6 and 0.42 against 1.03 on scenario 7's return. It loses the upward events, scenarios 7 and 4, where λ delays the first ramp, and loses scenario 6's evening ramp by more than a factor of two, the header falling to 2.926 MPa against 3.166 MPa. The plan holds w_u at its estimate for two minutes while the ramp keeps raising the draw, so after 480 s its header prediction 30 s ahead is 0.052 MPa high on average. It sends more steam to the condenser in most cases but less in scenario 9 and in scenario 10 without the runback. Scenario 8 depends on the run's length, 19.03 against 19.98 MPa s to 900 s and 33.18 against 22.63 to 2400 s. The store's discharge controller, offered the setpoint 0.04 MPa low, holds the header near 3.368 MPa from 900 to 1300 s while the plan walks the total demand up from 532.5 to 567.9 MW at its priced rate. No plan was rejected in any case.

9.7 Technology in Practice

The symbols of Figure 9.1 are, in a plant, pressure transmitters on the main steam header, power and flow transmitters on the generator and the process line, the reactors' own power control, a set of air-operated valves, and the controllers of a plant control system. Figure 9.6 draws the coordinated architecture with that hardware in place of the symbols.

What Measures It

Header pressure comes from pressure transmitters, which convert the pressure to a standard signal calibrated to a span, most often an analog current loop and, in newer smart transmitters, a digital signal that can also report the transmitter's own health [23]. The chapter carries each of the three as

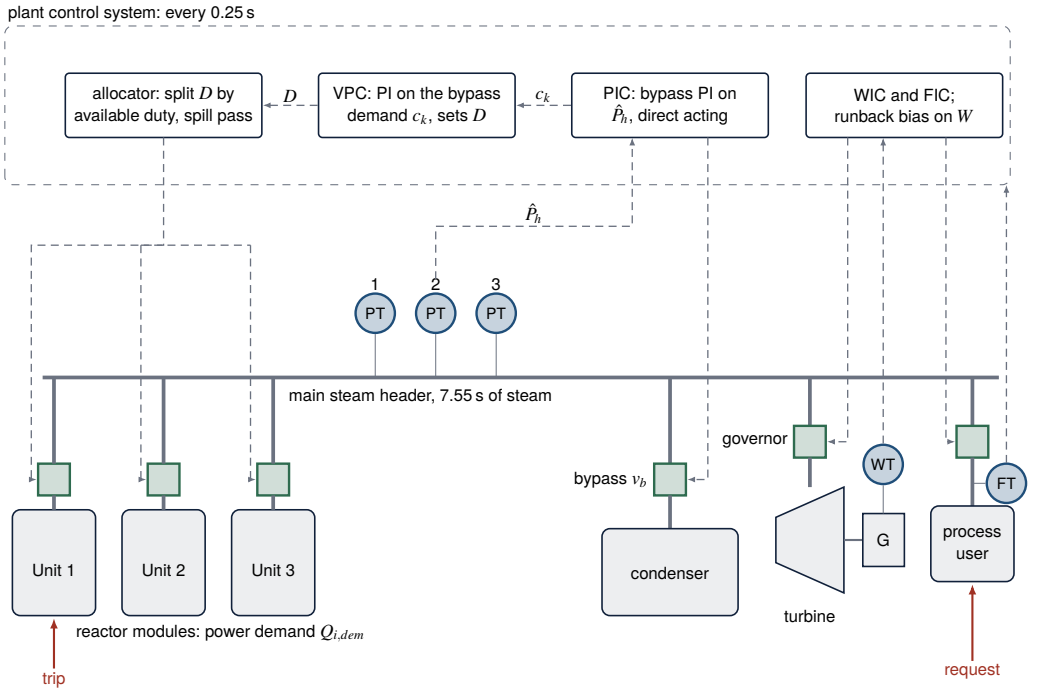

Figure 9.6 The coordinated architecture of Section 9.5 with plant hardware in place of the symbols of Figure 9.1. Sensors are circles, final elements are green squares, and dashed arrows are signals. The bypass PI block reads the header transmitters' unbiased value $\hat{P}_h$, equal to the reading of PT-2, and moves the bypass in seconds. The valve position block reads that block's unclamped demand and sets the total D, which the allocator divides into the three module power demands. The generator output (WT) and process flow (FT) loops move the governor and process valves. The storage valves and the relief are not drawn. Notice that one reading of the header feeds one pressure controller here, where the independent architecture gives each module a controller on its own transmitter.

a 0.4 s filter, Equation (9.19), and gives them the calibration spread of Equation (9.20). Generator output and process steam flow have their own transmitters, carried as 0.8 s and 1.2 s lags. None of the three measures what scenario 3 needs to see, the margin of each module, which comes from the coolant temperatures.

> **TECHNOLOGY** **Sensor: The Header Pressure Transmitter**
>
> A common electronic pressure sensor is the differential capacitance cell. A taut metal diaphragm sits between two fixed metal plates, and pressure flexes it toward one plate and away from the other, changing the two capacitances [24]. The transmitter turns the reading into a standard signal, most often 4 to 20 mA over a calibrated span, and the live zero at 4 mA reveals an open circuit [23]. The error that matters here is a zero shift, which moves every point of the range by the same amount [24]. On a trend it shows as nothing at all, only a steady and believable number. The chapter's three transmitters carry such shifts of $+0.015$, 0 and -0.015 MPa behind a 0.4 s lag.

What Moves It

The slow final element is each reactor's power control, which the chapter reduces to a rate-limited demand and a 14 s lag. In an island-mode dispatch optimization of small modular reactor units, most of the variation in demand was met by moving core thermal output with the control rods, and the turbine bypass was used below minimum power and for short periods when xenon limited the rods [9].

The fast final elements are valves: the governor, the process valve, the bypass and the two storage valves, which the chapter strokes as first-order lags of 0.6 to 5.0 s with no dead band and no stiction.

> **TECHNOLOGY** **Final Element: The Turbine Bypass Valve**
>
> The bypass to the condenser is an air-operated valve. Air pressure on a diaphragm pushes the stem against a spring, and a positioner adjusts the air until the stem reaches the position the controller asks for [25]. Valves that dump main steam to the condenser must stroke very quickly on a turbine load rejection, yet hold stable pressure control during startup and shutdown. Wear shows as air leakage, friction and dead band, which a digital positioner can find before it causes an upset. Safety analyses also consider the opposite failure, one turbine bypass valve opening by mistake, which raises steam flow and cools the reactor coolant [25]. In the chapter the bypass strokes as a 2.0 s lag and passes 1.00 kg/s per percent.

Where the Controller Runs

Plant control systems run supervisory and lower-level loops with cycle times from milliseconds to minutes, often nested so that one loop sets the setpoint of another [26]. The chapter's loops execute every 0.25 s and the predictive controller every 1 s. The coordinated structure is a chain of function blocks: a direct-acting PI block on header pressure that drives the bypass, a valve position PI block that reads that block's unclamped output and sets the total demand, and an allocation block that divides the demand by available duty with its spill pass. The runback adds a bias to the generator load setpoint.

> **IN PRACTICE** **Checking Three Transmitters Against Each Other**
>
> Where several instruments read the same quantity, operators perform a channel check: they compare the readings of the redundant channels on a regular basis and confirm that they agree within pre-defined criteria. On-line monitoring does the same from recorded data and can show that a sensor has not drifted since its last calibration. That matters because mis-calibration during maintenance is one of the most common modes of failure [23]. The three header transmitters of Equation (9.20) would pass such a check, since a spread of 0.015 MPa is a good calibration. Agreement within tolerance is not a license for three integrators on three readings.

What Is Modernizing It

Three developments reach this plant. Smart transmitters allow a calibration check from a remote location [23], and on-line monitoring of instrument channel calibration status uses data the control system already records [23], which is how a growing transmitter spread could be caught before a second integrator finds it. Digital positioners diagnose air leakage, friction and dead band on the valves, so that degradation is found before an upset [25]. And dispatch across units is becoming an optimization over days and years, with outage schedules and markets for grid services in the same problem [9], the long-horizon relative of the predictive controller of Section 9.5.5.

9.8 Limitations

The application is not a reactor model and cannot support a claim about a real plant's transient capability, safety margins or licensability. With no neutronics, nothing about reactivity, xenon or a rod-worth limit can be studied, and with no primary flow state the flow-limited power that constrains a natural-circulation module is absent. The steam generator is one lumped conductance where a real helical-coil once-through unit has moving section boundaries [12], and feedwater flow is not independent, so the coupled feedwater problem across several once-through generators is out of reach [11]. The header is one lumped volume with no pressure drop, which understates the disagreement between transmitters at opposite ends of a real header and therefore understates the failure of Figure 9.2. Repeatability is itself unrealistic, because a fixed 0.05 s step and a fixed 0.25 s scan reproduce one trajectory on every machine while a distributed control system does no such thing.

The four features that survived the simplification, a small shared inventory between rate-limited producers, a ceiling that moves with the shared variable, a fast expensive actuator mid-ranged against a slow free one, and a consumer that can be lost, all survived because a control system has to be designed around each of them. Where each assumption breaks is set out in Section 9.3.6.

The MPC of Section 9.5.5 holds the users' draw constant over its two-minute horizon and leaves out their loops, which is why its model overstates the header's open-loop response by about a third and why it loses the evening ramp by 111.73 against 47.15 MPa s, and it plans every module at its declared fouling, so a fouled module is planned against a ceiling it no longer has. The shared clamp and Coordinated's allocator, by contrast, estimate each ceiling from the measured coolant temperatures and the declared conductance, so for a fouled module their estimate can differ from the physical ceiling during a transient and approaches it as the measured temperature approaches its limit.

9.9 Discussion

The result that generalizes furthest is not the divergence of Figure 9.2 but the reason an operator cannot see it. Header pressure is the variable that appears on the control room display, carries the alarm limits and enters the scorecard, and through the whole of scenario 3 it holds 0.0162 MPa of a 0.16 MPa band, while the evidence that the plant is coming apart is on per-module power and margin trends and a margin-spread score, none of which is a controlled variable. The same asymmetry appears in the trip results, where the run that loses the customer ends with header pressure on setpoint just as the run that keeps it does. Any plantwide design that selects its displays and alarms from the list of controlled variables reproduces both blind spots, which is an argument for choosing the monitored set from the constraint list rather than from the loop list [1, 2].

The economics of the fast path are worth stating with numbers because they are frequently argued without them. A bypass resting at 3 % costs 7.0 MW thermal continuously, which is 99 $/h of electricity not sold at the default price, while losing the industrial customer costs 40 000 $ at once plus 9000 $/h until it returns. The one-off loss alone equals about 17 days of continuous bypass loss, and each hour the customer stays off adds almost four days more. On this plant, however, neither a 3 % resting position in scenario 10 nor a 6 % one in the challenge keeps the customer through the trip without the runback, so the money buys transient margin rather than the customer. The application exposes both the resting position and the four prices so that the trade can be computed rather than asserted.

The model also supports methods beyond the implemented loops. The ceiling of (9.22), the ramp limits of (9.3) and the customer's two limits are inequality constraints that a receding-horizon formulation enforces explicitly rather than through a spill pass and a runback latch, and the storage state of (9.17) is what makes the horizon matter. The controller of Section 9.5.5 enforces the ramp limits as hard bounds and the temperature limit as a priced one, and it leaves the store, the customer's two limits and the runback outside its plan. Coupling one heat source to several products is the arrangement that nuclear-renewable integrated energy system studies address with dynamic optimization over a day rather than with regulatory loops [15, 16], and allocation is then answered by the optimizer's constraint set rather than by a participation shape [7, 8].

⚠ OPEN CHALLENGE Dispatching Many Reactors Over Many Years

A multi-unit site can share load between its units, but it also shares services such as refueling equipment, and its units may follow a fixed sequence of fuel reloads. In one study, joining the markets for grid services lengthened the fuel cycles, and over a decade the refueling outages drifted by three months, into the summer. The cost of component wear was left for future work [9]. Progress would be an allocator that divides load across units under their fuel, outage and wear constraints as well as their thermal margin, which is all this chapter's allocator weighs.

9.10 Conclusion

This chapter studied three small modular reactors on one steam header as an interactive laboratory for load allocation on a shared variable. The simulator integrates eighteen states, three core powers on rate-limited demands, three average primary coolant temperatures, three once-through steam generator pressures, the header mass and temperature, a storage charge, five valve positions and a generator output lag. The header holds 1724 kg of steam and passes 228.2 kg/s for 7.55 s of residence against a primary loop time constant of 52.5 s and an aggregate ramp capability of 31.3 MW/min. Independent per-module pressure control separates the three modules by 91.7 MW in twenty minutes with no disturbance applied while header pressure stays inside 0.0162 MPa, master-slave holds a 23.2 MW spread and a 196 s recovery from an 8 MW load step, and coordinated allocation recovers from the same step in 110 s because a 2 s bypass valve is in the loop ahead of the reactors. An equal megawatt split leaves margins of 8.5, 4.9 and 0.0 °C where the margin-proportional rule leaves 5.41, 5.66 and 4.59 °C on the same total demand. A Unit 1 scram with no runback takes the header to 1.936 MPa and loses the industrial customer, after which the plant recovers to 3.400 MPa precisely because the customer is gone, while a turbine runback holds the nadir at 2.943 MPa and sheds 41.7 MW, and in the graded challenge that runback is the one change both necessary and sufficient. A model predictive controller planning the three ramps and the bypass ties coordinated allocation on header error, 71.61 against 72.56 MPa s over a day with a load drop, a process step and a trip, while spending 383 rather than 1262 module-seconds at the rate limits and 4.72 rather than 3.15 MWh on the condenser. A fuller model would add a sectioned steam generator with independent feedwater control, a primary flow state to bound module power, a tank thermocline, and a receding-horizon allocator that infers each module’s fouling from its coolant temperature rather than assuming the declared value. The plant shows, in a form a learner can watch, that a controlled variable on setpoint is not the same thing as a plant under control.

KEY POINTS

- A pressure on its setpoint is not proof of a plant under control: with no disturbance, three controllers on one pressure walk the reactors 91.7 MW apart while the header stays inside its band.
- One controlled variable gets one integrator, and the split is a separate decision made by a separate mechanism.
- A 2 s bypass valve ahead of the reactors recovers an 8 MW load step in 110 s, where master-slave needs 196 s.
- Splitting by margin leaves 5.41, 5.66 and 4.59 °C of margin where an equal split leaves 8.5, 4.9 and 0.0 °C.

Where else this applies. Plantwide structure decides the answer in Volume 6, Chapter 7 and Volume 6, Chapter 8, and Volume 4, Chapter 10 mid-ranges a fast, costly bypass against slower actuators in the same way.

References

[1] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, “Plantwide control design procedure,” *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

[2] S. Skogestad, “Control structure design for complete chemical plants,” *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[3] F. G. Shinskey, *Process Control Systems*, 4th ed. New York, NY: McGraw-Hill, 1996.

[4] F. Paparella, L. Dominguez, A. Cortinovis, M. Mercangoz, D. Pareschi, and S. Bittanti, “Load sharing optimization of parallel compressors,” in *2013 European Control Conference (ECC)*, 2013, pp. 4059–4064.

[5] A. Torzhkov, P. Sharma, C. Li, R. Toso, and A. Chakraborty, “Chiller plant optimization – an integrated optimization approach for chiller sequencing and control,” in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 2741–2746.

[6] J. M. Guerrero, M. Chandorkar, T.-L. Lee, and P. C. Loh, "Advanced control architectures for intelligent microgrids—part I: Decentralized and hierarchical control," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 4, pp. 1254–1262, 2013.

[7] A. Nuerlan, P. Wang, J. Wan, and F. Zhao, "Decoupling header steam pressure control strategy in Multi-Reactor and Multi-Load Nuclear Power Plant," *Progress in Nuclear Energy*, vol. 118, p. 103073, 2020.

[8] H.-H. Byun and M.-S. Yim, "Genetic algorithm-based non-synchronous unit-specific optimal load-following control of multi-unit small modular reactor," *Energy*, vol. 314, p. 134091, 2025.

[9] S. Alhadhrami, G. J. Soto, and B. Lindley, "Dispatch analysis of flexible power operation with multi-unit small modular reactors," *Energy*, vol. 280, p. 128107, 2023.

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

[11] P. Sun and X. Zhang, "Modeling and simulation of feedwater control system with multiple once-through steam generators in a sodium-cooled fast reactor (I): Dynamic model development," *Annals of Nuclear Energy*, vol. 138, p. 107208, 2020.

[12] F. U. Zaman, K. Qureshi, I. Haq, and W. Siddique, "Thermal hydraulics analysis of a helical coil steam generator of a small modular reactor," *Annals of Nuclear Energy*, vol. 109, pp. 705–711, 2017.

[13] J. Carter, "Small modular reactor deployment: Learning from the past and the present," *The Electricity Journal*, vol. 29, no. 8, pp. 1–8, 2016.

[14] D. T. Ingersoll, Z. J. Houghton, R. Bromm, and C. Desportes, "NuScale small modular reactor for Co-generation of electricity and water," *Desalination*, vol. 340, pp. 84–93, 2014.

[15] R. A. Jacob and J. Zhang, "Modeling and control of nuclear–renewable integrated energy systems: Dynamic system model for green electricity and hydrogen production," *Journal of Renewable and Sustainable Energy*, vol. 15, no. 4, p. 046302, 2023.

[16] A. Rigby, M. J. Wagner, and B. Lindley, "Dynamic modelling of flexible dispatch in a novel nuclear-solar integrated energy system with thermal energy storage," *Annals of Nuclear Energy*, vol. 204, p. 110534, 2024.

[17] 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.

[18] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

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

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

[21] 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.

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

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

[24] 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.

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

[26] 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

Three reactors feed one main steam header, which supplies a turbine, an industrial steam user, a condenser bypass and a thermal store. Which quantity is the single controlled variable that every module and every user shares?

(a) The fouling multiplier $m_{f,i}$ of each module's steam generator
(b) The calibration spread σ_b between the three pressure transmitters
(c) The main steam header pressure P_h, in MPa
(d) The mass of steam m_h held in the header, in kg

9.2 Concept check · 1 min

In scenario 3 each module runs its own well-tuned PI controller on header pressure, the transmitter spread is at its default of 0.015 MPa, and no disturbance is applied. The module powers walk 91.7 MW apart while the header stays inside its band. What keeps them walking?

(a) A slow load change on the grid that the three loops share unequally
(b) The coupling through saturation temperature, which moves every module's heat removal
(c) A gain three times too high on each loop, which makes the header oscillate
(d) Three integrators on readings that differ by a bias, so no pressure zeroes all three errors

9.3 Concept check · 2 min

The coordinated architecture has three layers: a bypass pressure loop, a valve position controller and an allocator. Which statement describes what the three layers achieve together?

(a) They hold the header on setpoint while dividing the reactor load within each module's thermal limit
(b) They send as much steam as the condenser accepts through the bypass to hold the header
(c) They hold each steam generator pressure at a separate setpoint of its own
(d) They keep the thermal store full whatever the header pressure does

9.4 Theory and calculation · 4 min

At the design point the header holds 1724 kg of steam and passes 228.2 kg/s. A module's steam side holds about 470 kg against an average 76.1 kg/s. Use Tables 9.2 and 9.1.

(a) Compute the residence time of the header and of one module's steam side. *Checkpoint:* 7.55 s and 6.18 s.

(b) Compute the primary loop time constant C_{p1}/UA_1 of Equation (9.4) for Unit 1. *Checkpoint:* 52.5 s.

(c) Compute the aggregate ramp capability of the three modules in MW/min, and the shortest time in which they can add 24 MW of core power together. *Checkpoint:* 31.25 MW/min and 46 s.

9.5 Theory and calculation · 8 min

At the design demand $D = 530$ MW the three ceilings are 232, 218 and 204 MW, and the conductances are $UA_i = 3.43\, m_{f,i}$ with fouling multipliers 1.00, 0.93 and 0.86. Treat the ceilings as fixed and use the margin of Equation (9.24).

(a) Divide D in the ratio of the ceilings and compute the three margins. *Checkpoint:* 12.8, 13.0 and 13.1 °C.

(b) Divide D equally instead and compute the three margins and their spread. *Checkpoint:* 16.1, 13.0 and 9.3 °C, a spread of 6.9 °C.

(c) Under the equal split, at what total demand does Unit 3 reach its limit, and at what total demand does the proportional split use up every margin at once? *Checkpoint:* 612 MW and 654 MW.

9.6 Theory and calculation · 6 min

The bypass rests at 3 % and passes 3.00 kg/s of header steam at 3.40 MPa and 296.8 °C to the condenser. Feedwater is at 149 °C, the turbine makes 783 kJ/kg, electricity sells for 42 $/MWh, and losing the industrial user costs 40 000 $ plus 9000 $/h.

(a) From Equation (9.1) with $c_p = 2.65\,\mathrm{kJ/kg/K}$, compute T_{sat}, the header enthalpy and $h_f(T_{fw})$, and the thermal power the bypass throws away. *Checkpoint:* 241.0 °C, 2951.9 kJ/kg, 631.9 kJ/kg, 6.96 MW.

(b) Compute the electricity that steam would have made and its value per hour. *Checkpoint:* 2.35 MW, 98.7 $/h.

(c) How many days of bypass loss equal the one-off penalty, and how many days does each hour of interruption add? *Checkpoint:* 16.9 days and 3.8 days.

9.7 App challenge · 8 min

Load scenario 3, Independent pressure control, fighting itself, and change no setting. Run to 1200 s at a speed of 60× at the default seed and export the CSV file. The file's time column is in minutes.

(a) From columns `Q1_MW`, `Q2_MW` and `Q3_MW`, read the three module powers and their spread at 10 and 20 minutes. *Checkpoint:* 131.5, 202.0 and 198.1 MW, spread 70.5 MW, then 120.1, 211.8 and 198.4 MW, spread 91.7 MW, each within 0.2 MW.

(b) From column `Ph_MPa`, find the lowest and highest header pressure over the run, and say what an operator watching only that trend would conclude. *Checkpoint:* 3.385 and 3.401 MPa, each within 0.001 MPa.

(c) Press Reset, switch to Advanced, set Pressure transmitter calibration spread to 0 on the Modules tab, and run to 1200 s again. Find the lowest and highest header pressure. *Checkpoint:* 2.693 and 3.917 MPa, each within 0.005 MPa.

9.8 App challenge · 8 min

Load scenario 4, Master–slave load sharing, and change no setting. The grid asks for 8 MW more at 90 s. Run to 600 s at a speed of 20× at the default seed and export the CSV file, whose time column is in minutes.

(a) From column `Ph_MPa`, find the lowest header pressure after the step, and the time after the step of the last row more than 0.02 MPa from the 3.40 MPa setpoint. *Checkpoint:* 3.277 MPa, within 0.002 MPa, and 196 s, within 2 s.

(b) Press Reset, set Pressure control architecture to Coordinated, and repeat the run and the two readings. *Checkpoint:* 3.339 MPa, within 0.002 MPa, and 110 s, within 2 s.

(c) Explain why the master-slave recovery cannot be made much faster by tuning, and what makes the coordinated run faster.

9.9 App challenge · 9 min

Load scenario 10, Unit 1 scrams, and change no setting: Coordinated, margin split, Runback target None. Unit 1 trips at 90 s. Use a speed of 60× and the default seed for every run, and export the CSV file of each.

(a) Run to 900 s. Find the lowest header pressure in column `Ph_MPa`, and read the Industrial user and Header tiles at the end. *Checkpoint:* 1.936 MPa, within 0.005 MPa, LOST, and 3.400 MPa.

(b) Press Reset, switch to Advanced, set Runback target to Turbine on the Dispatch tab before Run, and run to 900 s. Find the lowest header pressure and the generator load shed in the last row of column `shed_bias_MW`, and read the Industrial user tile. *Checkpoint:* 2.943 MPa, within 0.005 MPa, −41.7 MW, within 0.2 MW, and kept, 0 s low.

(c) Load scenario 12, Challenge: Ride through the trip. Change only Runback target to Turbine, run the full 900 s and read the verdict. *Checkpoint:* worst header error 0.042 MPa over the last third, within 0.002 MPa, user kept, relief 0 s.

9.10 Application area · 4 min

In scenario 10 the turbine runback keeps the industrial user by shedding 41.7 MW of generation, at the default electricity price of 42 $/MWh. Losing the user costs 40 000 $ plus 9000 $/h until it returns.

(a) What does the shed generation cost per hour, and after how many hours of shedding would it equal the one-off penalty alone? *Checkpoint:* 1751 $/h and 22.8 h.

(b) The process runback can cut the user only to a little above its 34 kg/s shutdown flow, 18.7 kg/s of relief against a shortfall of about 47 kg/s. What does that mean for the choice of product to shed?

9.11 Application area · 3 min

An operator display for this plant shows header pressure with alarms at ±0.08 MPa and nothing else. Using the Discussion and the Limitations, name the readings that would have revealed the divergence of scenario 3 and the lost user of scenario 10, and say whether a real header with a pressure drop along it makes the scenario 3 failure better or worse.

9.12 Review the agent · 5 min

An AI assistant is asked how to keep a separate pressure controller on each module and replies as follows. “Scenario 3 fails because each module loop has $K_c = 950$ against a master of 320, so three loops act at three times the master’s gain. Set Per-module pressure gain to 110 MW/MPa, a third of the master, and the three loops together give the master’s loop gain of 2.2. I ran scenario 3 for ten minutes with that gain: the spread is 7.2 MW and the header stays between 3.399 and 3.401 MPa. Independent control now works as well as master-slave.” Find the flaw and say how you would check it.

CHAPTER 10

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

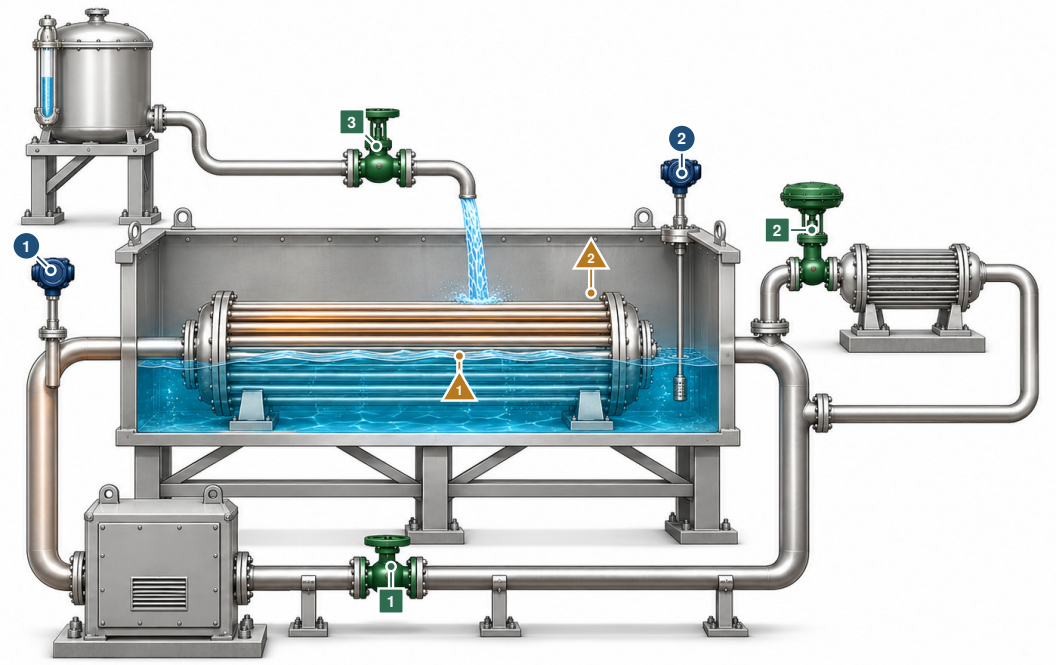

The controller moves 1 isolation valve on the cooled water returning to the reactor 2 auxiliary cooling train, which needs a diesel 3 makeup from a tank into the pool, opened by hand

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

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

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

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

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

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

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

THREE WAYS TO READ THIS CHAPTER

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

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

Learning objectives

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

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

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

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

10.1 Introduction

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

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

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

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

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

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

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

10.2 Decay Heat and the Passive Answer to It

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

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

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

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

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

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

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

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

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

10.3 Plant Model

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

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

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

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

10.3.1 Decay Heat and the Collapse of Fission

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

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

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

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

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

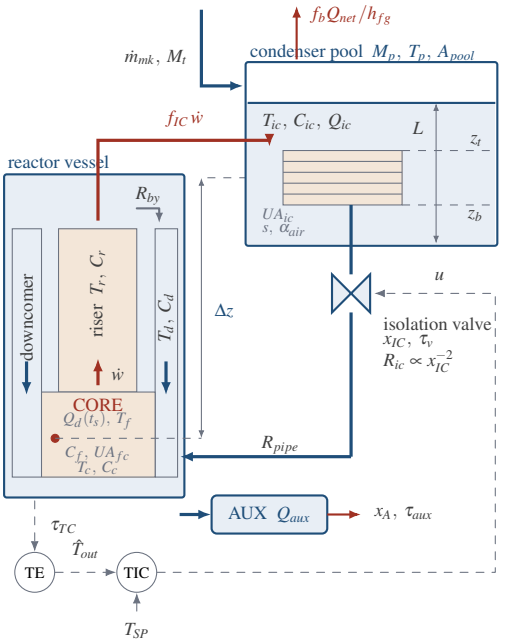

Figure 10.1 Natural-circulation loop, isolation condenser and condenser pool of the modeled plant. Decay heat Q_d raises the fuel node T_f, which passes heat to the core coolant node T_c through UA_{fc}, and the core outlet $T_{out} = 2T_c - T_d$ feeds the riser. A fraction f_{IC} of the circulation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

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

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

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

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

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

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

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

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

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

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

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

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

10.3.3 Circulation From a Momentum Balance

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

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

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

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

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

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

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

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

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

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

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

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

10.3.4 The Condenser and the Four Things That Spoil It

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

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

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

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

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

10.3.5 The Pool, in Three Regimes

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

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

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

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

10.3.6 Actuators and the Instrument Chain

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

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

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

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

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

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

10.3.7 Coping Time in Closed Form

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

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

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

10.3.8 Parameters, Initialization and Numerics

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

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

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

10.3.9 Assumptions

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

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

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

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

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

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

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

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

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

Table 10.2 Model Parameters. Hidden Constants Are Never Displayed.

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

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

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

10.4 Static Behavior of the Loop and the Valve

10.4.1 The Cube Root, Measured

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

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

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

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

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

10.4.3 The Arithmetic That Decides Everything

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

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

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

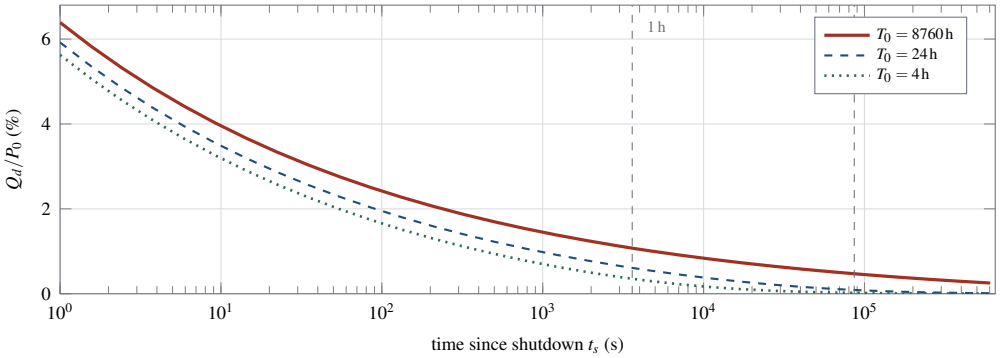

Figure 10.2 Decay heat fraction against time since shutdown for three operating histories: a year at power (solid), a day (dashed) and four hours (dotted). The year-long core falls only from 6.391 to 0.2525 % over six decades of time. Four hours at power supplies 32.9 % of the year-long fraction at one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

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

10.4.4 Operating History

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

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

10.5 Four Architectures and Two Controllers

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

10.5.1 The Implemented Algorithm

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

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

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

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

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

10.5.2 The Override Is a Selector in Everything but Name

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

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

Table 10.4 Scenario Sequence and the Mechanism Each Exercises

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

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

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

10.6 Measured Results

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

10.6.1 The First Two Minutes, and the Next Eight Hours

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

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

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

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

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

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

10.6.2 Fully Passive, and the Bill It Runs Up

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

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

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

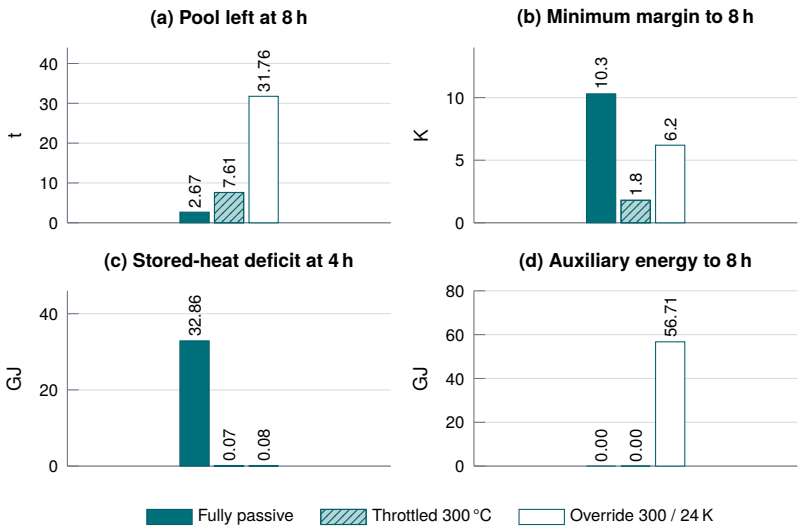

Figure 10.3 The three architectures of Table 10.5 on the identical trip transient: (a) pool water left at eight hours, (b) the minimum temperature margin over eight hours, (c) the stored-heat deficit at four hours and (d) the auxiliary energy taken over eight hours. Solid: fully passive. Hatched: throttled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

10.6.3 Throttling a Passive System on Purpose

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

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

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

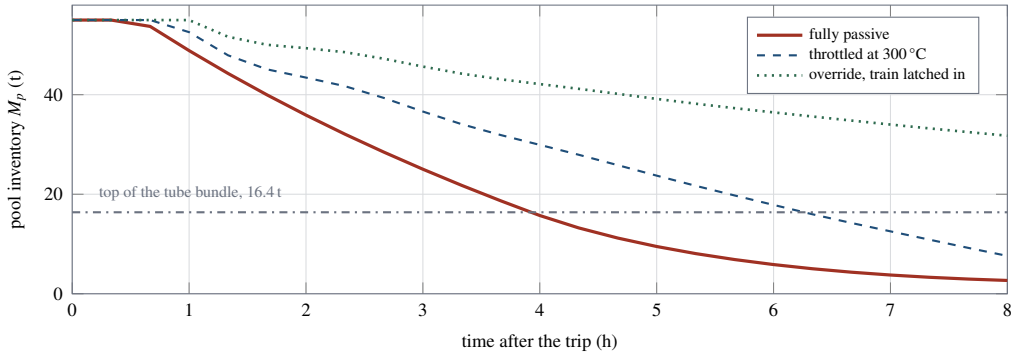

Figure 10.4 Condenser pool inventory for the three architectures of Table 10.5 on the identical trip transient: fully passive (solid), throttled (dashed) and the override with the train latched in (dotted). The dash-dotted line marks the top of the tube bundle. The passive trace crosses the top of the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

10.6.4 Inventory Trajectories

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

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

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

10.6.5 Three Ways the Plant Takes the Handle Away

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

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

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

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

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

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

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

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

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

10.6.6 The Failure That Shows Up in No Temperature

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

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

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

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

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

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

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

10.6.7 The Challenge, and Why Both Intuitions Fail

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

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

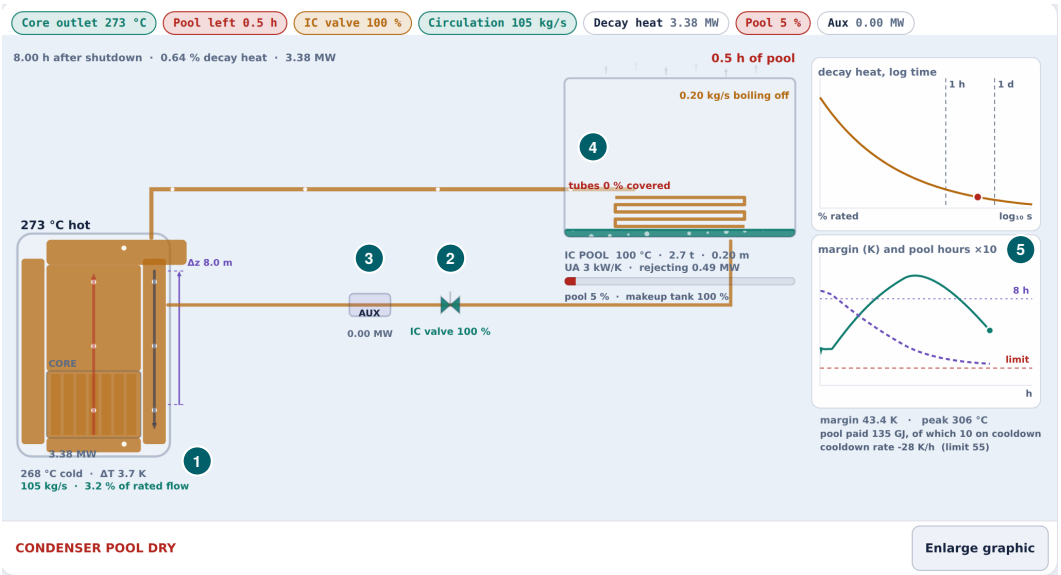

Figure 10.5 The fully passive architecture of scenario 4 at $t = 8.00$ h, with the trends spanning the whole run. Nothing has gone wrong with the temperature. Decay heat is down to 3.38 MW, which is 0.64 % of rated, the core outlet sits at 273 °C on 3.7 K around the loop, and the margin to the limit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

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

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

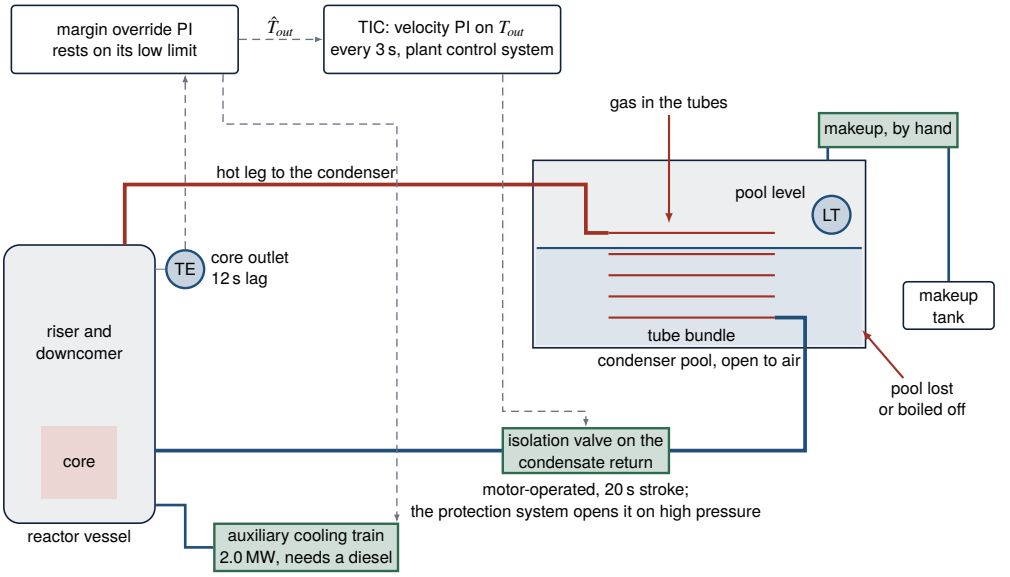

Figure 10.6 The loop of Figure 10.1 with the hardware of a plant. A core outlet thermocouple (TE) feeds the temperature controller TIC, which every 3 s moves the motor-operated isolation valve on the condensate return, and the margin override acts on the auxiliary cooling train from the same reading.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

10.7 Technology in Practice

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

What Measures It

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

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

What Moves It

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

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

Where the Controller Runs

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

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

What Is Modernizing It

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

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

10.8 Limitations

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

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

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

10.9 Discussion

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

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

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

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

10.10 Conclusion

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

KEY POINTS

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

10.11 Exercises

10.1 ○ Concept check · 2 min

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

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

10.2 Concept check · 2 min

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

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

10.3 Concept check · 1 min

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

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

10.4 Theory and calculation · 6 min

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

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

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

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

10.5 Theory and calculation · 6 min

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

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

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

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

10.6 Theory and calculation · 8 min

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

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

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

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

10.7 App challenge · 10 min

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

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

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

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

10.8 App challenge · 8 min

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

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

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

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

10.9 App challenge · 10 min

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

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

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

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

10.10 Application area · 4 min

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

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

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

10.11 Application area · 3 min

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

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

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

10.12 Review the agent · 5 min

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

CHAPTER 11

Waste Vitrification Melter Simulator for Control Across Four Time Scales

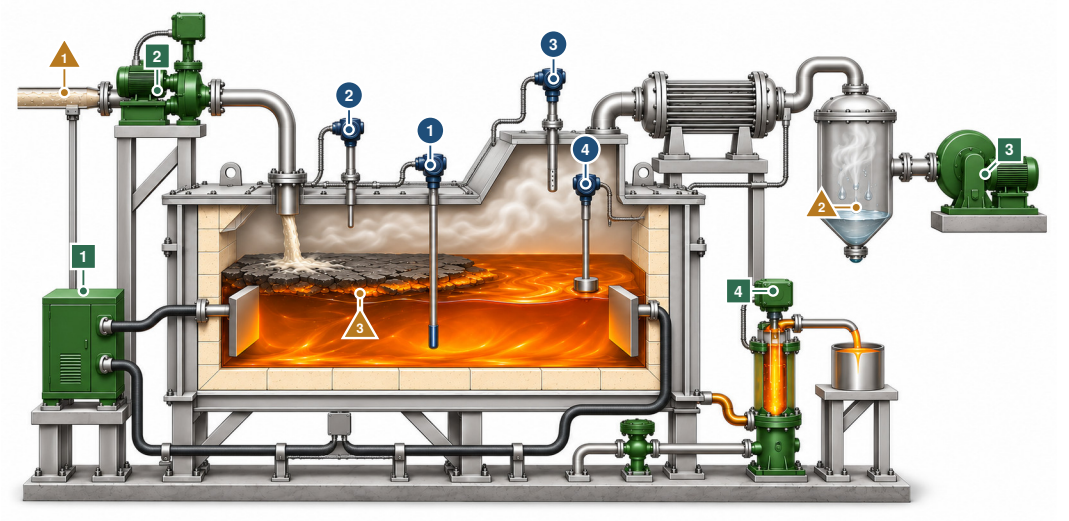

The controller moves 1 electrode power, from the supply cabinet 2 slurry feed pump 3 exhauster that holds the space above the glass under vacuum 4 airlift that pours glass into the canister

The controller reads 1 glass temperature 2 pressure above the glass 3 temperature above the crust, from which its coverage is inferred 4 glass level

What disturbs it 1 a wetter batch of slurry 2 deposits loading the off-gas scrubber 3 foam under the crust

Notice that all the glass is made inside the dark crust, which also blocks the glow of the pool from the space above, so the temperature up there tells how much of the pool the crust covers. Illustration, not to scale.

IN PLAIN TERMS **Making Glass From Waste Under a Floating Crust**

The main idea. A melter turns liquid nuclear waste into glass by spraying it onto a pool of molten glass kept hot by an electric current. The waste floats as a crust while it dries and melts, and all the glass is made inside that crust.

The challenge. The crust also works like a lid that keeps the pool warm. A wetter batch can then cut the output while the pool temperature hardly moves. Four controllers can each do their job well and still leave about a quarter of the output unused.

Why it matters. The glass locks the waste away for a very long time, and a faster melter lets a cleanup finish sooner.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Sections 11.2 and 11.3, the fact boxes of Sections 11.3 and 11.6, the try-it boxes of Section 11.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 11.6, Section 11.7, Section 11.8 and the open challenges

The control question A joule-heated melter makes its glass inside a floating cold cap that also insulates the pool and loads the off-gas train. Why do four correctly tuned loops leave a quarter of the production unused, and what does a feed rate under overrides and feedforward recover?

Learning objectives

- Compute the slurry enthalpy, the glass yield and the evaporation duty of a wetter batch, and explain why power does not set production.
- Compare the electrical gain of a fixed tap with the thermal damping of the pool, and choose the regulation.
- Estimate the time constant of the plenum gas and set it against the dead time of the tap.
- Explain how a low selector with external reset hands the feed between three controllers.
- Measure what coordination, feedforward and MPC each change in production and constraint exposure.

Prerequisites Sections 1.3 (loop action), 1.4.5 (protection layers), 1.6.2 (velocity-form PID), 1.6.3 (predictive control) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/waste-vitrification/

See also Volume 4, Chapter 4 for feed-rate overrides on a low selector, Volume 4, Chapter 6 for a push toward the first binding limit, Volume 4, Chapter 9 for valve-position control. Volume 6, Chapter 7 for the production handle and Volume 6, Chapter 4 for a throughput supervisor.

11.1 Introduction

A joule-heated ceramic melter has no burner and no heating element. Two pairs of Inconel plate electrodes hang in 25.2 t of molten borosilicate at 1150 °C and the current that heats the pool travels through the glass, so the glass is the resistor. Waste slurry mixed with glass formers is sprayed onto the surface, where it does not sink and dissolve. It floats, dries and calcines as a crust two or three centimeters thick called the cold cap, and every kilogram of glass the plant makes is made inside that crust [1, 2]. The crust is therefore the production mechanism. It is also the dominant thermal resistance between the pool and the space above it, because bare glass at 1150 °C radiates 232 kW/m^2 while a crust at 160 °C radiates almost nothing [3]. And it is where every cubic meter of steam and nitrogen oxide is born, so it is the load on the off-gas system as well.

That coincidence is what a block diagram cannot show. A transfer function from electrode power to glass temperature can be written down and tuned against, and it is correct, but it hides the fact that the resistance in its denominator is a state the feed rate creates. The misconception the simulator corrects is measured rather than asserted. Stepping the electrode power up by 140 kW at a fixed feed rate raises the pool by 25.9 °C and leaves glass production at 500 kg/h, approximately where it started, because at steady state the glass rate is bounded by the dry-solids feed multiplied by its oxide yield. What the extra power bought was a smaller crust, coverage falling from 90.8 % to 71.7 %, and a plenum 176 °C hotter carrying the difference away. A second measurement makes the same point from the other side. A feed batch thirteen points wetter costs 37 % of production while the glass temperature moves from 1150.0 °C to 1150.4 °C, because the thicker crust that absorbed the extra evaporation duty of 186 kW also stopped 105 kW of radiation escaping through the bare melt. Glass temperature stays near setpoint while about a third of the production rate is lost.

The class of plant is wider than vitrification. Any process whose dominant resistance is a layer that the control action itself deposits behaves this way, including batch-fed glass furnaces under a batch blanket, rotary kilns under a coating, freeze linings in electric smelting furnaces and fouling layers on any heated surface. Two general consequences follow. The first is that the time constants separate by orders of magnitude, because a gas inventory, a floating crust, a large thermal mass and an accumulating deposit have nothing in common but the plant they sit in. Time-scale separation of that kind is normally exploited to reduce a model [4], and here it has to be exploited structurally instead, because one manipulated variable reaches two loops three decades apart in speed. The second is that no single loop owns the constraint set, so the plant runs against whichever of several limits

is binding today. The standard answer is a production actuator driven by one controller and pulled back by constraint controllers through a selector [5, 6]. The feed-rate loop of the combine harvester in Volume 4, Chapter 4 is built the same way. A melter is an unusually clean example because the constraint that binds first is an actuator demand rather than a measurement.

A student who is handed the four loops of a melter as four independent single-input single-output problems will tune all four correctly and leave a quarter of the plant's output on the table. This chapter models an interactive joule-heated melter in enough detail to be rebuilt, quantifies what coordinating the feed rate against the constraints is worth against independent loops on the same disturbance sequence, and measures the two electrode gains that decide whether the supply is run on power or on voltage. Section 11.2 gives the background of the glass and the machine. Section 11.3 writes out the eleven-state model with its instrument chain, its actuator envelope and its supervisor, to the constant, and lists its assumptions. Section 11.4 measures the four time scales, the two electrode gains, a power step and the pressure loop. Section 11.5 describes the six PI loops, the two classical architectures, the feedforward inversion and the optional predictive controller. Section 11.6 follows the scenario sequence from the electrodes to a graded shift. Its central comparison is 2444 kg of glass against 1968 kg over the same 300 min run, with 0 min above the cold-cap coverage ceiling against 188.3 min, and it shows that the feedforward path is structural rather than cosmetic because it takes the worst pressure excursion from 90.0 Pa to 54.2 Pa. Section 11.7 sets the loops in their plant hardware. Sections 11.8 to 11.10 state the limitations, discuss what transfers to other plants and conclude, and Section 11.11 holds the exercises.

11.2 Vitrification and Its Control Problem

Reprocessing spent fuel for plutonium left the United States with tens of millions of gallons of high-level liquid waste, neutralized with sodium hydroxide and stored in underground carbon steel tanks at Hanford and Savannah River built in the 1940s and 1950s. Tank storage was always an interim answer, and the search for a durable solid began in the 1950s with calcination, with a range of ceramic hosts and with several glass families [7]. Borosilicate glass won for chemical rather than mechanical reasons. It accepts a wide and variable mixture of waste oxides into a single amorphous phase, it is durable in groundwater on a geological time scale, and it tolerates the composition swings a tank farm delivers [8, 9]. Alternatives with better waste loading, including crystalline ceramics and glass-ceramic composites, have never matched it on tolerance to a feed nobody can specify in advance.

The machine that makes the glass took longer to settle than the glass did. Early work used in-can melting and induction-heated metallic melters, both of which are limited by the corrosion of a metal wall by molten borosilicate at 1150 °C. The joule-heated ceramic melter removed the wall from the heating path by making the melt its own resistor, passing alternating current between electrodes immersed in the pool and lining the vessel with refractory rather than metal [7]. The Defense Waste Processing Facility at Savannah River started radioactive operation in 1996 on that design and has run on it since, and the Waste Treatment and Immobilization Plant at Hanford was designed around the same principle at larger scale and with air bubblers added under the melt surface [10]. Bubbling is the single largest step change in throughput the technology has had, and in this simulator it is worth a factor of 2.5 on the melt-to-cap heat transfer coefficient, the parameter that melting-rate correlations fitted to melter feeds and operating conditions treat as the principal lever [11].

Cold-cap behavior became the operating constraint once the melters were large enough that throughput, rather than glass quality, decided whether a cleanup mission would finish. Production is not set by the electrode power, and it is not set by the pool temperature. It is set by the heat that crosses into the crust and by the enthalpy needed to convert what is sitting up there, so the crust is the throughput bottleneck by construction. Laboratory and pilot work has resolved the crust into layers, a wet slurry pool on top, then drying, then calcination, then a reacting foam layer, then the melting front at about 980 °C, and has measured the temperature profile across those few centimeters

directly [2, 12]. The same work identified foaming as the mechanism that destroys melting rate, because a raft of gas-filled glass forms exactly at the interface the heat has to cross [13]. What none of it produced is an instrument. There is no cold-cap sensor on an operating melter, and coverage is inferred from plenum thermocouples through the radiation contrast between bare melt and crust, a relation that computational studies of a pilot-scale melter show to be curved and emissivity dependent [3]. Coverage is consequently an advisory variable on every operating plant, which is precisely the situation that makes it interesting to close a loop on in a simulator.

The electrical side carries its own history. Molten borosilicate conducts by alkali ion migration, so its resistivity follows an Arrhenius law and falls steeply with temperature, and correlations for the viscosity and electrical conductivity of low-activity waste glass melts are fitted over composition and temperature for exactly this reason [14]. A supply holding constant voltage on a falling resistance delivers rising power, which is a positive feedback path around the fastest element in the plant. That is why every operating joule-heated melter regulates power and treats the tap changer as a slave, and Section 11.4.2 measures both gains involved so the design decision can be read off numbers instead of accepted.

> ⧖ HISTORY **Two Melters in Twenty-One Years (1996 to 2017)**
>
> The Defense Waste Processing Facility has vitrified high-level tank waste since 1996 in joule-heated melters whose upper and lower electrodes, each set controlled on its own, pass current through the glass. The first melter served from 1996 to November 2002. The second ran from March 2003 to February 2017 and made over 4.5 million kilograms of glass, twelve years longer than its conservative design life of two years, which rested on experiments that predicted fast wear of the refractory in contact with the glass. Feeding stopped in January 2017 when a pour spout heater failed, and a few days later a lower electrode broke at its stem and leaked cooling water into the glass pool. The stem sat inside the melter shell, so the melter could not be repaired, and a third melter took its place [15].

11.3 Process Model

Melter MLTR-101 is a refractory box with a 10 m^2 melt surface holding 25.2 t of borosilicate glass at a design depth of 0.95 m. Two pairs of plate electrodes on opposite walls carry 1260 kW into the pool at the design point, waste slurry enters through a nozzle in the lid at 1425 kg/h, glass leaves through an airlift into a canister, and the space above the melt is held at −250 Pa gauge by an exhauster pulling through a film cooler and a scrubber train. Figure 11.1 shows the arrangement, the instrument tags and every symbol used in the equations that follow. Reported time and controller periods are in minutes. The energy, mass and pressure balances use seconds before multiplication by 60 in the integrator, while scrubber loading and actuator time constants use minutes. Temperatures are in degrees Celsius, feed and production in kg/h, gas flows in kg/s, pressure in pascals gauge, and electrode quantities in ohms, volts and amperes.

> ✦ DID YOU KNOW? **Ten Square Meters of Molten Glass**
>
> The low-activity waste melters described for Hanford in 2011 have a melt surface of 10 m^2, the size of this chapter's melter, and a design rating of 1.5 t of glass per square meter per day [16]. The chapter's design point makes 1.20 t/m^2 a day.

Table 11.1 classifies what the control system interacts with. Four alarms bound normal operation, a containment limit at −25 Pa above which the plenum can leak outward, a glass temperature floor at 1050 °C, a refractory corrosion ceiling at 1210 °C and a cold-cap coverage ceiling at 97 % beyond which the crust is thickening rather than spreading. None of the four stops the run. Each latches an annunciation and accumulates minutes on the scorecard, so a test that carried the melter outside its limits leaves a record. Eleven states are integrated, seven of them physical and four of them actuator lags.

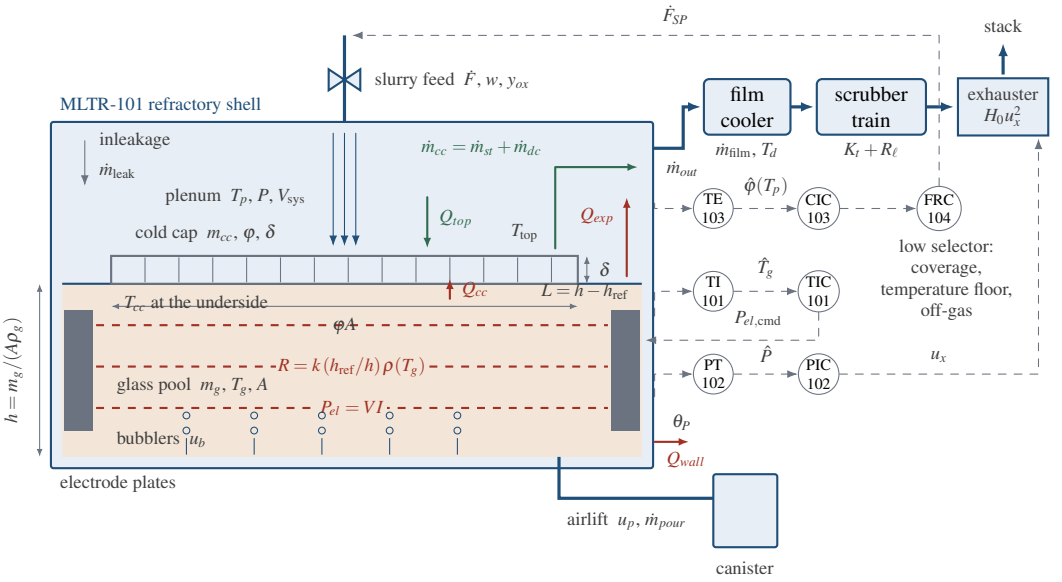

Figure 11.1 Joule-heated ceramic melter MLTR-101 with its off-gas train. Slurry of moisture w enters at rate $\dot{F}$ and forms the cold cap of mass m_{cc}, which covers a fraction φ of the melt area A at mean thickness δ. Current passes between the plate electrodes through the glass itself, so the circuit resistance R carries the pool temperature T_g and the glass depth h. Heat crosses into the cap as Q_{cc}, leaves the uncovered melt as radiation Q_{exp}, returns from the plenum as Q_{top} and leaves the shell as Q_{wall}. The gas the cap makes joins film cooler air and shell inleakage and is drawn through the train resistance $K_t + R_\ell$ by the exhauster at demand u_x. TIC-101 writes electrode power, PIC-102 writes exhauster demand against a pressure tap carrying transport delay θ_P, LIC-105 writes the airlift u_p, and FRC-104 writes the feed rate from the coverage inferred from the plenum temperature.

11.3.1 Glass Pool and the Cold Cap

The pool is one well-mixed inventory heated by the electrodes and cooled by four paths, as shown in Equation (11.1).

$$m_g c_p \frac{\mathrm{d}T_g}{\mathrm{d}t} = P_{el} - Q_{cc} - Q_{exp} - Q_{wall} - Q_{sens} \tag{11.1}$$

where m_g is the glass inventory in kg, c_p is the glass heat capacity in kJ/(kg K), P_{el} is the delivered electrode power in kW, Q_{cc} is the duty crossing into the cold cap, Q_{exp} is radiation from the uncovered melt into the plenum, $Q_{wall} = UA_m(T_g - T_a)$ is the shell loss and Q_{sens} is the sensible heat needed to raise newly made glass from the cap underside to pool temperature. The product $m_g c_p$ is 31.5 MJ/K at the design inventory, and the four loss terms together shed 6.84 kW/K about the design point, which is measured directly in Section 11.4.2.

The cold cap is one lumped mass whose coverage and thickness follow from it, as shown in Equation (11.2).

$$\varphi = 1 - \mathrm{e}^{-m_{cc}/M_{\mathrm{ref}}}, \qquad \delta = \frac{m_{cc}}{A\rho_{cc}\varphi} \tag{11.2}$$

where M_{ref} is the spreading mass in kg, A is the melt area and ρ_{cc} is the bulk density of the crust. The crust spreads before it thickens, which is the choice that gives the layer its characteristic asymmetry. Coverage responds quickly when the crust is small and slowly when it is nearly complete, and Section 11.4 measures the factor of five between the two.

Table 11.1 Classification of Variables. Nominal Values Are Those of the Design Operating Point at Which Every Scenario Starts.

Symbol	Role	Nominal	Range
T_g (TI-101)	CV, state	1150.0 °C	300–1600 °C
P (PT-102)	CV, state	−250 Pa	−6000 to 3000 Pa
φ	CV, inferred	90.8 %	60–96 % SP
L	CV, measured	0 mm	−50 to 50 mm SP
P_{el}	MV	1260 kW	0–1800 kW
V	MV (volt mode)	664 V	0–750 V
u_x	MV	66.6 %	0–100 %
$\dot{F}$	MV or DV	1425 kg/h	0–2600 kg/h
u_p	MV	13.9 %	0–100 %
w	DV, measured	55 %	35–75 %
y_{ox}	DV, unmeasured	0.78	0.60–0.90
u_b bubblers	DV, measured	1×	0–1×
s sulfate	DV, unmeasured	0×	0–1×
η_b exhauster	DV, measured	1×	0.2–1×
η_s supply	DV, measured	1×	0.5–1×
κ train	DV, measured	1×	1–3×
λ inleakage	DV, unmeasured	1×	0.4–3×
m_{cc}	state	215 kg	0–40 000 kg
T_p	state	436 °C	25–1400 °C
m_g	state	25 175 kg	8000–40 000 kg
f foam	state	0	0–1
R_ℓ	state	0 Pa s^2/kg^2	0–200 000

The heat crossing into the crust drives everything downstream of it, and it is written in Equation (11.3).

$$Q_{cc} = h_{cc}\left[\beta_0 + (1-\beta_0)u_b\right]\left(1 - k_f f\right) A\,\varphi\,\max(0, T_g - T_{cc}) \tag{11.3}$$

where h_{cc} is the melt-to-cap heat transfer coefficient in kW/(m^2 K), u_b is the bubbler air fraction, β_0 is the unbubbled fraction of the coefficient, k_f is the loss at full foam, f is the foam index and T_{cc} is the cap underside temperature at which the batch dissolves. Bubbling is therefore worth $1/\beta_0 = 2.5$ on the coefficient and a full foam raft takes 55 % of it back, both of which follow the direction of the effects reported for melter feeds [11, 13].

> ⧖ HISTORY **Four Bubblers for a Bottleneck (2010)**
>
> Until 2010 the same facility had never filled more than 250 canisters of glass in a year, and most years averaged 215. A new contract to close 22 waste tanks in eight years called for 400, a melt rate the melter as built could not reach. In the late summer of 2010 four bubblers were lowered through nozzles in its top head, two of them doubling as the glass thermowell and the level probe. Started on argon on 28 September 2010, they raised the instantaneous melt rate from about 59 to about 91 kg/h. They also changed how the pool behaved. With both electrode sets in manual at fixed output, the first argon flows dropped the lower pool temperature by about 140 °C in less than an hour, so the electrode power is now adjusted whenever a bubbler flow changes [17].

11.3.2 Conversion Inside the Crust

The order of the two conversion steps matters more than either rate. Slurry lands on a crust whose top is already at a couple of hundred degrees, so the water flashes within seconds, and only what is left of the heat afterward calcines and melts the solids. Equation (11.4) writes it that way.

$$\begin{aligned} \dot{m}_{st} &= \min\left(\dot{F}w,\ \frac{Q_{cc}+Q_{top}}{h_w}\right) \\ \dot{m}_{cons} &= \frac{\max(0,\ Q_{cc}+Q_{top}-\dot{m}_{st}h_w)}{h_d} \\ \dot{m}_{glass} &= y_{ox}\,\dot{m}_{cons}, \qquad \dot{m}_{dc} = (1-y_{ox})\,\dot{m}_{cons} \\ \frac{\mathrm{d}m_{cc}}{\mathrm{d}t} &= \dot{F}(1-w) + \left(\dot{F}w - \dot{m}_{st}\right) - \dot{m}_{cons} \end{aligned} \tag{11.4}$$

where $\dot{F}$ is the slurry feed in kg/s, w is the mass fraction of water, h_w is the enthalpy to boil and superheat feed water, h_d is the enthalpy to calcine and melt dry solids, y_{ox} is the oxide yield of the dry solids, $\dot{m}_{st}$ is the steam rate, $\dot{m}_{dc}$ is the carbon dioxide and nitrogen oxide released and $\dot{m}_{glass}$ is the glass production rate. Tying the steam rate to the melting front instead, which is the obvious lumped simplification, makes a wetter batch produce less steam because less material is melting, and that is the opposite of what a melter does. The enthalpy and yield of a kilogram of slurry follow from Equation (11.5).

$$\Delta h_{feed} = w\,h_w + (1-w)\,h_d, \qquad Y = (1-w)\,y_{ox} \tag{11.5}$$

At 55 % water the melter needs 2788 kJ per kilogram of slurry and gets 0.351 kg of glass from it, while at 68 % it needs 3028 kJ and gets 0.250 kg. A wetter slurry costs more energy per kilogram and contains less glass per kilogram, which is why one feed change moves production so far.

11.3.3 Radiation and the Plenum

The two radiation terms trade against each other as coverage moves, and they are given in Equation (11.6).

$$\begin{aligned} Q_{exp} &= A(1-\varphi)\,\varepsilon_m\,\sigma\left[(T_g+T_0)^4-(T_p+T_0)^4\right] \\ Q_{top} &= A\varphi\,\varepsilon_{cc}\,\sigma\left[(T_p+T_0)^4-(T_{\mathrm{top}}+T_0)^4\right] \end{aligned} \tag{11.6}$$

where ε_m is the effective emissivity of the exposed melt, ε_{cc} is the absorptivity of the cap top, T_{top} is the cap top surface temperature, $T_0 = 273.15\,\mathrm{K}$ converts to absolute temperature and σ is the Stefan-Boltzmann constant. Both are in kW after dividing by 1000. The magnitude is the point. Black-body emission at 1150 °C is 232.6 kW/m^2, so a receding crust appears on the plenum thermocouples within a minute, and Q_{top} is roughly a tenth of the melting energy at design coverage, which is why plenum temperature and melting rate are not independent [3]. The plenum itself is a refractory heat capacity

with radiation in, radiation out, gas enthalpy and a shell loss, as shown in Equation (11.7).

$$C_p \frac{\mathrm{d}T_p}{\mathrm{d}t} = Q_{exp} - Q_{top} - UA_p(T_p - T_a) + \dot{m}_{cc} c_{pv}(T_{ex} - T_p) + \dot{m}^{+}_{leak} c_{pa}(T_a - T_p) \tag{11.7}$$

where C_p is the plenum refractory heat capacity in kJ/K, c_{pv} and c_{pa} are the cap off-gas and air heat capacities, T_{ex} is the temperature of the gas leaving the cap, $\dot{m}^{+}_{leak}$ is inleakage when it is inward and UA_p is the lid loss.

11.3.4 Plenum Pressure and the Off-Gas Train

Pressure is a gas inventory worked against an exhauster, and it is the fastest state in the plant, as shown in Equation (11.8).

$$\frac{\mathrm{d}P}{\mathrm{d}t} = \frac{R_g(T_d + T_0)}{V_{\mathrm{sys}}} \left(\dot{m}_{cc} + \dot{m}_{leak} + \dot{m}_{\mathrm{film}} - \dot{m}_{out}\right) \tag{11.8}$$

$$\dot{m}_{out} = \sqrt{\frac{\max(0,\ P + \eta_b H_0 u_x^2/10^4)}{\kappa\,(K_t + R_\ell)}}$$

$$\dot{m}_{leak} = -\operatorname{sign}(P)\,\lambda\, C_{\mathrm{leak}} \sqrt{|P|} \tag{11.9}$$

where R_g is the specific gas constant of the off-gas, T_d is the duct temperature after the film cooler, V_{sys} is the gas volume of the plenum, cooler, duct and scrubber together, H_0 is the exhauster head at full demand, u_x is the exhauster demand in percent, K_t is the clean train resistance, R_ℓ is the accumulated scrubber loading and C_{leak} is the shell and nozzle inleakage coefficient. Three disturbance multipliers appear here and all three default to unity, namely the exhauster availability η_b, the train restriction κ over and above the accumulated loading, and the inleakage multiplier λ. The coefficient $R_g(T_d + T_0)/V_{\mathrm{sys}}$ evaluates to 1524 Pa/kg. Both sides of (11.8) carry units of Pa/s before the conversion to minutes. The sign in (11.9) reverses when the plenum goes positive, which is contamination leaving the melter rather than air entering it, and at the design vacuum the melter draws 541 kg/h of air in through its own seams. That inward flow supplies most of the loop's self-regulation, because $\partial \dot{m}_{leak}/\partial P$ is -3.00×10^{-4} kg/(s Pa) against 9.38×10^{-5} kg/(s Pa) for the exhauster.

Scrubber loading is the slowest state on the plant and it has no measurement of its own, as shown in Equation (11.10).

$$\frac{\mathrm{d}R_\ell}{\mathrm{d}t} = k_{\mathrm{foul}} \frac{\dot{m}_{cc}}{\dot{m}_{cc,0}} \left(1 + k_c f\right) \gamma \tag{11.10}$$

where k_{foul} is the loading rate at design carryover, $\dot{m}_{cc,0} = 0.2083\,\mathrm{kg/s}$ is the design carryover, k_c is the extra carryover at full foam and γ is an operator-set acceleration used so that a week of loading can be watched in an hour. At design carryover and $\gamma = 1$ the train reaches the 6500 Pa s^2/kg^2 of scenario 8 in 6.1 days.

11.3.5 Electrodes and the Power Envelope

The glass is the resistor, and its resistivity is Arrhenius in absolute temperature, as shown in Equation (11.11).

$$\rho(T_g) = \rho_0\, \mathrm{e}^{E_a/R_u(T_g + T_0)}, \qquad R = k\,\frac{h_{\mathrm{ref}}}{h}\,\rho(T_g) \tag{11.11}$$

$$P_{el} \le \min\left(\eta_s P_{\max},\ \frac{(\eta_s V_{\max})^2}{10^3 R},\ \frac{(\eta_s I_{\max})^2 R}{10^3}\right)$$

where ρ_0 is the pre-exponential in Ω m, E_a/R_u is the activation temperature in K, k is the electrode cell constant at the design level in 1/m, h_{ref} is the design glass depth, $h = m_g/(A\rho_g)$ is the actual

depth, and $V_{\max}$, $I_{\max}$ and $P_{\max}$ are the transformer voltage, electrode current and supply ratings, derated together by the supply availability η_s. The cell constant carries h_{ref}/h because the conducting cross-section between plates on opposite walls is proportional to the submerged depth, so a canister pour is an electrical disturbance as well as a hydraulic one. At the design point ρ is 0.0501 Ω m and R is 0.3504 Ω, which places the supply at 664 V and 1896 A for 1260 kW against an envelope of 1605 kW. The Arrhenius form used here is a single lumped law calibrated for a low-activity waste borosilicate, and correlations fitted over composition as well as temperature are available [14].

11.3.6 Inventory, Foam and the Actuators

Glass inventory, the foam raft and the four actuators complete the state vector, as shown in Equation (11.12).

$$\begin{aligned} \frac{\mathrm{d}m_g}{\mathrm{d}t} &= \dot{m}_{glass} - \dot{m}_{pour}, \qquad L = 10^3\left(\frac{m_g}{A\rho_g} - h_{\mathrm{ref}}\right) \\ \tau_f \frac{\mathrm{d}f}{\mathrm{d}t} &= \frac{s}{1+\mathrm{e}^{-(T_g - T_{\mathrm{foam}})/w_{\mathrm{foam}}}} - f \\ \tau_i \frac{\mathrm{d}a_i}{\mathrm{d}t} &= a_{i,\mathrm{cmd}} - a_i, \qquad i \in \{P_{el},\ u_x,\ \dot{F},\ u_p\} \end{aligned} \tag{11.12}$$

where $\dot{m}_{pour}$ is the airlift discharge, L is the level deviation from the design depth in mm, s is the sulfate loading, T_{foam} is the foaming onset temperature, w_{foam} is the onset width, τ_f is the foam time constant and τ_i is the lag of actuator i. The airlift discharge is $\dot{m}_{pour} = (u_p/100)\dot{m}_{pour,max}$, with $\dot{m}_{pour,max} = 1\,\mathrm{kg/s}$ (3600 kg/h). The level is a pure integrator with no self-regulation whatever, and the measured ramp is −0.0222 mm/min of level for every percent of airlift demand.

11.3.7 Instruments

Glass temperature, pressure, coverage, moisture and level reach the controllers. Plenum temperature and glass level are available without measurement noise or an additional instrument lag. The glass thermocouple is a first-order lag of 0.60 min with optional colored noise, and feed moisture comes from an analyzer with a 0.05 min lag. The pressure transmitter is a transport delay of 0.07 min followed by a lag of 0.010 min, and the delay is real, because the tap is an air-purged dip tube through a hot lid with a long impulse line to a transmitter outside the shielded cell. Coverage has no instrument and is inferred from the plenum thermocouple by a soft sensor through the two-point calibration of Equation (11.13).

$$\hat{\varphi} = \mathrm{clamp}\left(\frac{T_{p,\mathrm{bare}} - T_p}{T_{p,\mathrm{bare}} - T_{p,\mathrm{full}}},\ 0,\ 1\right) \tag{11.13}$$

where $T_{p,\mathrm{full}}$ is the plenum temperature under a complete cap and $T_{p,\mathrm{bare}}$ is the temperature over a bare melt. The result is filtered with a time constant of 3.5 min. The underlying relation is curved and this reads it as a straight line, so the soft sensor carries a real bias, which is the reason coverage is advisory on a real melter [3] and the reason a scenario exists in which the inference is fooled.

11.3.8 Initialization, Numerics and Constraints

Table 11.2 lists the constants. The physical states are initialized from a steady-state solve for the scenario's feed, power and disturbances. This avoids marching the slow pool to equilibrium, but does not put every automatic controller at equilibrium. Coordinated control sets its temperature target during the initial scan, which can create an immediate power demand. The steady state is solved as a nested pair. The inner problem is a damped two-by-two Newton iteration on the glass and plenum temperatures with the cap mass held fixed, multi-started from 1150 °C, 1280 °C and 980 °C when the warm start fails, and the outer problem is a bisection in $\log m_{cc}$ until what the melter melts equals what is being fed to it. Consumption is monotone in cap mass, so the bisection cannot fail, and failure at the top of the bracket is itself the answer. The melter is over-fed, no steady state exists, and the

Table 11.2 Model Parameters. Selected Physical and Instrument Constants. Hidden Entries Are Unavailable in the Model Panel.

Symbol	Meaning	Value	Unit	Symbol	Meaning	Value	Unit
Adjustable, shown in the model panel							
A	melt surface area	10.0	m^2	T_{ex}	cap off-gas temperature	600	°C
h_{ref}	design glass depth	0.95	m	V_{sys}	off-gas gas volume	150	m^3
ρ_g	glass density	2650	kg/m^3	T_d	duct temperature	380	°C
c_p	glass heat capacity	1.25	kJ/(kg K)	K_t	clean train resistance	8115	$Pa\,s^2/kg^2$
UA_m	shell heat loss	0.090	kW/K	H_0	exhauster head at full demand	8455	Pa
E_a/R_u	activation temperature	12085	K	$\dot{m}_{film}$	film cooler air	0.25	kg/s
k	electrode cell constant	7.0	1/m	C_{leak}	inleakage coefficient	0.0095	$kg/(s\,Pa^{0.5})$
V_{max}	transformer voltage limit	750	V	C_p	plenum heat capacity	420	kJ/K
I_{max}	electrode current limit	2600	A	UA_p	plenum shell loss	0.120	kW/K
P_{max}	supply rating	1800	kW	k_{foul}	train loading rate	0.60	$Pa\,s^2/(kg^2\,min)$
M_{ref}	cap spreading mass	90	kg	k_c	carryover at full foam	2.5	—
ρ_{cc}	cap bulk density	1100	kg/m^3	$\tau_{P_{el}}$	power supply lag	0.10	min
h_{cc}	melt-to-cap coefficient	0.66	$kW/(m^2\,K)$	τ_{u_x}	damper travel	0.04	min
T_{cc}	cap underside temperature	980	°C	$\tau_{\dot{F}}$	feed transport lag	0.25	min
T_{top}	cap top temperature	160	°C	τ_{u_p}	airlift response	0.30	min
ε_m	exposed-melt emissivity	0.55	—	τ_{T_g}	thermocouple lag	0.60	min
ε_{cc}	cap top absorptivity	0.75	—	θ_P	pressure tap delay	0.07	min
β_0	unbubbled fraction	0.40	—	$\tau_{\hat{\varphi}}$	coverage filter	3.5	min
k_f	heat-transfer loss at full foam	0.55	—	$T_{p,full}$	plenum, complete cap	365	°C
τ_f	foam time constant	12	min	$T_{p,bare}$	plenum, bare melt	1160	°C
T_{foam}	foaming onset temperature	1120	°C	*Hidden, never displayed*			
w_{foam}	foaming onset width	28	°C	ρ_0	resistivity pre-exponential	1.027×10^{-5}	Ω m
h_w	feed water enthalpy	3620	kJ/kg	T_a	ambient temperature	25	°C
h_d	dry solids enthalpy	1770	kJ/kg	R_g	off-gas gas constant	350	J/(kg K)
				c_{pv}	cap off-gas heat capacity	2.10	kJ/(kg K)
				c_{pa}	air heat capacity	1.05	kJ/(kg K)

crust grows until it bridges the plenum. That boundary is the frontier drawn on the operating map of Figure 11.5. The engine of Section 1.6.1 integrates the model with classical fourth-order Runge-Kutta at a fixed step of 0.005 min (0.3 s), the controllers execute every 0.01 min (0.6 s), and the historian declares one sample every seven physics steps (2.1 s), giving approximately 28.6 samples per minute. The default seed is 20 260 118.

The constraints the supervisor and the operating map enforce are a glass temperature floor of 1050 °C, a refractory ceiling of 1210 °C with a guard band of 12 °C below it, coverage between 55 % and 97 %, a plenum temperature limit of 750 °C, an exhauster demand ceiling of 95 % and the supply envelope of (11.11).

11.3.9 Assumptions

Six simplifications carry the model, and each one fails somewhere specific.

One lumped cold cap. A real crust is a layered structure whose relative zone thicknesses set the melting rate [2, 12], and here all of it sits inside one heat transfer coefficient and two enthalpies. The model can therefore show what a thicker cap does, but not why one batch forms a thicker cap than another at the same water content.

Coverage is exponential in cap mass. The shape in (11.2) gives the layer the right asymmetry, fast when small and sluggish when nearly complete, but the constant of 90 kg for 63 % of a 10 m^2 surface is a calibration rather than a measurement.

One Arrhenius resistivity and one electrode pair. A real melter has several electrode pairs on separate transformers, current that leaks into the cold cap and along the refractory, and a resistivity that depends on alkali content as much as on temperature [14]. Equation (11.11) would mispredict the current split on a multi-pair machine entirely.

One resistance and one gas volume for the whole off-gas train. The real train is a film cooler, a submerged-bed scrubber, a wet electrostatic precipitator, mist eliminators, filters and a caustic

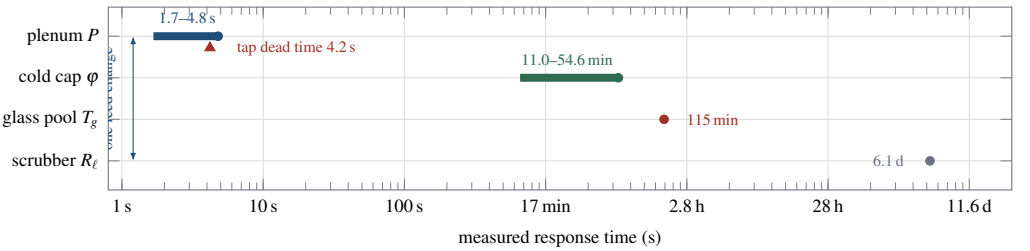

Figure 11.2 Response times measured with damper, feed and power tests on the current model. A moisture disturbance couples to all four states. The scrubber entry is a time to reach the loading of a week-old train rather than a first-order constant, and the cold-cap bar runs between 57.8 and 90.8 % coverage.

scrubber, each with its own volume and condensation. The single R_ℓ of (11.10) reproduces the control-relevant behavior, authority disappearing from the pressure loop over days, and nothing else. Off-gas is also one ideal gas at a fixed duct temperature, whereas the steam fraction really swings from a third to two thirds with feed moisture.

Foaming is one first-order state. Real foaming depends on redox state, sulfate solubility, glass composition and bubbler action together, is strongly hysteretic and is not reliably predictable [13]. The sigmoid in (11.12) reproduces the onset and the insulating consequence and nothing about the chemistry.

Volatilization is absent. An exposed melt on a real machine loses cesium, technetium and halides to the off-gas, which is much of why the coverage floor exists. Here the floor is imposed as a constraint rather than emerging from a mass balance, and the model cannot be used to argue about volatile retention.

11.4 Dynamic Analysis

11.4.1 Four Time Scales on One Plant

The four principal dynamics were measured on the shipped model by stepping one input and reading the 63 % time against a stated late observation. The feed-step baselines were held for 500 min and the responses were observed for another 500 min. Plenum pressure reaches 63 % of an open-loop damper step in 4.8 s, against a gas-inventory time constant of 1.67 s computed from V_{sys} and the self-regulation of (11.8). Cold-cap coverage reaches 63 % of a +50 kg/h feed step in 54.6 min at the design coverage of 90.8 % and, for a +100 kg/h step, 11.0 min at 57.8 % coverage, a factor of 5.0 from the same equation at two operating points. Glass temperature reaches 63 % of a +140 kW power step in 115 min. Scrubber loading reaches the resistance of a loaded train in 6.1 days at design carryover. Figure 11.2 places all four on one logarithmic axis.

The span from 1.7 s to 6.1 days is a factor of 3.1×10^5, and what makes it a control problem rather than a modeling curiosity is that one disturbance excites all four. A batch thirteen points wetter raises the steam rate within one scan, thickens the crust over two hours, changes the pool energy balance over the same two hours and loads the scrubber for the rest of the campaign.

11.4.2 Joule Heating Is Open-Loop Unstable on Voltage Control

The Arrhenius resistivity of (11.11) makes the electrode circuit a feedback path in its own right, and whether that path is positive depends entirely on which quantity the supply regulates. Two gains decide the question, and both were measured on the shipped model rather than estimated.

The first is the electrical gain. Resistance falls by a factor of 1.82 between 1100 °C and 1200 °C, so at a fixed tap of 664 V the delivered power V^2/R changes by 7.52 kW/K about the design point. The second is the thermal gain. Stepping the delivered power by 10 kW on power regulation and

solving for the new steady state moves the pool by 1.462 °C, so the melter sheds 6.84 kW/K. The electrical feedback is larger than the plant's steady-state thermal damping by a factor of 1.10. The fixed-voltage and fixed-power trajectories below demonstrate the resulting difference in open-loop behavior.

The trajectories distinguish the two regulation modes. Taking 34 kW out on power regulation moves the pool from 1150.0 °C to 1145.3 °C and it stops there. Lowering the voltage tap from 664 V to 655 V starts a walk that never stops, and the pool reaches 1129.3 °C after five hours, 1104.1 °C after ten, 1053.3 °C after seventeen and 1020.0 °C after twenty-five, with the delivered power falling from 1258 kW to 338 kW and glass production reaching zero. Nobody would identify that as an instability inside a single shift, which is why the design decision is made structurally instead of by observation. Every operating joule-heated melter regulates power and runs the tap changer as a slave.

🕮 **Pause.** Try Exercise 11.3 before reading on.

The same law has a second consequence that runs the wrong way. The deliverable power of (11.11) is bounded by the transformer voltage at low temperature, where the glass is resistive, and by the electrode current at high temperature, where it is conductive. Because resistance rises as the pool cools, the envelope collapses when the melter needs it most, falling from 1605 kW at 1150 °C to 1178 kW at 1100 °C, 845 kW at 1050 °C and 590 kW at 1000 °C. A cold melter cannot heat itself back up, and the recovery procedure is to stop the feed rather than to turn up the power. The application trends the envelope on the same strip as the power demand for that reason, and reports the margin between them on the scorecard.

11.4.3 Electrode Power Does Not Set Production

The steady-state consequence of (11.4) is worth measuring explicitly because it reverses the intuition a heated reactor teaches. Stepping the electrode power from 1260 kW to 1400 kW at a fixed feed rate of 1425 kg/h raises the pool by 25.9 °C, or 0.185 °C/kW, read 800 min after the step. The native schedule removes the step at minute 400, where the rise is 24.4 °C. The extended test restores 1400 kW immediately after that removal and holds it through minute 820. It is not the unmodified scenario trajectory. The first five minutes run at 0.253 °C/min, which is exactly what 140 kW does to a 31.5 MJ/K pool with nothing else happening, and the remaining hundred minutes are the cold cap catching up. Coverage falls from 90.8 % to 71.7 % and the plenum rises from 436 °C to 612 °C, while glass production moves from 500.2 kg/h to 500.4 kg/h. The transient briefly increases melting as stored cold-cap material is consumed. Once the cap inventory settles, production returns to the rate set by the feed and oxide yield.

Feed rate sets production. Power, temperature and coverage between them decide only how much feed the melter can get away with, and every control decision in Section 11.5 follows from that sentence. In plantwide terms the feed pump is the throughput manipulator, the same choice of production handle that Volume 6, Chapter 7 makes for a reactor and its recycle. It also explains why the temperature setpoint has to move when production is maximized. Melting rate is $h_{cc} A \varphi (T_g - T_{cc})$ by (11.3), so holding both the temperature setpoint and the coverage target fixed pins the melting rate and leaves the feed rate free to decide only how thick the crust gets. At 1150 °C with the cap exactly on its coverage ceiling the absolute ceiling is 1484 kg/h of feed and 521 kg/h of glass, and everything beyond that comes from running hotter.

🕮 **Pause.** Try Exercise 11.1 before reading on.

11.4.4 The Pressure Loop and Its Dead Time

Pressure is the opposite extreme and the loop that keeps containment. A closed-loop setpoint step from −250 Pa to −400 Pa moves the damper by 5.01 %, giving a process gain of −29.9 Pa/%, and the pressure reaches 63 % of the step in 9.3 s and settles inside 10 Pa in 38.7 s. The small-signal steady derivative taken directly on the model is −26.8 Pa/%, and the difference between the two readings is the curvature of the square-root characteristic in (11.8). What limits the tuning is neither of these. It

Table 11.3 Implemented Loops. Direction $a = +1$ Is Reverse Acting and $a = -1$ Is Direct Acting. All Six Execute Every 0.01 min With $\tau_D = 0$.

Loop	CV → MV	a	K_c	τ_I (min)
TIC-101	$\hat{T}_g \to P_{el}$	+1	25 kW/°C	45
PIC-102	$\hat{P} \to u_x$	−1	0.015 %/Pa	0.13
CIC-103	$\hat{\varphi} \to \dot{F}$	+1	1800 kg/h/φ	25
LIC-105	$L \to u_p$	−1	4.5 %/mm	40
OVR-T	$T_{g,\min} \to \dot{F}$	−1	60 kg/h/°C	25
OVR-X	$u_{x,\max} \to \dot{F}$	+1	200 kg/h/%	6

is the 4.2 s of transport delay in the purged tap, which is of the same order as the entire closed-loop response, so the high controller gains require verification against this delay and the chosen integral time. Feedback cannot beat a dead time, and Section 11.5.3 shows what feedforward does instead.

11.5 Control Structure

11.5.1 The Six PI Controllers as Implemented

Every loop uses the velocity-form PID of Section 1.6.2. All six have $\tau_D = 0$, so in this chapter's symbols the proportional and integral recursions are Equations (11.14) and (11.15).

$$
\begin{aligned}
e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k)\,, \quad P_k = K_{c,k} e_k \\
s_k &= s_{k-1} + P_k - P_{k-1} \\
c_k &= u_{\mathrm{bias},k} + s_k + I_{k-1} \\
u_k &= \mathrm{clamp}\,(c_k,\ u_{\min},\ u_{\max})
\end{aligned}
\tag{11.14}
$$

$$
I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c \tag{11.15}
$$

where $a = \pm 1$ selects the loop direction, K_c is the controller gain in units of actuator per measurement, τ_I is the integral time in minutes, u_{bias} is the feedforward bias described below, T_t is the back-calculation tracking constant defaulting to τ_I, and Δt_c is the 0.01 min controller period. The running sum s_k is S_k of Equation (1.14) and the output before limiting c_k is v_k of Equation (1.15). Back-calculation anti-windup and manual tracking are those of Section 1.6.2. Tracking makes a loop's output agree with its manual station, although a simultaneous change of control architecture can also change the supervisor's target. Table 11.3 lists the six loops with their directions and default settings in the convention of Section 1.3.

The two override directions repay a sentence each, because both are easy to get backwards. OVR-T is direct acting on the feed, because the pool being below its floor has to reduce the feed and the process gain from feed to temperature is negative on a cool pool, where the floor can bind. Near the design point that gain is small and changes sign, as the identified path of Table 11.4 shows. OVR-X is reverse acting, because more feed is more off-gas is more damper, so an exhauster demand above its ceiling has to reduce the feed. Actuator limits are 0 kW to 1800 kW on power, 0 % to 100 % on the damper and the airlift, and 0 kg/h to 2600 kg/h on the feed pump. No loop carries a rate limit, because the four actuator lags of (11.12) supply the physical rate limiting.

11.5.2 Classical Architectures

MANUAL and AUTO decide whether the loops are in service. The architecture selector decides how they are wired to each other once they are. Independent loops and coordinated control remain available alongside the MPC described in Section 11.5.4. In MANUAL every actuator is whatever the operator typed and every controller tracks it, so the transfer either way is bumpless.

Independent loops are what a melter is commissioned with. Temperature holds itself on power, pressure on the exhauster, level on the airlift, and the feed rate is a number an operator typed. It is stable, slow and conservative, because the operator picks a rate that is safe on the worst batch in the tank and the melter then runs at that rate on every other batch as well.

Coordinated control makes the feed rate a manipulated variable. The coverage controller CIC-103 pushes the feed up toward its target, and the two override controllers pull it back down through a low selector, as shown in Equation (11.16).

$$\dot{F}_{SP} = \min\Big(\mathrm{PI}_{\varphi},\ \mathrm{PI}_{T},\ \mathrm{PI}_{x},\ \dot{F}_{\mathrm{op}}\Big) \tag{11.16}$$

where PI_{φ} is the coverage controller output, PI_{T} is the temperature-floor override, PI_{x} is the exhauster-demand override and $\dot{F}_{\mathrm{op}}$ is the operator's feed target, which becomes a hard ceiling rather than a command. The card on screen names which of the four is writing the feed rate at every instant. The same supervisor also walks the temperature setpoint toward the refractory guard band at 0.6 °C/min, clamped between the operator's own setpoint and $T_{g,\max} - 12\,^{\circ}\mathrm{C}$, because Section 11.4 showed that a fixed temperature pins the melting rate. The electrode dryer of Volume 4, Chapter 6 pushes its line speed toward the first binding limit by a similar model-based supervisor. The target it walks toward is the most productive point that satisfies every constraint, found by a 320-point sweep of the steady-state solver over the feed and power plane. That sweep is re-solved at most every 6 min and only when a rounded key built from the measured disturbances and the parameters has changed, and it runs on measured quantities rather than true ones, so it does not know that the exhauster has tripped.

External reset drains the losing controllers' windup through finite back-calculation. After each PI has accounted for its own output clamp, the low selector adds the correction

$$\Delta I_{\mathrm{selector}} = \frac{u_{\mathrm{app}} - u_{\mathrm{PI}}}{T_t}\Delta t, \qquad T_t = \tau_I. \tag{11.17}$$

where u_{PI} is the clamped output u_k of a losing controller and u_{app} is the feed the selector applies. The proportional response remains available while the integral follows the selected feed over a finite time. Manual operation instead tracks each controller to the operator's applied feed for a bumpless return to Auto. OVR-X reads delivered damper travel, not the preceding command, because that command already contains feedforward from the selector's own feed output. Reading the command would close a high-gain loop around the selector and cause scan-to-scan chatter.

11.5.3 Feedforward as an Exact Model Inversion

Feed rate reaches plenum pressure in about two seconds and glass temperature in about two hours, so a supervisor that moves the feed to chase production hands the pressure loop a step disturbance it can only discover after the plenum has told a transmitter 4.2 s downstream of it. The feedforward path is the exact inverse of (11.8) evaluated on measured quantities, as shown in Equation (11.18).

$$\hat{\dot{m}}_{in} = \frac{\dot{F}_{\mathrm{cmd}}}{3600}\Big[\hat{w} + (1-\hat{w})(1-y_{ox})\Big] + \dot{m}^{+}_{leak} + \dot{m}_{\mathrm{film}}$$
$$u_{x,ff} = 100\sqrt{\frac{\hat{K}_t\,\hat{\dot{m}}_{in}^{2} - P_{SP}}{H_0}} \tag{11.18}$$

where $\dot{F}_{\mathrm{cmd}}$ is the feed rate the pump is being asked for, $\hat{w}$ is the moisture the analyzer reports and $\hat{K}_t$ is the train resistance computed from the duct venturi and the differential-pressure cell across the scrubber. The result enters PIC-102 as u_{bias} in (11.14), so it moves the damper in the same scan the feed moves. An equivalent inversion of (11.1) is available as a bias on TIC-101 and is worth much less, because that loop has two hours to notice. Nothing in (11.18) knows that the exhauster

Table 11.4 Nominal Predictive Paths. Time Constants and Delays Are in Minutes. Gains Use Physical Output per kW or per kg/h.

Path	Gain	Lags	Delay[a]	RMS[b]
$P \to \hat{T}_g$	0.144762	77.491	0	0.00350
$F \to \hat{T}_g$	0.003291	28.403	3	0.00218
$P \to \hat{\varphi}$	−0.001498	84.593, 38.547	5	$2.46\,10^{-5}$
$F \to \hat{\varphi}$	0.001043	43.521, 4.625	5	$5.92\,10^{-5}$
$F \to u_x$	0.008380, 0.008236	0.1, 36.901	0	0.00203

[a]Implemented delay. Fitted delays are 0.40, 2.58, 5.20 and 5.15 min. [b]Physical output units. Exhauster paths add in parallel, fitted after the first minute. Other two-lag paths are cascades.

has tripped or that the bubblers have failed, which is deliberate. A feedforward that is right about everything is a simulation rather than a feedforward, and the interesting question is what the feedback has to do when it is wrong.

11.5.4 Model Predictive Control

The optional model predictive controller (MPC) plans electrode power and slurry feed together [18]. It follows the general form of Section 1.6.3, with the plant-specific model, observer and objective below. Its fixed local model is identified from symmetric ±25 kW and ±25 kg/h steps about 1260 kW and 1425 kg/h, following 30 min of settling and 360 min of observation. The outputs are the filtered glass thermocouple, inferred coverage and the pressure controller's exhauster command. Table 11.4 gives the implemented coefficients and independently reproduced fit residuals. The small positive feed-to-temperature gain includes the cap's insulating response. These local coefficients differ from the large-step response times in Section 11.4.

Commands and outputs are normalized as

$$v = \frac{u - u_0}{1000}, \qquad z = D_y^{-1}(y - y_0),$$
$$y_0 = (1150.00648, 0.9102041, 66.6326),$$

where $u = (P_{\mathrm{cmd}}, F_{\mathrm{cmd}})$ is the command pair, $u_0 = (1260, 1425)$ in kW and kg/h, and $y = (\hat{T}_g, \hat{\varphi}, u_x)$ holds the outputs, so the components of y_0 are in °C, coverage fraction and percent. The diagonal scaling D_y has the entries 50 °C, 0.1 and 10 %. Each lag uses $a = \exp(-1/\tau)$ and the sampled recursion $x^+ = ax + (1-a)Kq$, where q is its normalized input and K its scaled gain. The normalized path gain is $1000K/D_{y,ii}$. Cascades use the preceding sample of the upstream lag, with unit gain in the second lag. Two five-sample command histories supply the delays. Eight lag states, ten delay states and three random-walk output biases give 21 states in $x^+ = Ax + Bv$, $z = Cx$. Each output sums its corresponding paths and one bias. Short-delay rounding and cascade sampling add RMS errors of 0.00610 and 0.000245 °C on the two temperature paths, and $3.78\,10^{-5}$ and $5.24\,10^{-5}$ on coverage, relative to the continuous fits.

A steady-state Kalman correction estimates the output biases for offset-free regulation [19]:

$$\bar{x}_k = A\hat{x}_{k-1} + Bv_{k-1}, \qquad \hat{x}_k = \bar{x}_k + L(z_k - C\bar{x}_k). \tag{11.19}$$

The normalized process covariance has diagonal 10^{-8} on the 18 dynamic states and 10^{-4} on the three biases, with measurement covariance $0.0004I_3$. Initialization fills the histories from known commands and puts initial measurement mismatch into the biases. The last sampled applied inputs approximate the preceding minute as a zero-order hold. Envelope clipping or fallback can vary the actual inputs within that minute. This estimator does not recover true coverage from a biased inference.

The decision period is 1 min, while the existing 0.01 min loops retain their scan. A six-hour prediction has 360 samples with block lengths $[1, 1, 2, 4, 8, 16, 24, 36, 48, 60, 72, 88]$ samples. The 24

command variables and three nonnegative horizon-wide slacks minimize

$$J = \sum_{i=1}^{360} \left[(z_{T,i} - r_T)^2 + (z_{\varphi,i} - r_\varphi)^2 - 0.1 v_{F,i} \right] + 5 \sum_{j=1}^{12} \|\Delta v_j\|_2^2 + 10^5 \|\varepsilon\|_2^2, \tag{11.20}$$

where r_T, r_φ are the normalized current supervisory temperature and coverage targets, held over the horizon. All terms are dimensionless. The feed reward is a throughput surrogate, with no energy price or terminal penalty. The shared solver adds a numerical ridge to the QP Hessian. Hard command boxes are $0 \le P_{\text{cmd}} \le P_{\max}$ and $0 \le F_{\text{cmd}} \le \min(F_{\max}, F_{\text{op}})$, with changes at each block boundary limited to 60 kW and 60 kg/h. The first move is anchored to the applied input clipped into the current box. Soft prediction bounds are $T_{g,\min} - 50\varepsilon_T \le T_g \le T_{g,\max} + 50\varepsilon_T$, $\hat{\varphi} \le \varphi_{\max} + 0.1\varepsilon_\varphi$ and $u_x \le u_{x,\max} + 10\varepsilon_x$. Default limits are 1050 and 1210 °C, 0.97 coverage and 95% exhauster demand. Each slack is shared across all 360 samples, including both temperature bounds.

The first solve occurs at initialization. Accepted command pairs hold until the next decision. No future batch information enters either controller. The pressure PI, selected feedforward, level PI, timed pours and physical power envelope remain active. Replaced PIs track the applied commands after clipping. A nonoptimal or invalid plan returns authority to Coordinated until a valid plan is available. Validation checks shapes, finite costs and physical predictions, nonnegative slacks, first-move consistency and every planned command's box and rate limits. The nonlinear plant, its instrument chain and this nominal model remain separate. True foam, cap inventory and future disturbances are evaluation quantities unavailable to the predictive controller.

11.6 Results

The app offers twelve scenarios. They are 1, the glass is the resistor, 2, a power step test, 3, a pressure step test, 4, lined out at the design point, 5, independent loops meet a wetter batch, 6, coordinated control and its feedforward, 7, sulfate foaming, 8, a loaded scrubber, 9, a canister pour, 10, the exhauster trip, 11, MPC and Coordinated on a loaded scrubber, and the eight-hour challenge, listed last. Each fixes the starting mode, the architecture, the six sets of tuning constants, the disturbances and any timed events, and each states its objective before the run. A completion rule judges each one, as Section 1.6.4 describes, and some steps keep an observation across a Reset. A confirmation records an observation that the rule cannot inspect directly. The exercises of Section 11.11 are built on these cases.

11.6.1 The Design Point and What Lining Out Costs

The nominal design initialization uses 1425 kg/h of 55 % slurry and 1260 kW, which lines out at 1150.0 °C of glass and 436 °C of plenum, with the cold cap covering 90.8 % of the 10 m^2 surface at 21.5 mm mean thickness and making 500.2 kg/h of glass at 2.52 kW h/kg, of which 1.58 kW h/kg is spent boiling water that arrived with the waste. The electrode circuit sits at 0.350 Ω, 664 V and 1896 A against an envelope of 1605 kW, and the exhauster holds −250 Pa at 66.6 % demand against 2366 kg/h of total gas, of which the cold cap makes 925 kg/h, the film cooler adds 900 kg/h and the shell leaks 541 kg/h inward. That is 1.20 t/m^2 of glass per day, the number these machines are sold on.

Nothing is misbehaving at that point and a third of the output is missing. Sweeping the feed and power plane against every constraint puts the best feasible point at 662 kg/h of glass, reached at 1885 kg/h of slurry and 1631 kW with the pool at 1204 °C, the cap at 92.2 % coverage and the exhauster at 74.5 %. The binding constraint there is the refractory ceiling, which is why the guard band below it is an engineering decision about melter life rather than a safety factor.

Table 11.5 Measured Outcomes Over the Same 300 min Run With Feed Moisture Stepped to 68 % at $t = 60$ min and Back at $t = 240$ min. Pressure Error Is Accumulated Over $2 \le t \le 300$ min.

	Indep. no FF	Coord. no FF	Coord. FF $\to P$	Coord. FF $\to P, T$
Glass produced (kg)	1968	2444	2444	2454
Above coverage ceiling (min)	188.3	0.0	0.0	0.0
Worst pressure error (Pa)	73.1	90.0	54.2	53.8
Pressure IAE (Pa min)	61.3	136.4	75.0	77.5
Peak glass temperature (°C)	1150.4	1200.7	1200.7	1205.0
Production at 300 min (kg/h)	577	620	620	624
Feed rate at 300 min (kg/h)	1425	1850	1850	1850

11.6.2 The Wet Batch That Nothing Reports

Scenario 5 holds the design point under independent loops and takes the feed from 55 % water to 68 % at $t = 20$ min. Four things move and they move at four different speeds. Within one scan the steam rate goes from 784 kg/h to 969 kg/h, the plenum rises 66.0 Pa before the damper catches it, and the exhauster settles at 70.9 % against 66.6 %. Over the following two hours the crust grows from 90.8 % coverage at 21.5 mm to 99.6 % at 44.9 mm, well past the bridging limit. Production falls from 500.2 kg/h to 314.9 kg/h, a loss of 37 %. And the glass temperature does not move, reading 1150.0 °C before the batch, a maximum of 1150.4 °C during it and 1150.0 °C at the end, with the electrode power going slightly down from 1260 kW to 1242 kW.

> **TRY IT IN THE APP** **5 · Independent Loops Meet a Wetter Batch**
>
> Load scenario 5 and press Run at 10 or 20 times speed. The batch turns wetter early in the run. Watch the glass temperature: it reads 1150.0 °C before the batch and never more than 1150.4 °C during it. Then open the Production tab: glass production has fallen from 500.2 to 314.9 kg/h.

The thermal compensation explains the small temperature excursion. The extra evaporation duty is 186 kW and it does come out of the pool, but the crust that grew to absorb it took the exposed-melt radiation of (11.6) from 110 kW to 5 kW, so the pool lost 105 kW of its own heat leak at the same time. The cold cap absorbed the disturbance and the temperature loop reported nothing wrong. This modeled disturbance illustrates the reason for inferring coverage from the plenum thermocouples and putting it on the screen even though the inference is biased.

Pause. Try Exercise 11.2 before reading on.

11.6.3 Coordination Against Independent Loops

Table 11.5 runs the identical disturbance sequence under four configurations for 300 min. Feed moisture goes from 55 % to 68 % at $t = 60$ min and back at $t = 240$ min. The independent case runs at the design feed rate of 1425 kg/h, and the three coordinated cases are given 2200 kg/h as a ceiling and find their own rate underneath it.

Coordination is worth 2444 kg of glass against 1968 kg, 476 kg or an increase of 24.2 % from the same equipment on the same feed, and it spends 0 min above the cold-cap coverage ceiling where independent control spends 188.3 min. The margin depends on the feed rate chosen for the independent case, and 1425 kg/h is the rate at which this melter makes its nameplate 500 kg/h, so the comparison is against a competent choice. Neither result comes from tuning. The independent loops are correctly tuned and each one holds its own setpoint throughout. What changes is that the coverage controller sheds feed when the crust starts to grow instead of letting a fixed feed rate push it past the bridging limit, and that the supervisor walks the pool to 1198 °C on the dry batch where the operator's setpoint held it at 1150 °C. On a dry batch after twelve hours the coordinated melter

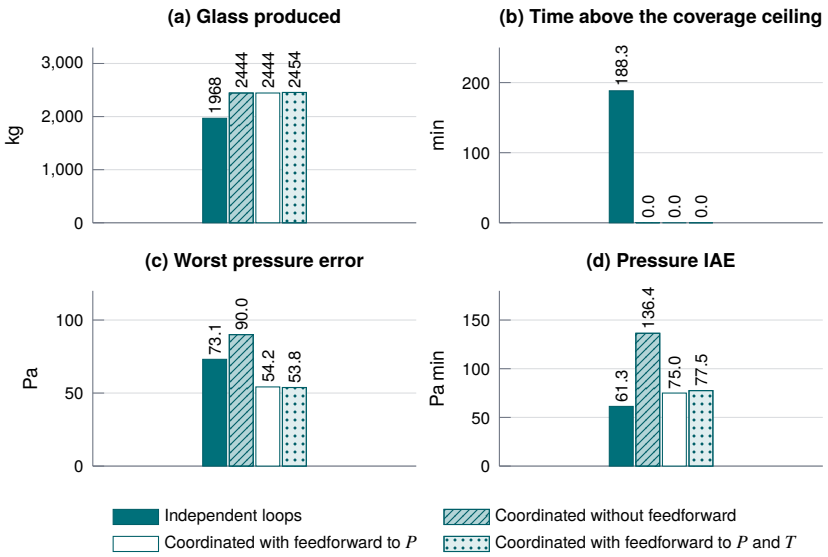

Figure 11.3 The four configurations of Table 11.5 on the same 300 min run with a wetter batch: (a) glass produced, (b) minutes above the cold-cap coverage ceiling, (c) the worst pressure error and (d) the accumulated absolute pressure error. Solid: independent loops. Hatched: coordinated without feedforward. White: coordinated with feedforward to the pressure loop. Dotted: feedforward to both loops. Notice that coordination buys the glass and the coverage, while the pressure feedforward changes only the two pressure panels.

settles at 646.6 kg/h of glass at 1842 kg/h of feed, which is 97.7 % of the best feasible 662 kg/h. The independent column of Table 11.5 reads 577 kg/h at $t = 300$ min rather than 500 kg/h because the oversized crust the wet batch left behind is still melting through faster than the nozzle is feeding it. Figure 11.3 draws four rows of Table 11.5 as bars.

The two middle columns isolate the feedforward. Both produce approximately the same glass, 2444 kg, with differences below one kilogram. The pressure feedforward has only a small indirect effect on the production loops in this comparison. What it changes is the pressure loop, where the worst excursion falls from 90.0 Pa to 54.2 Pa and the accumulated absolute pressure error falls from 136.4 Pa min to 75.0 Pa min. The independent case has the smallest integrated error of the four at 61.3 Pa min, and that is not a point in its favor, because it moved the feed rate exactly once and gave the pressure loop almost nothing to reject. Extending the feedforward to the temperature loop adds 10 kg of glass over five hours with only a small change to the pressure error, which is the expected result when a disturbance path is already two hours slow.

11.6.4 Constraints That Bind at the Actuator

Scenario 8 opens on an off-gas train that has run for a week. Pressure is on setpoint, temperature is on setpoint, production is 500 kg/h and nothing alarms. What changed is the exhauster demand, which reads 88.1 % against the clean train's 66.6 %, because the train now costs 6311 Pa of pressure drop for the same flow instead of 3504 Pa. Twenty-two points of actuator authority have gone and no measurement on the melter reported it. The wet batch that cost 66.0 Pa and took the damper to 70.9 % on a clean train now costs 72.3 Pa and takes it to 92.5 %, within two points of the 95 % ceiling the supervisor is allowed to use and close enough to the stop that a second surge would not be absorbed at all. Turning the feedforward on brings the same excursion back to 45.5 Pa.

A constraint that binds at an actuator has to be relieved at an actuator, and the only actuator that can relieve this one is the feed pump. Holding an actuator position below a ceiling with a slower

manipulated variable is also the idea behind the valve-position controller of Volume 4, Chapter 9. Scenario 10 is the extreme case. The main exhauster trips to a standby unit with 45 % of the head, and the plenum reaches +34.8 Pa and spends 33 s above the −25 Pa containment limit, which no control system could have prevented because the damper travels in two seconds and the head halved. What follows is the point. The exhauster demand crosses its ceiling, the low selector hands the feed rate to OVR-X, and the feed comes down from 1577 kg/h to a minimum of 1052 kg/h. The plenum recovers to −250 Pa without scan-to-scan command chatter. From minutes 100–150 the feed command changes by less than 0.002 kg/h per scan. The standby unit carries the reduced off-gas load, and production at minute 200 is 430 kg/h. The melter produces 1236 kg over the 160 min trip. Under independent loops the feed does not move, the damper sits at 100 %, and the crust bridges completely. That run produces 785 kg, 36.5 % less. Override control lets the plant give up production to relieve the actuator constraint. The coordinated melter still ends at 48.3 % coverage, below the imposed floor, and at a plenum temperature of 788 °C, above its 750 °C limit. The model does not quantify volatilization.

> **TRY IT IN THE APP** **10 · The Exhauster Trips and the Feed Pump Becomes the Pressure Controller**
>
> Load scenario 10 and press Run. The main exhauster trips to a standby unit with 45 % of the head. Watch the plenum pressure jump above its containment limit, then the binding line on the FRC card: the low selector hands the feed rate to OVR-X, the feed falls from 1577 to a minimum of 1052 kg/h, and the plenum recovers to −250 Pa.

11.6.5 Foam, the Pour and the Graded Shift

Scenario 7 sends a sulfate-bearing batch to the melter, which forms a raft between the melt and the cap and cuts the melt-to-cap coefficient of (11.3) by up to 55 %. With independent loops and the feed unchanged, production falls from 500 kg/h to 84 kg/h and the melter produces 542 kg in the 200 min the batch takes to clear, with the electrode power falling too, because a pool that cannot give heat away needs less of it. Two remedies are usually offered and only one works. Dropping the temperature setpoint to 1100 °C does reduce the foam index, but it also takes the driving difference $T_g - T_{cc}$ from 170 °C to 120 °C, and the second effect is larger, so the melter produces 135 kg instead of 542 kg. Cutting the feed target to 900 kg/h produces 1151 kg, and letting the coordinated supervisor find the rate for itself produces 1427 kg, which is 2.6 times what independent loops managed with nobody touching a setpoint.

Scenario 9 requests full airlift for 25 min at minute 20. By minute 50 the cumulative discharge is 1867 kg, including normal level-controlled discharge before and after the request. The level reaches −54.7 mm and resistance rises from 350.4 mΩ to 371.8 mΩ. Power regulation raises the voltage from about 664 V to 684 V. In Auto, voltage regulation also recalculates the tap from the temperature controller's power demand at each scan, so this experiment gives essentially the same delivered power in both modes. It does not hold a fixed voltage. A held tap would lose power as resistance rises, as the Manual experiment in Section 11.4.2 demonstrates. In the added Manual pour comparison, a 664 V tap gives about 1258 kW at minute 20 and 1183 kW at minute 45, while Manual Power holds 1260 kW. The fixed-tap trial also changes temperature, so the difference is not solely a level derivative.

The challenge runs an eight-hour shift with three batches, a partially loaded scrubber accumulating deposits three times as fast as usual, and a view port seal failure at $t = 380$ min. To pass, the melter must produce at least 4100 kg by minute 480 with no time above −25 Pa, less than 5 min more than 100 Pa off the pressure target, less than 10 min below 1050 °C and less than 90 min above the coverage ceiling, and the logged temperature error from $t = 0.66(480)$ min onward must remain within 12 °C. The scenario parameters and disturbances must stay unchanged. The rule judges cumulative glass formation rather than airlift discharge. As delivered, the melter produces 3.880 t and discharges 3.637 t, with 289.4 min above the coverage ceiling and 22.4 min more than 100 Pa off the pressure target.

Table 11.6 Fixed-Setting Comparison Over 480 min. E_T Is Temperature IAE (°C min), $e_{T,\max}$ the Peak Measured Error (°C), t_x Time at or Above 95% Exhauster Demand (min), D_x Its Excess Integral (% min), $x_{\max}$ Peak Demand (%), and t_c True Coverage Above 0.97 (min). n/f Counts Solves/Rejected Plans. Mean Times Are Headless Wall-Clock Observations.

Case	Law	Glass (kg)	Energy (kWh)	E_T	$e_{T,\max}$	t_x	D_x	$x_{\max}$	t_c	n/f	Mean (ms)
Loaded scrubber	C	4038.6	12480	1668.8	40.9	341.275	7.8324	96.01	0.000	0/0	—
	MPC	4013.4	11676	8644.9	40.9	325.735	23.2778	99.98	0.000	480/0	19.05
	C-SIMC	4045.8	12471	1315.3	40.9	358.525	8.8199	96.16	0.000	0/0	—
Wet batch	C	4373.8	12897	1918.3	48.0	0.000	0.0000	77.05	0.000	0/0	—
	MPC	4423.4	12936	2248.6	48.0	0.000	0.0000	78.12	0.000	480/0	18.05
	C-SIMC	4374.3	12883	1732.8	48.0	0.000	0.0000	77.02	0.000	0/0	—
Challenge	C	4413.6	12065	4028.9	48.0	79.125	11.9713	96.69	72.210	0/0	—
	MPC	4425.0	11660	5369.2	48.0	89.500	34.2925	100.00	0.000	480/0	18.34
	C-SIMC	4426.6	12061	2090.4	48.0	70.730	9.9220	96.37	15.120	0/0	—
Sulfate	C	3824.5	10949	2943.7	48.0	0.000	0.0000	68.55	0.000	0/0	—
	MPC	3716.3	9645	14233.0	50.3	0.000	0.0000	68.56	19.185	480/0	18.54
	C-SIMC	3811.3	10957	1437.7	48.0	0.000	0.0000	68.65	0.000	0/0	—
Blower trip	C	4446.4	12405	2527.9	48.0	109.560	10.5789	97.05	0.000	0/0	—
	MPC	4323.8	11391	14794.7	67.5	152.635	76.4851	100.00	0.000	480/0	19.02
	C-SIMC	4450.4	12421	1780.6	48.0	115.710	12.2878	97.30	0.000	0/0	—

Applying suggested settings before Run or applying them and then resetting gives the same trajectory: 4.414 t produced and 4.418 t discharged, with 72.2 min above the coverage ceiling and 0.41 min off pressure. Production rises by 533 kg, or 13.7 %. These settings use power regulation, coordinated control, both feedforward paths, a feed ceiling of 2100 kg/h and a coverage target of 88 %. Both startup paths pass the completion rule. Its temperature test uses the largest logged error from minute 316.8 onward, the final 34 percent of the shift.

11.6.6 Predictive and Selector Control Compared

The added loaded-scrubber case starts with $R_\ell = 6500$ and receives an unannounced moisture change from 55 to 68% at minute 60, returning at minute 240. All five comparisons run 480 min with seed 20260118 and identical instruments. Neither structure receives a batch preview. Coordinated (C) retains shipped tuning, except for the unchanged suggested challenge settings. MPC retains the six-hour settings of (11.20) in every case. A third law, C-SIMC, uses one fixed retuning of the temperature and coverage PIs from the same nominal diagonal paths [20]. The half-rule reduction gives $\tau = \tau_1 + \tau_2/2$ and $\theta = \theta_{\text{fit}} + \tau_2/2$ for coverage. The tuning rule is $K_c = \tau/[K(\lambda+\theta)]$, $\tau_I = \min[\tau, 4(\lambda+\theta)]$. Selecting the highest loaded-scrubber glass total among $\lambda = 5, 10, 20, 40, 80$ min gives $\lambda = 5$ min, with $(K_{c,T}, \tau_{I,T}) = (99.1327, 21.5994)$ and $(K_{c,\varphi}, \tau_{I,\varphi}) = (3528.8971, 45.8335)$ in the units of Table 11.3. This selection is limited to that tested tuning family.

On the loaded train, MPC produces 4013.4 kg against C's 4038.6 kg, while using 11676 against 12480 kWh. Its temperature IAE is 8644.9 against 1668.8 °C min. Exhauster demand is at or above 95% for 325.735 against 341.275 min, but the excess integrals are 23.2778 and 7.8324 % min. The longer C exposure is much shallower. C-SIMC gives 4045.8 kg and an IAE of 1315.3 °C min, with 358.525 min and 8.8199 % min of exhauster excess. During minutes 60–240, C and MPC produce 1144.4 and 1124.6 kg. During minutes 240–480 they produce 2329.6 and 2345.2 kg. Duration, depth, energy and production describe different losses.

MPC also produces less glass on sulfate and blower, at 3716.3 and 4323.8 kg against C at 3824.5 and 4446.4 kg. All 15 runs have zero time below the glass floor, above the glass ceiling or outside the level band. These observations do not establish physical safety. On blower, C's true coverage stays below 0.55 for 142.190 min and its plenum exceeds 780 °C for 137.195 min. MPC avoids both, but containment loss lasts 0.580 against 0.555 min. Containment excess integrals are 15.4941 and 17.0726 Pa min, and the peaks are 59.7903 and 65.6884 Pa beyond the −25 Pa limit. Pressure-error exposure is 1.960 and 4.285 min. On sulfate, MPC instead spends 19.185 min above true coverage

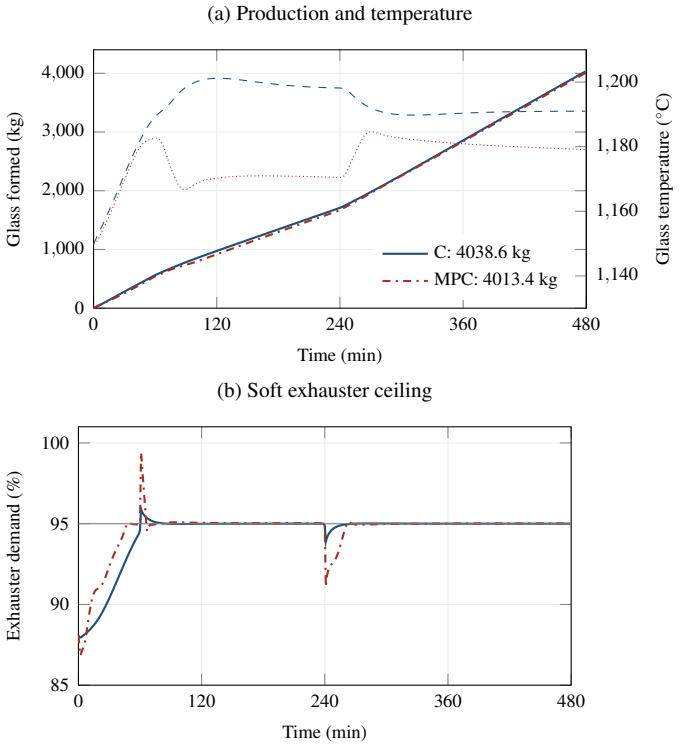

Figure 11.4 Loaded-scrubber comparison with identical axes for both control laws. Coordinated is drawn solid and MPC dash-dotted. Panel (a) plots accumulated glass as thick lines on the left scale and glass temperature as thin lines on the right scale, dashed for Coordinated and dotted for MPC. Panel (b) plots exhauster demand and its 95% ceiling as a thin gray line. The wet batch occupies minutes 60–240. MPC uses less energy but produces slightly less glass, tracks temperature less closely and incurs a larger exhauster-excess integral.

0.97. Time, depth and peak were recorded for every evaluated limit, from physical truth used solely for evaluation. Requested power is clipped for 27.440 and 23.770 min on the loaded train, and 27.440 and 79.780 min on blower. This is unavailable requested power, not a violation of delivered-power capacity.

> ✦ DID YOU KNOW? **The Predictive Controller Wins One Case of Five**
>
> From the simulator: of the five cases of Table 11.6, MPC makes the most glass only on the wet batch, 4423.4 kg against 4374.3 for C-SIMC and 4373.8 for Coordinated. On the loaded scrubber, the challenge, sulfate and the blower trip, a selector law makes more glass than MPC.

Suggested C passes the unchanged challenge with 7.7 °C worst error in its scored interval. MPC produces 4425.0 against 4413.6 kg, but fails at 28.4 °C against the 12 °C band. C-SIMC passes at 2.2 °C. Requested-power travel is 3500 kW for C-SIMC, 2832 for C and 3041 for MPC. The local model, soft output bounds and throughput reward leave losses that output-bias correction cannot remove. These runs do not isolate their individual causes.

Each MPC run accepts 480 plans without fallback. A separate browser measurement gave a warmed mean of 22.10 ms over 219 solves, p95 22.5 ms and maximum 29.0 ms. At the application's 1× rate of one simulated minute per wall second, this is 2.210% of one core. Requested 20× playback

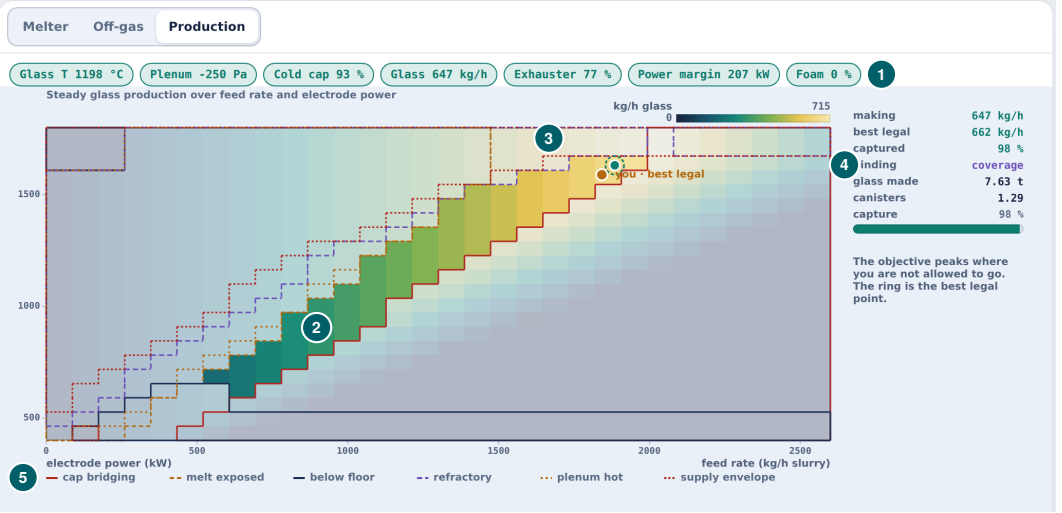

Figure 11.5 Coordinated control at $t = 12.00$ h on the Production view, with the feed ceiling at 2200 kg/h. The operating map plots steady glass production over feed rate and electrode power, with each constraint traced as its own outline, infeasible cells dimmed rather than hidden, the live point trailing behind itself and the best legal point ringed. The melter is making 647 kg/h against a best feasible 662 kg/h, which is 98 % captured, with the coverage controller holding the selection, the pool at 1198 °C, the plenum on setpoint at −250 Pa and the exhauster at 76.6 % of which 10.3 % is feedforward. The view is cut to the Production panel of the application. Callouts: (1) the status line, (2) the operating map with each constraint outline, (3) the live point and the ringed best legal point, (4) the production readings with the binding constraint, (5) the key to the constraint outlines.

measured 18.147 simulated min/s and 49.366 frames/s with zero skipped reached decisions. A 25 ms frame-work guard bounds pending frame work while retaining simulation steps and decisions. A short guard-off observation of 18.238 min/s and 50.332 frames/s does not establish a frame-rate benefit. Trajectory reproducibility excludes wall-clock solver timing.

Figure 11.5 shows the separate dry coordinated run at 720 min, with the objective on the map peaking where the melter is not allowed to go and the best legal point sitting on a corner where two constraints meet.

11.7 Technology in Practice

The tags of Figure 11.1 stand, in a plant, for thermocouples in thermowells, purged dip tubes, transformers on the electrodes, a slurry pump, an exhauster and an airlift. Figure 11.6 draws the loops with that hardware in place of the schematic's symbols.

What Measures It

Glass and vapor-space temperatures are read by thermocouples in thermowells lowered through the top head of the melter, beside the level probe and the feed tube [17]. A thermocouple is a junction of two dissimilar metals whose voltage is roughly proportional to temperature, and the well around it adds a lag [21]. The chapter's glass thermocouple TI-101 is that lag, 0.60 min. The plenum thermocouples carry a heavier job. With no instrument for the cold cap, they are the reading from which coverage is inferred in Equation (11.13), and operating melters set their feed by them [3]. Pressure and level come from purged dip tubes.

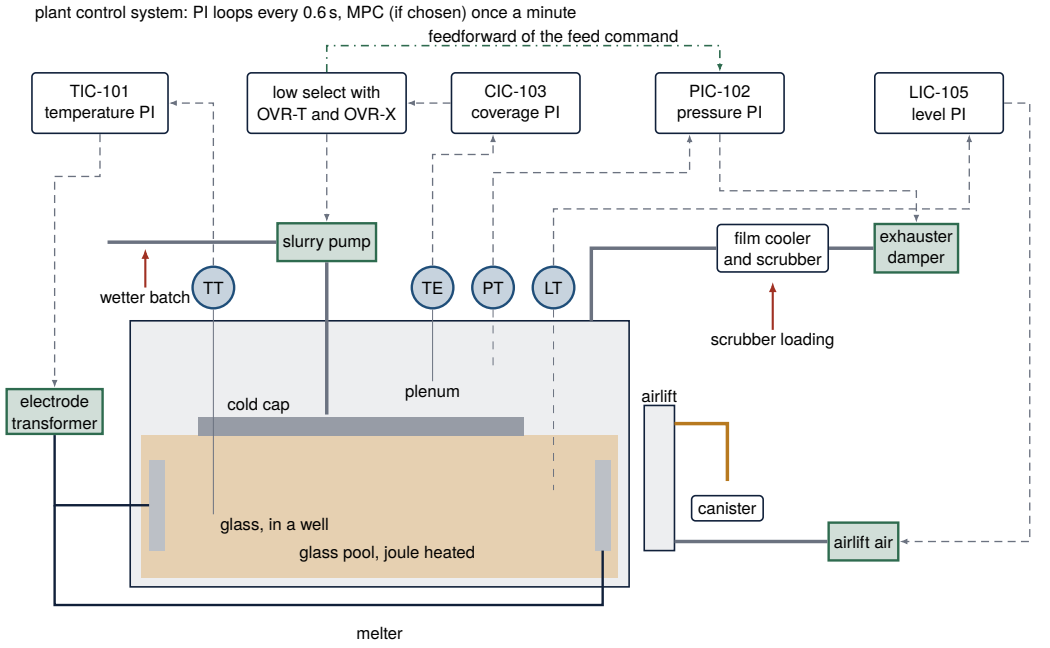

Figure 11.6 The loops of Figure 11.1 with the hardware of a plant. A thermocouple in a well (TT) feeds TIC-101, which sets the power of the electrode transformer. A plenum thermocouple (TE) gives the coverage that CIC-103 holds with the slurry pump, through a low select shared with the overrides OVR-T and OVR-X, where external reset keeps the losing controllers from winding up. Purged dip tubes read the plenum pressure (PT) for PIC-102, which moves the exhauster and receives the feed command as feedforward (dash-dot), and the glass level (LT) for LIC-105, which sets the airlift air. Thick gray lines carry slurry and off-gas, thin dashed arrows signals.

> **TECHNOLOGY** **Sensor: Purged Dip Tubes for Pressure and Level**
>
> A dip tube is a pipe whose open end sits in the process, fed with a purge gas so slowly that only a few bubbles escape. The gas pressure then equals the pressure at the tube's end, and a transmitter anywhere along the line reads it without touching a corrosive or plugging fluid. If the purge stops, the reading is no longer accurate and the tube may plug, too much purge reads high from friction, and each bubble that breaks away leaves a small ripple on the trend [21]. The Savannah River melter reads plenum pressure and glass level from a short and a long pipe of one argon-purged probe, and bubbling moved its indicated level by less than 0.25 cm [17]. The chapter's 4.2 s tap delay stands for this line.

What Moves It

Four final elements carry the six loops. Electrode power comes from a transformer on each electrode set. Slurry enters through feed tubes, and one low-activity design spreads it through six nozzles to give a cap of uniform thickness. Glass leaves by an airlift: air bubbled into the bottom of a riser lifts the glass over a weir into a trough and the canister [16]. Air or argon bubbled into the pool itself raises the heat transfer to the cap [16, 17], which is the multiplier u_b of Equation (11.3). The exhauster at the end of the train holds the plenum below atmospheric pressure, and the chapter models it as a fan whose head grows with the square of its demand u_x.

TECHNOLOGY Final Element: The Electrode Power Supply

Single-phase current reaches the glass through opposing pairs of nickel-alloy plate electrodes, and the glass pool temperature sets the electrode power through a feedback loop [16]. Each electrode set has its own transformer, and the hardware fixes its limits. At Savannah River a change of tap switch position raised the lower transformer's maximum current from 1920 to 2560 A, and the upper controller's current limit went from 71 to 90 % of its 1920 A output, after the allowed current density was doubled. A melter needs both sets: it cannot run on one [15]. The chapter lumps them into one pair with the voltage, current and power envelope of Equation (11.11).

Where the Controller Runs

The loops run in a plant control system, whose loops cycle in times from milliseconds to minutes [22]. The chapter's PI loops execute every 0.01 min (0.6 s) and its predictive controller decides once a minute, both inside that range. The coordinated structure is the override pattern of process control. Several controller outputs meet at a low-select function, only the lowest moves the feed pump, and each de-selected controller must be kept from winding up while it waits [21]. The external reset of Equation (11.17) is that protection, built from one PI block per loop and one selector block.

IN PRACTICE Feeding by the Plenum Temperature

Cameras installed to watch the cold cap at two melters stopped working after very limited time inside, so both plants learned to set the feed rate from the plenum temperature instead. One targeted a plenum between 475 and 525 °C, and the other keeps it at or above 493 °C to control flammable gas. Feeding too fast bridges the crust to the side wall, builds up feed in places, and brings large pressure swings and unstable temperatures [3]. The plenum itself is held slightly below atmospheric pressure to capture the off-gas and keep radiological control [16].

What Is Modernizing It

Two developments address the readings the chapter has to infer. Millimeter-wave sensors view the glass through refractory waveguides from a receiver that can sit outside the biological shield. Their thermal emission signal is linear with temperature, and tests have reached 1500 °C [23]. Computational fluid dynamics models of a pilot melter reproduce its glass and plenum thermocouples on average within 8 and 91 °C, and such models may one day inform the melter's feedback loops [3]. The nonlinear predictive control and moving horizon estimation that Section 11.9 proposes would use the chapter's own equations in the same way.

11.8 Limitations

The model is not a melter design tool. With one lumped cold cap, a single-pair electrode circuit and a one-resistance off-gas train it cannot support a claim about a real machine's throughput, refractory life, off-gas loading or sulfate limit, and the constants that set those answers are calibrations chosen to put the design point at 500 kg/h rather than measurements. Volatilization is absent, so the coverage floor is imposed rather than derived. Glass chemistry does not appear at all, which puts waste loading, durability and the composition constraints that actually size a melter out of reach. The pour is a timed transfer from a well-mixed inventory, with no riser freeze, no canister changeout and no pour spout heating, which are the things that stop a real melter from pouring. Foaming is one first-order state driven by a temperature sigmoid, whereas the real phenomenon is hysteretic and not reliably predictable, which is why sulfate limits on melter feed are set conservatively in practice [13]. Repeatability is itself unrealistic, because freezing the integration step and the controller period reproduces the physical trajectory, excluding wall-clock solver timings, while a distributed control system does no such thing. What survived the simplification is what a controller has to be designed around, namely the self-insulating layer, the four time scales, the positive electrical feedback, the

shrinking power envelope, the biased soft sensor and the actuator constraint, and Section 11.3.9 says where each assumption breaks.

> **OPEN CHALLENGE** **Seeing Inside a Hot, Radioactive Melter**
>
> Waste glass melters lack sophisticated diagnostics because they are hot, corrosive and radioactive. The temperature profiles of the glass, the cold cap and the plenum gas, and the conductivity, density and viscosity of the melt are all wanted, and with them in real time the process could run under feedback control [23]. Readings that disagree with conditions inside are a known operating problem, and the temperature profile matters more as melters grow, because local hot spots have breached refractory [16]. The chapter's one glass temperature and one inferred coverage are that gap in miniature. Progress would be sensors that survive inside the melter and report its profiles.

The predictive model is fixed at the nominal design point. It omits foam, changes in the electrical envelope and future batch composition. Its output-bias estimator corrects measured offsets without identifying those mechanisms or removing the coverage instrument bias. Soft constraints and finite-horizon optimization provide no guarantee of true plant limits or closed-loop stability. The sampled applied-input approximation also omits within-minute variation during clipping and fallback. Halving the physics step from 0.005 to 0.0025 min changes glass by less than 0.005 kg and temperature IAE by less than 0.08 °C min in the loaded-scrubber and suggested-challenge comparisons. Controller scans remain 0.01 min and MPC decisions remain one minute apart.

11.9 Discussion

The melter is a plant where the objective and the constraints live on different actuators. Production comes from the feed pump, the temperature constraint is relieved at the feed pump, the containment constraint is relieved at the damper until the damper runs out and then at the feed pump as well, and the electrode power sets none of it directly. That structure is what the operating map of Figure 11.5 exists to make visible, and reading such a map is a skill that transfers to any plant run against several limits at once.

> **OPEN CHALLENGE** **Coverage From a Thermocouple**
>
> A full-scale melter will rarely let the cold cap be watched, so its feed rate will rest on plenum thermocouples. Simulations of a pilot melter put a swing of up to 606 °C on the plenum from cold-cap coverage and up to 179 °C from the emissivity of the crust alone, and test feeds often build uneven caps that leave plenum temperatures differing from place to place [3]. The soft sensor of Equation (11.13) reads one straight line between two calibration points. Progress would be a reading that tells a lopsided or darker crust from a smaller one, quickly enough for operators to act [3].

The economics change what a controller should optimize. A melter campaign is measured in years, the equipment cannot be inspected or repaired, and the number that decides whether a cleanup mission finishes on schedule is tonnes of glass per day. Every degree of setpoint the supervisor takes toward the refractory limit is throughput bought with melter life, which is why the guard band is adjustable rather than buried. The difference between the operator's 1150 °C setpoint and the supervisor's 1198 °C is 146 kg/h of glass, and the matching cost in refractory corrosion is outside this model entirely.

The model's equations also support methods beyond the implemented PI hierarchy and fixed linear MPC. An eleven-state model with a nonconvex feasible region and an actuator envelope that moves with a state is a natural nonlinear model predictive control problem, and the selector of (11.16) is exactly the structure a constrained optimizer replaces [24]. Moving horizon estimation on the same equations would estimate the cold-cap mass from the plenum thermocouple and the off-gas flow together, which could help distinguish a changing cap from the inleakage disturbance that

biases (11.13) in the final challenge. The model transfers to the Gekko Optimization Suite in Python with the five hidden constants of Table 11.2 recast as estimated parameters [25].

11.10 Conclusion

This chapter studied a joule-heated waste vitrification melter as an interactive laboratory for a plant whose dominant resistance is a state the control action creates. The simulator integrates eleven states whose four principal dynamics span a factor of 3.1×10^5 between the computed gas-inventory time constant of 1.67 s and 6.1 days of scrubber loading. The pressure reaches 63 % of a damper step in 4.8 s, with a cold cap between 11.0 min and 54.6 min depending on coverage and a glass pool at 115 min. Two measured gains decide the electrode regulation, 7.52 kW/K of positive electrical feedback at a fixed tap against 6.84 kW/K of thermal damping, so a 9 V tap change walks the pool from 1150 °C to 1020 °C over twenty-five hours while the same 34 kW removed on power regulation settles at 1145.3 °C. A 140 kW power step raises the pool 25.9 °C, thins the crust from 90.8 % to 71.7 % coverage and leaves production at 500 kg/h, and a batch thirteen points wetter costs 37 % of production while the glass temperature moves 0.4 °C. Making the feed rate a manipulated variable under a coverage controller with temperature-floor and off-gas overrides on a low selector produces 2444 kg of glass in 300 min against 1968 kg for independent loops, with 0 min above the coverage ceiling against 188.3 min, and the feedforward inversion from feed rate to exhauster demand takes the worst pressure excursion from 90.0 Pa to 54.2 Pa with less than a kilogram difference in production. The additional MPC uses less loaded-scrubber energy, 11676 against 12480 kWh, but produces slightly less glass, 4013.4 against 4038.6 kg, with a higher temperature IAE, 8644.9 against 1668.8 °C min, and a larger exhauster-excess integral. No controller wins every constraint. A fuller model would add a layered cold cap, a multi-pair electrode circuit with its own current balance, a volatilization mass balance behind the coverage floor, and a condensing off-gas train. The melter remains a plant for learning to coordinate loops that share one actuator and disagree about time.

KEY POINTS

- A melter makes its glass inside a floating crust, and the same crust insulates the pool and loads the off-gas train.
- Feed rate sets production. Power, temperature and coverage decide only how much feed the melter can get away with.
- On a fixed voltage tap the electrical feedback of 7.52 kW/K exceeds the thermal damping of 6.84 kW/K, so the supply regulates power.
- A coverage controller with two overrides on a low selector makes 2444 kg of glass in 300 min against 1968 kg for independent loops.

Where else this applies. Overrides on a low selector also set the feed of Volume 4, Chapter 4, and a production handle pushed toward the first binding limit drives Volume 4, Chapter 6.

References

[1] R. Pokorny and P. Hrma, “Model for the conversion of nuclear waste melter feed to glass,” *Journal of Nuclear Materials*, vol. 445, no. 1-3, pp. 190–199, 2014.

[2] P. Ferkl, P. Hrma, J. Kloužek, A. A. Kruger, and R. Pokorný, “Cold-cap structure in a slurry-fed electric melter,” *International Journal of Applied Glass Science*, vol. 15, no. 1, pp. 73–87, 2024.

[3] A. W. Abboud, D. P. Guillen, and R. Pokorny, “Effect of cold cap coverage and emissivity on the plenum temperature in a pilot-scale waste vitrification melter,” *International Journal of Applied Glass Science*, vol. 11, no. 2, pp. 357–368, 2020.

[4] P. Kokotović, H. K. Khalil, and J. O’Reilly, *Singular Perturbation Methods in Control: Analysis and Design*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999.

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

[6] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[7] C. M. Jantzen, "Historical development of glass and ceramic waste forms for high level radioactive wastes," in *Handbook of Advanced Radioactive Waste Conditioning Technologies*. Woodhead Publishing, 2011, pp. 159–172.

[8] I. W. Donald, B. L. Metcalfe, and R. N. J. Taylor, "The immobilization of high level radioactive wastes using ceramics and glasses," *Journal of Materials Science*, vol. 32, no. 22, pp. 5851–5887, 1997.

[9] M. I. Ojovan and W. E. Lee, "Glassy wasteforms for nuclear waste immobilization," *Metallurgical and Materials Transactions A*, vol. 42, no. 4, pp. 837–851, 2011.

[10] 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.

[11] S. Lee, D. A. Cutforth, D. Mar, J. Klouzek, P. Ferkl, D. R. Dixon, R. Pokorny, M. A. Hall, W. C. Eaton, P. Hrma, and A. A. Kruger, "Melting rate correlation with batch properties and melter operating conditions during conversion of nuclear waste melter feeds to glasses," *International Journal of Applied Glass Science*, vol. 12, no. 3, pp. 398–414, 2021.

[12] D. R. Dixon, M. J. Schweiger, B. J. Riley, R. Pokorny, and P. Hrma, "Temperature distribution within a cold cap during nuclear waste vitrification," *Environmental Science & Technology*, vol. 49, no. 14, pp. 8856–8863, 2015.

[13] S. Lee, P. Hrma, R. Pokorny, J. Klouzek, B. J. VanderVeer, D. R. Dixon, S. A. Luksic, C. P. Rodriguez, J. Chun, M. J. Schweiger, and A. A. Kruger, "Effect of melter feed foaming on heat flux to the cold cap," *Journal of Nuclear Materials*, vol. 496, pp. 54–65, 2017.

[14] P. Ferkl, X. Lu, A. A. Kruger, and J. D. Vienna, "Temperature and composition dependence modeling of viscosity and electrical conductivity of low-activity waste glass melts," *Journal of Non-Crystalline Solids*, vol. 640, p. 123119, 2024.

[15] 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

[16] S. E. Kelly, "A joule-heated melter technology for the treatment and immobilization of low-activity waste," Washington River Protection Solutions, for the U.S. Department of Energy Office of River Protection, Richland, WA, Tech. Rep. RPP-48935, Rev. 0, Mar. 2011.

[17] M. E. Smith and D. C. Iverson, "Installation of bubblers in the Savannah River site defense waste processing facility melter," in *WM2011 Conference*, Phoenix, AZ, 2011, report SRR-STI-2010-00784, OSTI 1001181. [Online]. Available: https://www.osti.gov/biblio/1001181

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

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

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

[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.

[23] P. P. Woskov, J. S. Machuzak, P. Thomas, S. K. Sundaram, and W. E. Daniel, Jr., "Millimeter-wave monitoring of nuclear waste glass melts: An overview," Massachusetts Institute of Technology, Plasma Science and Fusion Center, Cambridge, MA, Tech. Rep. EMSP-81897-2002, Jun. 2002.

[24] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[25] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, "GEKKO optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.

11.11 Exercises

11.1 Concept check · 2 min

The melter runs in Manual at the design point with the feed fixed at 1425 kg/h of slurry at 55 % water. The electrode power is stepped from 1260 to 1400 kW and held until the pool and the cold cap have settled. What does the model predict for the glass production rate?

(a) It returns to about 500 kg/h, the dry-solids feed times its oxide yield

(b) It rises by about 11 %, in proportion to the extra electrode power

(c) It rises, because the melt-to-cap heat flow grows with $T_g - T_{cc}$

(d) It falls, because the thinner cap lets more heat escape as radiation

11.2 Concept check · 2 min

In scenario 5 a batch thirteen points wetter reaches the melter under independent loops at the design feed rate. Production falls by 37 %, yet the glass temperature never moves more than 0.4 °C from 1150.0 °C. Why does the temperature loop see almost nothing?

(a) The temperature loop is tuned fast enough to reject the extra evaporation duty within minutes

(b) The thicker crust cuts the radiation from the bare melt by about as much as the extra duty takes from the pool

(c) The extra water boils in the plenum, so its evaporation never draws heat from the pool

(d) The thermocouple lag of 0.6 min filters the temperature change out of the measurement

11.3 Concept check · 2 min

The melter is switched from power regulation to a fixed voltage tap at its design point, and the tap is then lowered from 664 to 655 V. What happens to the pool, according to the model?

(a) It settles a few degrees cooler, as it does after the same cut on power regulation

(b) It oscillates with growing amplitude, because the supply lag adds phase to the thermal loop

(c) It keeps cooling for many hours, because the power at a fixed tap falls faster with temperature than the losses do

(d) It warms, because a lower voltage drives a larger current through the glass

11.4 Theory and calculation · 6 min

Slurry enters at 1425 kg/h. The feed water needs $h_w = 3620\,\mathrm{kJ/kg}$, the dry solids need $h_d = 1770\,\mathrm{kJ/kg}$, and the dry solids have an oxide yield $y_{ox} = 0.78$. Use Equation (11.5).

(a) Compute the enthalpy per kilogram of slurry, Δh_{feed}, and the glass per kilogram of slurry, Y, at $w = 0.55$ and at $w = 0.68$. *Checkpoint:* 2788 kJ/kg and 0.351; 3028 kJ/kg and 0.250.

(b) Compute the steady glass production at this feed rate for each batch, and the fractional loss. *Checkpoint:* 500.2 kg/h and 355.7 kg/h, a loss of 28.9 %.

(c) Compute the extra evaporation duty of the wetter batch in kilowatts. Scenario 5 reads 314.9 kg/h just before the batch clears, a 37 % loss. Why is the run below the steady value of part (b)? *Checkpoint:* 186 kW.

11.5 Theory and calculation · 7 min

The glass resistivity is $\rho_0\,\mathrm{e}^{E_a/R_u(T_g+T_0)}$ with $E_a/R_u = 12085\,\mathrm{K}$. At the design point $T_g = 1150\,^\circ\mathrm{C}$, $P_{el} = 1260\,\mathrm{kW}$ and $R = 0.3504\,\Omega$. The pool holds $m_g c_p = 31.5\,\mathrm{MJ/K}$ and sheds 6.84 kW/K about the design point. The supply limits are 750 V, 2600 A and 1800 kW.

(a) Compute the ratio of the resistance at 1100 °C to that at 1200 °C, at a fixed glass depth. *Checkpoint:* 1.82.

(b) At a fixed tap, $P = V^2/R$. Show that $\mathrm{d}P/\mathrm{d}T_g = P\,(E_a/R_u)/(T_g + T_0)^2$ and evaluate it. Subtract the thermal damping and estimate the time scale $m_g c_p/(\text{net gain})$ on which the pool walks away. *Checkpoint:* 7.52 kW/K; net 0.68 kW/K; about 13 h.

(c) Compute the power envelope of Equation (11.11) at 1150 °C and at 1100 °C, and name the limit that binds. *Checkpoint:* 1605 kW and 1178 kW, both on the voltage limit.

11.6 Theory and calculation · 6 min

The plenum, cooler, duct and scrubber hold $V_{\mathrm{sys}} = 150\,\mathrm{m}^3$ of off-gas with $R_g = 350\,\mathrm{J/(kg\,K)}$ at $T_d = 380\,^\circ\mathrm{C}$. At the design vacuum of −250 Pa the shell inleakage follows Equation (11.9) with $\lambda = 1$ and $C_{\mathrm{leak}} = 0.0095\,\mathrm{kg/(s\,Pa^{0.5})}$, and the exhauster pulls $\dot{m}_{out} = 0.6571\,\mathrm{kg/s}$ through the clean train, $K_t = 8115\,\mathrm{Pa\,s^2/kg^2}$.

(a) Compute the coefficient $R_g(T_d + T_0)/V_{\mathrm{sys}}$, the inleakage in kg/h and its slope $\partial\dot{m}_{leak}/\partial P$. *Checkpoint:* 1524 Pa/kg; 541 kg/h; -3.00×10^{-4} kg/(s Pa).

(b) From Equation (11.8), compute the slope $\partial\dot{m}_{out}/\partial P$ at a fixed damper. *Checkpoint:* 9.38×10^{-5} in kg/(s Pa).

(c) Compute the time constant of the gas inventory, and the share of the self-regulation that the inleakage supplies. Compare the time constant with the 4.2 s transport delay of the pressure tap. *Checkpoint:* 1.66 s; 76 %.

11.7 App challenge · 7 min

Load scenario 1, The glass is the resistor (MANUAL). Before Run, switch to Advanced, set Electrode regulation to Voltage and set Electrode voltage tap to 655 V. Run to $t = 600$ min at 20× (30 seconds) at the default seed and export the CSV file. Its time column is in hours for a run this long.

(a) Read `glass_T_C` at 5 and 10 h, and `power_kW` at 10 h. *Checkpoint:* 1129.3 °C and 1104.1 °C, each within 0.3 °C; 869 kW, within 3 kW.

(b) Reload the scenario, keep Electrode regulation on Power, set Electrode power to 1226 kW before Run and run the same 600 min. Read the glass temperature at the end. *Checkpoint:* 1145.3 °C, within 0.1 °C.

(c) Both runs remove about 34 kW at the start. Explain with the two gains of Section 11.4.2 why one settles and the other does not.

11.8 App challenge · 8 min

Load scenario 6, Coordinated control, and why it needs feedforward, as shipped: Coordinated, Feedforward from feed on To pressure and a feed target of 2200 kg/h. Each run goes to $t = 300$ min at 20× (15 seconds) at the default seed. Read the Glass made tile and the IAE of the melter plenum pressure loop on the scorecard.

(a) Run the case as shipped. *Checkpoint:* 2.44 t; IAE 75.9 Pa min, within 0.5 Pa min.

(b) Pause, press Reset, switch to Advanced and set Feedforward from feed to Off before Run. Run the same 300 min and give the ratio of the two pressure IAE values. *Checkpoint:* 2.44 t; IAE 137.5 Pa min, within 0.5 Pa min; ratio 0.55.

(c) Why does the feedforward change the pressure loop and not the glass? Table 11.5 gives 75.0 and 136.4 Pa min. Why do your readings differ slightly?

11.9 App challenge · 8 min

Load scenario 10, The exhauster trips and the feed pump becomes the pressure controller, as shipped: Coordinated, To pressure feedforward and a feed target of 1800 kg/h. The standby exhauster takes over at $t = 40$ min and the main one returns at 200 min. Run to $t = 200$ min at 20× (10 seconds) at the default seed and export the CSV file. Its time column is in hours.

(a) Read the Containment tile and the peak of `plenum_P_Pa`. *Checkpoint:* 0.6 min; about +35 Pa, within 1 Pa.

(b) From `metric_glass_kg` at the last rows at or before 40 min and 200 min, compute the glass made during the trip, and read the minimum of `feed_kgh` after the trip. *Checkpoint:* 1236 kg, within 5 kg; 1052 kg/h, within 5 kg/h.

(c) Pause, press Reset, set Control architecture to Independent loops before Run and repeat. Give the glass made during the trip and the feed rate. *Checkpoint:* 785 kg, within 5 kg; the feed stays at 1800 kg/h.

11.10 Application area · 3 min

At the end of the exhauster trip of scenario 10 the coordinated melter has held containment and kept producing, but its cold cap covers only 48.3 % of the melt and its plenum is at 788 °C.

(a) What does an exposed melt do on a real machine that this model does not show, and why does the coverage floor exist?

(b) The feed cannot rise, because the exhauster is on its ceiling. Which actuator could an added override use to protect the cap, and what would it cost?

11.11 Application area · 4 min

The supervisor of the coordinated architecture walks the temperature setpoint from the operator's 1150 °C to 1198 °C, 12 °C below the refractory ceiling of 1210 °C, and the chapter puts the difference at 146 kg/h of glass.

(a) Express that difference in tonnes of glass per day and as a share of the 500 kg/h design production. *Checkpoint:* 3.5 t per day; 29 %.

(b) Why is the guard band adjustable rather than fixed in the code, and what must an engineer bring from outside the model to set it?

11.12 Review the agent · 5 min

An AI assistant reviews scenario 11 and writes: "On the loaded scrubber the MPC uses 11 676 kW h against 12 480 kW h for Coordinated, about 6.4 % less. Its exhauster demand is at or above the ceiling for 325.7 min rather than 341.3 min. The predictive controller is better and should replace the selector for all operation of this melter." Find the flaw and explain how to check it in the app.

CHAPTER 12

Tokamak Vertical Stabilization Simulator for Gain Limits and the Recoverable Region of an Unstable Plant

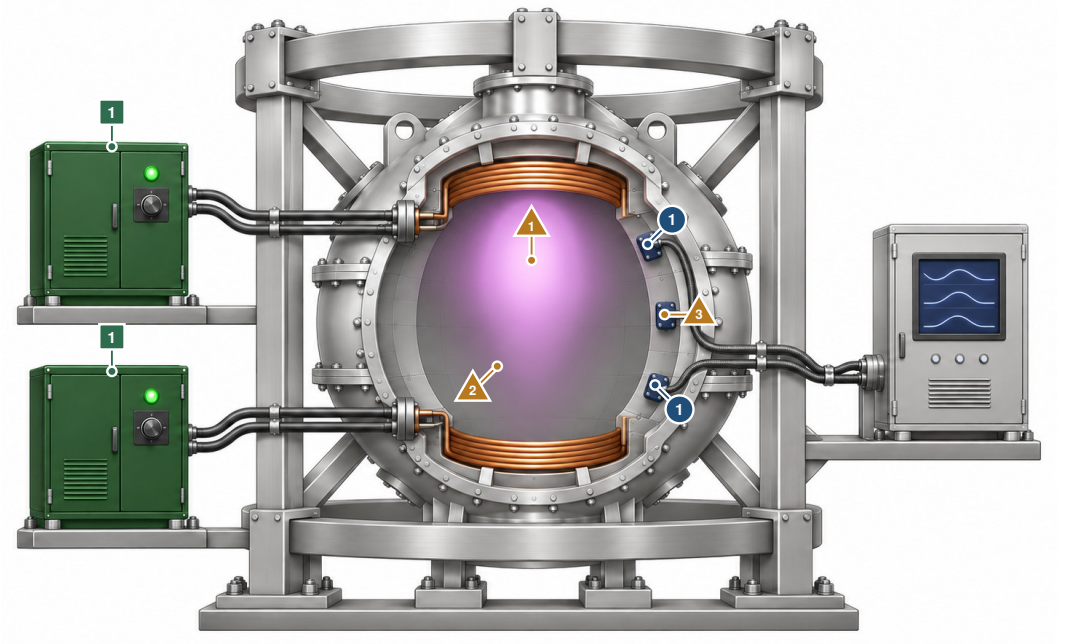

The controller moves 1 voltage of the upper and lower coil supplies

The controller reads 1 magnetic pickups on the vessel that sense the plasma's height and radius

What disturbs it 1 sudden kicks from edge bursts that jolt the plasma up or down

2 a drop in plasma pressure that shifts it inward 3 noise on the magnetic readings

The coils above and below the plasma are the controller's only push against a plasma that, once off center, falls faster and faster toward the wall. Illustration, not to scale.

IN PLAIN TERMS Holding Up a Plasma That Wants to Fall

The main idea. Inside a fusion machine, magnets stretch a ring of very hot gas into a tall shape that falls toward the wall once it moves off center. Coils above and below it push it back thousands of times each second, guided by magnetic sensors.

The challenge. Too gentle a push lets it drift away, and too hard a push, applied a little late, makes it shake. The coils also have a voltage limit, so a big enough jolt cannot be undone by any tuning.

Why it matters. If control is lost, the hot gas strikes the wall with its full electric current and puts large forces on the machine.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 12.2, the fact boxes of Sections 12.3 and 12.5, the try-it boxes of Section 12.5, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 12.5, Section 12.6, Section 12.7 and the open challenges

The control question An elongated plasma falls toward the wall within milliseconds, and a pair of coils with a voltage limit and a latency must catch it. Which limits does the tuning set, which does the supply set, and when does a controller that knows the voltage limit recover less than one that does not?

Learning objectives

- Compute the wall stability margin and the growth rate from the force balance and the wall and coil circuits.
- Locate the gain floor and the gain ceiling of the position loop on runs, and explain how elongation, the wall and the latency move them.
- Compare a filtered differentiator with a model observer as the velocity estimate under magnetics noise.
- Predict the voltage each allocation delivers to the vertical loop when the radial loop saturates the shared supply.
- Explain why a supply limit inside a short predictive plan does not enlarge the recoverable region.

Prerequisites Sections 1.4.6 (floor, region and ceiling), 1.6.2 (velocity form) and 1.6.4 (running a case). Section 1.6.3 is needed only for the explicit MPC. Volume 7, Chapter 6 is a slower unstable plant to meet first.

The app https://apmonitor.com/apps/tokamak-vertical/

See also Volume 7, Chapter 6 for an unstable suspension bounded below by its gain inequalities and above by saturation and noise. Volume 7, Chapter 7 for the largest recoverable push on a finite rail. Volume 4, Chapter 13 for allocation against hard limits that switch. Volume 4, Chapter 12 for several actuators sharing one budget.

12.1 Introduction

Control courses usually treat an unstable plant as a problem of placing an unstable pole in the left half plane, after which the loop is considered solved. Real unstable plants are held against three limits that a pole placement calculation hides. There is a minimum gain, below which the actuator cannot push back as hard as the plant runs away. There is a maximum gain, set by the latency and the measurement noise. And because the actuator saturates, there is a region of states outside which the plant cannot be recovered whatever the gains, so that the size of the disturbance matters as much as the tuning [1, 2]. Section 1.4.6 derives the three limits for the reduced vertical mode of the example plasma P-1.

The vertical position of an elongated tokamak plasma is a clean industrial home for all three. The shaping field that elongates the plasma produces a destabilizing force proportional to the displacement, the conducting vessel wall resists any motion with induced eddy currents, and the instability grows at a rate set by the wall time constant and a wall stability margin [3, 4]. The control coils act through a supply limited in voltage and a digital loop with a finite cycle, and losing control produces a vertical displacement event that drives the plasma into the wall with its full current and large forces on the structure [5].

The simulator addresses the belief that an unstable plant is safe once it has a stabilizing gain. The class of problem is general: magnetic bearings and levitation (Volume 7, Chapter 6), balancing robots (Volume 7, Chapter 7), aircraft flown with relaxed static stability and rockets landing on thrust all have a gain floor, a gain ceiling and a recoverable region, and the design of each begins with the growth rate and the actuator authority rather than with a tuning rule [6].

This chapter works with a vertical stabilization simulator in which the elongation, the wall, the supply, the latency and the noise can be changed and their consequence read as a growth rate, a recoverable margin and a loss or survival. The rigid massless model, position and velocity control, priority allocation and explicit predictive control are all established methods. What the simulator adds for study is a two-circuit model whose open-loop growth rate is exact and whose unstable mode and recoverable region are computed live, with a model-based velocity observer beside the usual filtered differentiator, and an explicit model predictive controller whose voltage limit follows the

allocation, compared with the loop on identical runs. The cases find the floor, the ceiling and the region one at a time.

Section 12.2 traces vertical control from the first elongated machines to the controllable-displacement specification of ITER. Section 12.3 sets out the force balance, the wall and coil circuits, the growth rate and the recoverable region. Section 12.4 describes the position and velocity loop, the two velocity estimates, the radial loop with its two allocations and the explicit MPC. Section 12.5 reports the measured floor, ceiling and region, the offset and allocation cases, the graded challenge and the comparison of the MPC with the loop. Section 12.6 turns to the sensors, supplies and control systems of real machines, Section 12.7 states what the model leaves out, and Section 12.8 summarizes the results.

12.2 Vertical Control of Elongated Tokamaks

Early tokamaks had circular cross-sections and were vertically stable. As machines of the late 1970s and 1980s elongated their plasmas to raise the current and the pressure a given toroidal field could hold, vertical stability became a design constraint, and active control of the vertical position with coils and fast power supplies became standard. A review of that period established the elements still used: a passive conducting structure that slows the instability, an active loop that stabilizes it, and a stability margin that measures how much slowing the passive structure provides [3]. The relation between elongation, aspect ratio and the achievable vertical stability was mapped soon after [7].

> ⧖ HISTORY **Letting the Plasma Drift to Measure It (1982)**
>
> Reactor designs of the early 1980s chose elongated plasmas for a higher limit on plasma pressure and a larger volume, yet the forces that stretch the plasma also make it move unstably upward or downward. A 1982 study on a divertor tokamak measured how strong that instability was. Each plasma was formed in a stable shape and moved radially into the unstable range with the vertical loop still holding it. The vertical feedback was then removed, not by switching off the amplifier but by holding its input at zero, so the currents induced by the motion still flowed. The plasma drifted up or down at 0.2 to 1.0 m/s, followed over 30 to 350 ms, and the drift gave the growth rate. A feedback amplifier of only 100 kW, aided by passive structure, held plasmas of 170 to 310 kA [8].

Control-oriented plasma response models followed. Linearized models of the coupled plasma, vessel and coil circuits made it possible to design controllers for current, position and shape together [9], and practice on large machines grew into integrated magnetic control systems [10]. Definitions of the stability margin were examined carefully, because a margin computed with a rigid displacement model and one computed with a deformable plasma can differ [4].

> ⧖ HISTORY **A Recoverable Displacement Becomes a Requirement (2008 to 2010)**
>
> By 2008 the design of ITER asked how far its plasma could drift before full voltage could no longer bring it back. Experiments on several tokamaks disabled vertical control for varying intervals to measure this maximum controllable displacement, and the measured values came out close to the calculated ones or well below them. The machines' experience implied that ITER would need a recoverable displacement above about 5 % of its minor radius [11]. A pair of in-vessel coils for vertical stabilization was then designed from a simulation in which the feedback is turned off, the plasma drifts 16.5 cm upward, and the feedback is restored. The coils had to work behind the blanket in a radiation field of about 3000 MGy and to deliver 30 000 such pulses over the life of the machine [12, 13].

The control-theoretic side matured in parallel. The conditions under which a massless plasma model is adequate for vertical stabilization were established [2], the multivariable structure of the vertical instability and its feedback stabilization were analyzed [6], and the effect of actuator saturation on the set of recoverable disturbances was treated directly, with a DIII-D application [1]. For ITER,

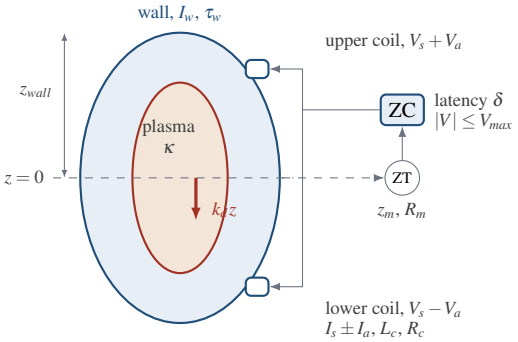

Figure 12.1 Vertical stabilization of an elongated plasma. The plasma of elongation κ, displaced by z from the midplane, feels the destabilizing force $k_d z$ and is held by the wall eddy current I_w with time constant τ_w and by the antisymmetric current I_a in the upper and lower in-vessel coils. The coils are driven with $V_s \pm V_a$, each within $\pm V_{max}$, the symmetric part V_s moving the plasma radially. The magnetics ZT measure z and R, the controller ZC acts after a latency δ, and the plasma is lost when it reaches the first wall at z_{wall}.

Table 12.1 Variables at the Operating Point

Symbol	Role	Nominal	Range
z	CV, vertical position	0 mm	SP $\pm$20
R	CV, radial position	0 mm	SP $\pm$30
V_a	MV, antisymmetric voltage	0 V	$\pm$10
V_s	MV, symmetric voltage	0 V	$\pm$10
I_w, I_a, I_s	states, currents	0	$\le$ 6 kA
R	state, radial position	0	—
κ	DV/parameter, elongation	1.80	1.50–2.02
n_z	DV, magnetics noise (1σ)	0.05 mm	0–1
—	DV, kicks (ELM, large, minor)	3, 20, 4 mm	events
—	DV, β drop	−15 mm	event

the maximum controllable displacement became a design requirement, as the second history box of this section tells. The forces on the vessel when control is lost are large enough that vertical displacement events are modeled in their own right [5].

12.3 Plasma, Wall and Coil Model

The model follows the rigid, massless displacement idea used in control analysis [2]. At elongation 1.8 the unstable growth rate with the coil voltage held is 356 s^{-1}, one e-fold every 2.8 ms, and the first wall is 150 mm from the plasma. The arrangement is shown in Figure 12.1.

The variables are classified in Table 12.1. Five states are carried: the wall current I_w, the antisymmetric and symmetric coil currents I_a and I_s, the radial position R, and a flag set at the loss of the plasma.

12.3.1 Force Balance and Circuits

The plasma is treated as a rigid current filament with negligible mass on these time scales, so the vertical forces on it balance at every instant, Equation (12.1).

$$k_d z + g_w I_w + g_c I_a = 0 \tag{12.1}$$

where k_d is the destabilizing stiffness, $g_w = 1.0\,\mathrm{N/A}$ and $g_c = 0.8\,\mathrm{N/A}$ the forces per ampere of wall and coil current. The position is therefore algebraic, $z = -(g_w I_w + g_c I_a)/k_d$. The wall and the coils are L-R circuits with mutual coupling, and a moving plasma induces a voltage in each, Equations (12.2) and (12.3).

$$L_w \dot{I}_w + M\dot{I}_a + g_w \dot{z} = -R_w I_w \tag{12.2}$$

$$M\dot{I}_w + L_c \dot{I}_a + g_c \dot{z} = V_a - R_c I_a \tag{12.3}$$

with $R_w = L_w/\tau_w$ and $M = k_M\sqrt{L_w L_c}$. Substituting $\dot{z}$ from the time derivative of Equation (12.1) gives a two-state linear system $\dot{x} = Fx + GV_a$ in $x = [I_w, I_a]^T$, solved at each step as a 2×2 linear system for the current rates.

12.3.2 Growth Rate and the Stability Margin

With the coils open-circuited the system has one unstable root, and its growth rate is exact, Equation (12.4).

$$\gamma\tau_w = \frac{1}{m}, \qquad m = \frac{g_w^2}{L_w k_d} - 1 \tag{12.4}$$

The wall stability margin m measures how much stiffer a perfectly conducting wall would be than the destabilizing push. The simulator ties it to elongation by $m(\kappa) = \max[0.02,\, 1.2(2.05 - \kappa)]$, so that $k_d = g_w^2/[L_w(1+m)]$ and a minor disruption scales m by $(1+\Delta m)$. At $\kappa = 1.8$, $m = 0.30$ and the wall alone would give $667\,\mathrm{s^{-1}}$. The coil circuit held at constant voltage also resists the motion, so the growth rate shown by the application, the unstable eigenvalue λ of F, is $356\,\mathrm{s^{-1}}$. It is $168\,\mathrm{s^{-1}}$ at $\kappa = 1.5$ and $1070\,\mathrm{s^{-1}}$ at 2.0.

12.3.3 The Recoverable Region

With w the left eigenvector of F for λ and $\xi = w^T x$ the unstable mode, the mode obeys Equation (12.5).

$$\dot{\xi} = \lambda\xi + bV_a, \qquad b = w^T G, \qquad |\xi| < \xi_{max} = \frac{|b|\,V_{max}}{\lambda} \tag{12.5}$$

Full voltage applied in the right direction can reverse the mode only while $|\xi| < \xi_{max}$, so the recoverable margin $1 - |\xi|/\xi_{max}$ is shown live. It grows smaller as λ rises and larger with the supply voltage, and it does not depend on the controller. It is Equation (1.10) written for the mode of the two-circuit model, with ξ in place of the position.

> **DID YOU KNOW? A Few Centimeters to Spare**
>
> On a research tokamak whose plasma grew vertically at $410\,\mathrm{s^{-1}}$, near the $356\,\mathrm{s^{-1}}$ of this chapter's plasma, the largest displacement that full control effort could still reverse was calculated at 0.8 cm, 3.7 % of the minor radius. That case was judged marginal, with a high likelihood of losing vertical control [11].

12.3.4 Radial Position, Limits and Disturbances

The symmetric coil current follows $L_c\dot{I}_s = V_s - R_c I_s$ and the radial position follows it through a first-order lag, $\tau_R\dot{R} = k_R I_s + R_{dist} - R$, with $k_R = 10\,\mathrm{mm/kA}$ and $\tau_R = 50\,\mathrm{ms}$. The coil currents $I_s \pm I_a$ are limited to 6 kA by scaling after each step, and the plasma is lost, with the simulation frozen, when $|z|$ reaches 150 mm. The magnetics add nearly white noise of 0.05 mm to z and half as much, correlated over 0.5 ms, to R. A kick of size Δz, which represents an edge localized mode (ELM) or a sawtooth crash, changes the wall current by $-k_d\Delta z/g_w$ at once, so the plasma is displaced instantly while the force balance still holds. A minor disruption cuts the margin by 40 % for 50 ms with a 4 mm kick, and a β drop moves the plasma and its radial equilibrium 15 mm inward.

The parameters are listed in Table 12.2.

Table 12.2 Plasma, Wall and Coil Parameters, With Hidden Ones Marked

Symbol	Meaning	Value	Class
τ_w	wall time constant	5.0 ms	visible
L_w	wall mode inductance	10 µH	hidden
g_w	wall force per ampere	1.0 N/A	hidden
g_c	coil force per ampere	0.8 N/A	visible
L_c, R_c	coil inductance, resistance	20 µH, 1.0 mΩ	visible
k_M	wall-coil coupling	0.30	visible
V_{max}	supply voltage limit	10 V	visible
—	coil current limit	6 kA	visible
δ	control and supply latency	100 µs	visible
k_R	radial shift per ampere	10 mm/kA	hidden
τ_R	radial response	50 ms	visible
z_{wall}	distance to first wall	150 mm	visible

12.3.5 Assumptions

The plasma is rigid and massless. Real plasmas deform as they move and the vertical mode couples to the shape, so a rigid-displacement margin can differ from a deformable one [4]. The massless idealization is adequate for stabilization analysis under conditions that have been stated [2].

One wall mode and one coil circuit represent the vessel and the coils. A real vessel has many eddy-current patterns, and the faster ones add delay to the coil action [9].

The margin is a fixed linear function of elongation. On a real machine it depends on the current profile, the pressure and the wall geometry, which is what a minor disruption changes.

The radial loop is a single lag coupled to the vertical loop only through the shared supply. The magnetics carry additive noise only, without pickup from the coils themselves or calibration drift.

12.4 Implemented Controller

Control executes each 100 µs and the equations are integrated with a 20 µs fourth-order Runge-Kutta step on the engine of Section 1.6.1. The voltages reach the coils after the latency δ through a delay line.

12.4.1 Position and Velocity Loop

The vertical loop uses the shared velocity-form PID of Section 1.6.2 on the measured position z_m, with the velocity term supplied as a live output bias recomputed at every scan, Equation (12.6).

$$\begin{aligned} e_k &= z_{SP} - z_{m,k}, \qquad s_k = s_{k-1} + K_p\,(e_k - e_{k-1}) \\ c_k &= s_k + I_{k-1} - K_d\,\hat{\dot{z}}_k \\ V_{a,k} &= \mathrm{clamp}\,(c_k, -V_{max}, V_{max}) \\ I_k &= I_{k-1} + \frac{K_p}{\tau_I}\,e_k\,\Delta t_c + \frac{\Delta t_c}{\tau_I}\,(V_{a,k} - c_k) \end{aligned} \tag{12.6}$$

where s_k accumulates the proportional increments without limit, c_k is the demanded voltage, I is the integral state, used at scan k as accumulated through scan $k-1$, and $\Delta t_c = 100\,\mu s$. In the symbols of Equations (1.13) to (1.16), s_k is S_k with no derivative increment, $-K_d\hat{\dot{z}}_k$ is the output bias u_b, c_k is v_k and $V_{a,k}$ is u_k. Anti-windup is the back-calculation term with a tracking time equal to τ_I (Section 1.6.2). Away from the voltage limit the recursion equals the positional law $V_a = K_p e - K_d\hat{\dot{z}} + (K_p/\tau_I)\int e\,dt$. The loop is reverse acting in the ISA sense, $a = +1$ in the convention of Section 1.3, because a positive antisymmetric voltage pushes the plasma up, so a plasma above the setpoint receives a negative voltage. The defaults are $K_p = 2.0\,\mathrm{V/mm}$ and $K_d = 1.0\,\mathrm{V}$ per m/s, and the integral is off, which the

interface represents by an integral time of 200 ms. The integral update and its back-calculation are then skipped. Transfer between manual and automatic is bumpless as Section 1.6.2 describes.

12.4.2 Two Velocity Estimates

The differentiator estimate is the first difference of z_m passed through a first-order filter of $\tau_f = 0.3\,\text{ms}$. The model observer avoids differentiating noise. It reads the wall current from the force balance, $\hat{I}_w = -(k_d z_m + g_c I_a)/g_w$, and computes the drift velocity from the circuit equations with the applied voltage set to zero, Equation (12.7).

$$\hat{\dot{z}} = -\frac{g_w \dot{\hat{I}}_w + g_c \dot{I}_a}{k_d}, \qquad [\dot{\hat{I}}_w, \dot{I}_a] = F\,[\hat{I}_w, I_a]^T \tag{12.7}$$

The constants in the force balance and in F are the declared values at the selected elongation, so a plant changed in the parameters, a faster wall or a minor disruption is not followed by the observer. Leaving out the instantaneous response to the coil's own voltage keeps the velocity term from chasing its own output through the latency. A counterfactual observer that includes it, run on the same engine with the noise at zero, sustains an oscillation after an ELM for velocity gains above about 1.3 V per m/s.

12.4.3 Radial Loop and Allocation

The radial loop is a PI on R_m with $K_c = 1.0\,\text{V/mm}$ and $\tau_I = 60\,\text{ms}$ setting V_s. The upper and lower coils receive $V_s \pm V_a$, and two allocations are offered, Equation (12.8).

$$\begin{aligned} \text{clip:} \quad & V_u = \text{clamp}(V_s + V_a),\; V_l = \text{clamp}(V_s - V_a) \\ \text{priority:} \quad & V_a' = \text{clamp}(V_a) \\ & V_s' = \text{clamp}\big(V_s, \pm(V_{max} - |V_a'|)\big) \end{aligned} \tag{12.8}$$

Clipping each coil on its own lets a saturated radial demand pin both supplies at the same limit, which reduces the antisymmetric voltage, half the difference of the two, to a fraction of what the vertical loop asks for. Priority allocation serves the vertical demand first and gives the radial loop what remains, on the reasoning that a lost vertical position is the costlier failure [5]. Sharing a request among actuators that each have a hard limit is the subject of Volume 4, Chapter 13. The back-calculation of each loop sees only its own $\pm V_{max}$ clamp, not the allocation or the latency that follow it, so an active integral, such as the radial one while priority allocation cuts V_s back, can keep winding while the coils receive less than its loop demanded.

12.4.4 Explicit Model Predictive Control

A second vertical controller can replace the position and velocity loop in automatic mode, while the radial PI and the allocation stay as they are. It is a model predictive controller (MPC) on the vertical channel whose quadratic program has been solved in advance, unlike the online MPC option of Section 1.6.3. Because the program's linear term and constraint limits move affinely with the measured state and the box, its minimizer is an affine function of them on each of a finite number of polyhedral regions, one for each optimal combination of active limits, and a decision reduces to finding the region and evaluating its law [14]. An explicit MPC of this kind has been proposed for the vertical stabilization of ITER [15], and Volume 7, Chapter 6 uses one on a suspended ball. The loop remains the default, and only case 11 selects the MPC.

The controller's model is Equations (12.1) to (12.3) with the declared constants of Table 12.2 and the stiffness k_d of the elongation the operator sets, without the margin drop of a minor disruption, and it is rebuilt whenever the elongation changes. Eliminating I_w with the force balance writes it in the measured position and coil current, $x = [z, I_a]^T$, as $\dot{x} = F_z x + G_u(V_a + d_c) + G_w d_w$, where d_w and d_c are constant voltage errors on the wall and coil circuits. The declared latency equals one control period,

so the voltage applied over the next 100 µs is the command u_p decided at the previous execution, and u_p joins the state. An exact zero-order hold over $T_s = \Delta t_c = 100\,\mu s$ gives Equation (12.9).

$$\xi_{k+1} = A\,\xi_k + B\,u_k, \qquad A = \begin{bmatrix} \Phi & \Gamma_u \\ 0 & 0 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{12.9}$$

where $\xi = [z, I_a, u_p]^T$, $\Phi = e^{F_z T_s}$ and Γ_u is the held voltage's column. At $\kappa = 1.8$ the channel's roots are $+356$ and $-77\ \mathrm{s}^{-1}$, the first being the growth rate of the application. The model is written from the chapter's own linear equations rather than identified, and an independent derivation from the simulator's circuit function reproduces A to 10^{-12} and the controller gain below to a relative 3×10^{-10} at $\kappa = 1.8$.

Each execution plans three single-period voltage moves, a 0.3 ms horizon, and applies the first, as shown in Equation (12.10).

$$\begin{aligned} &\min_{u_0,u_1,u_2} \sum_{k=1}^{2} Q_z\,\tilde{z}_k^2 + \tilde{\xi}_3^T P\,\tilde{\xi}_3 + R\sum_{k=0}^{2}\tilde{u}_k^2 \\ &\text{subject to } \tilde{\xi}_{k+1} = A\tilde{\xi}_k + B\tilde{u}_k, \quad \tilde{\xi}_0 = [z_m, I_a, u_p]^T - \xi_t \\ &\qquad -b \le u_t + \tilde{u}_k \le b, \quad k = 0,1,2 \end{aligned} \tag{12.10}$$

where the tilde marks a deviation from the target $\xi_t = [z_{SP}, I_{a,t}, u_t]^T$, b is the half-width of the voltage box and P is the solution of the discrete algebraic Riccati equation for A, B, $\mathrm{diag}(Q_z, 0, 0)$ and R. Charging the infinite-horizon cost after the last move, without a terminal set, is the usual stabilizing choice for a receding horizon [16]. The weights follow Bryson's rule on the 0.5 mm band of the vertical loop and the declared 10 V supply, $Q_z = 4 \times 10^6\,\mathrm{m}^{-2}$ and $R = 0.01\,\mathrm{V}^{-2}$, and were fixed before any comparison was run. With no limit active the first move is the LQR's, $\tilde{u}_0 = -K\tilde{\xi}_0$, with $K = (11.5\,\mathrm{V/mm}, 0.0104\,\mathrm{V/A}, 0.785)$ at $\kappa = 1.8$ and $(9.90\,\mathrm{V/mm}, 0.0110\,\mathrm{V/A}, 0.891)$ at 1.9, so its position gain is five to six times the loop's default K_p and the third entry feeds back the voltage still on its way. The closed loop of $A - BK$ has a spectral radius of 0.958.

The only limits are the box on every move, all hard, and because only inputs are constrained the program is feasible whenever $b \ge 0$. The box is what the allocation leaves the vertical channel after the radial command V_s of the same scan. Under priority allocation $b = V_{max}$. Under clipping it is the largest antisymmetric voltage the coil-by-coil clip can deliver with a command inside $\pm V_{max}$, $b = V_{max} - |V_s|/2$, and a planned voltage u is commanded as u while $|u| \le V_{max} - |V_s|$ and as $\mathrm{sign}(u)(2|u| - V_{max} + |V_s|)$ beyond, which puts one coil at its limit and delivers exactly u. On 20 000 random supplies and radial commands the clip's reach equals $V_{max} - |V_s|/2$ to 7×10^{-15} V, and in closed loop the antisymmetric voltage reaching the coils one period later equals the plan's first move to 7×10^{-11} V. A first version that gave the MPC $V_{max} - |V_s|$ under clipping was lost in four clipped cases that both forms of the loop hold, because it claimed less than the clip itself delivers. The clip, the priority rule, the coil current limit and the first wall act on the MPC's command exactly as on the loop's.

The parameter of the program is $\theta = (z_m - z_{SP},\, I_a - I_{a,t},\, u_p - u_t,\, -b - u_t,\, b - u_t)$, so the box enters only the right-hand sides and one table serves every supply limit, radial share and allocation rule. Three moves, each free or at one of two edges, give $3^3 = 27$ combinations, and all 27 are stored at every elongation from 1.50 to 2.02. The table is complete, since every θ with a nonempty box lies in one of the regions, so a state leaves it only through a number that is not finite. Restricting the pending voltage to one the supply could have delivered, 21 regions have an interior on the default plant at every elongation (21 to 23 over every parameter setting), and 13 were met in 112 321 closed-loop decisions. Each lookup checks the stored regions in sequence, starting with the unconstrained law, against the signs of its multipliers and the six box rows. The row test is relative, 10^{-12} times the largest of 1, $|b|$

and the unconstrained moves, because far from the midplane the unconstrained move reaches 10^3 V, and a fixed 10^{-12} V test misses 2 of those 112 321 decisions. On 199 321 states, 87 000 with a limit active and 15 000 at region boundaries, the lookup agrees with the apps' shared dual active-set solver, run online on the identical program, to 1.2×10^{-11} V in the first move. Rebuilding the design and the table for a new elongation takes 0.38 ms at the median in Node.js, inside the execution that sees the change.

The MPC is offset-free by a target shift [17]. At each execution the measured z_m and I_a are compared with their one-step prediction, the two errors are mapped onto d_w and d_c through the disturbance columns of the discretized model, and each estimate is moved through a first-order filter of 50 ms, ten declared wall times, and clamped to ± 4 V and ± 5 V. The target is the declared equilibrium that holds the setpoint against them, Equation (12.11).

$$I_{w,t} = \frac{d_w}{R_w}, \quad I_{a,t} = -\frac{k_d z_{SP} + g_w I_{w,t}}{g_c}, \quad u_t = R_c I_{a,t} - d_c \tag{12.11}$$

With the coil force per ampere at 0.7 on case 9's 10 mm step and no noise, the MPC ends 0.000 mm from the setpoint and the loop with the differentiator 0.58 mm away.

A region not found, or an online solve that is not optimal, rejects the decision. The loop, which tracks the applied command at every execution, then takes the coils, its offset cleared and its voltage ramped from the last command to its own law over 20 decisions (2 ms), the record is marked, and after 20 decisions found in a row the MPC returns over another 2 ms. A switch to the loop during a run is handed over the same way. Misses forced with the noise off give steps of 0.27 V in the recovery from an ELM and 0.49 V at the reversal after a large kick, and no measured run missed. The plant integrates the run's parameters and the disruption's margin drop, while the MPC and the observer hold the declared model, so any parameter a learner edits becomes model mismatch for both. The MPC's estimator runs under every controller and writes nothing the loop reads, so adding the MPC changed no earlier run under the loop.

12.5 Results

Table 12.3 lists the twelve cases in the order the application presents them, with the point each teaches and what completes it. A case is credited when the learner has performed the listed actions and the simulation has reached the listed state, and the credits accumulate toward certificates of progress and of completion. The challenge is judged on the case as given, so a change to a model parameter or disturbance setting, or firing an extra disturbance, voids the attempt.

The values in Table 12.4 were read from runs at the default seed of the published app (Section 1.6.4), each case selected, changed and then run from $t = 0$. An rms position is the application's running value for the true z, sampled at every control scan from the start of a run of the stated length, except in case 5, where it covers the last 0.3 s of 0.6 s runs. Peaks and settling times are taken at the integration step. Rerunning each closed-loop configuration with its settings carried through a reset gives the same run. With other noise seeds the loss times of cases 3, 4 and 8 move by up to 47 ms. The suggested challenge settings pass at all six seeds tested, while priority allocation alone passes at three and fails at three.

12.5.1 Growth Rate and the Gain Floor

Held at the midplane with 0.05 mm of magnetics noise, the true position wanders 0.023 mm rms over 0.2 s, and nothing in the trends suggests an instability growing at 356 s^{-1}. With the coil voltage held at zero, a 0.1 mm displacement reaches the wall in 20.0 ms at elongation 1.8, 42.3 ms at 1.5 and 6.7 ms at 2.0, about seven e-folding times in each case. Choosing the elongation before Run or carrying it through Reset initializes the same 0.1 mm displacement.

Table 12.3 Scenarios and Completion Requirements

#	Scenario and kind	Teaching point	Required to complete
1	Held at the midplane *Experiment*	a quiet unstable loop	After the first 0.1 s, fire the 3 mm ELM kick, and the loop holds the plasma within 0.5 mm of the setpoint for 20 ms.
2	Let go (manual) *Demonstration*	the growth rate	With the coil voltage held at zero, let the plasma reach the first wall at $\kappa = 1.8$, then reset and repeat with κ set to 1.5 and to 2.0 before each run.
3	Too little gain *Demonstration*	the gain floor	Lose the plasma at $K_p = 0.3\,\mathrm{V/mm}$, then reset and run 0.2 s with the plasma held at 0.5 V/mm.
4	Stretch it *Demonstration*	the floor moves with κ	Run $K_p = 0.5\,\mathrm{V/mm}$ through the change from $\kappa = 1.8$ to 2.0 at 0.1 s until the plasma is lost or $t = 0.2$ s, then reset and run 0.2 s at 1.0 V/mm.
5	Too much gain *Experiment*	the ceiling from latency	With 300 µs of latency, run 0.1 s at $K_p = 10\,\mathrm{V/mm}$, then 50 ms at 5 and 50 ms at 2 V/mm, then 0.1 s at 10 V/mm with the latency set back to 100 µs.
6	Velocity from a noisy position *Experiment*	differentiator against observer	With 0.3 mm of noise and $K_d = 2$ V per m/s, run 0.1 s each on the differentiator, on the model observer, and with $K_d = 0$.
7	The recoverable region *Experiment*	the size of the region	Run past the 20 mm kick at 50 ms (to loss or 0.15 s) at $\kappa = 1.9$ as given, with $K_p = 10\,\mathrm{V/mm}$, at $\kappa = 1.8$, and at $\kappa = 1.9$ with $V_{max} = 15$ V, each change made before its run.
8	A faster wall *Demonstration*	the passive structure	Run $K_p = 0.6\,\mathrm{V/mm}$ through the halving of τ_w at 0.1 s until the plasma is lost or $t = 0.2$ s, then reset and run 0.2 s at 2 V/mm.
9	Move it up *Experiment*	offset and integral	Run to 0.15 s through the setpoint step to 10 mm at 50 ms with no integral action, then reset and run 0.15 s with $\tau_I = 20$ ms.
10	Two jobs, one supply *Experiment*	allocation	Run with clipped allocation past the β drop and kick until the plasma is lost or $t = 0.15$ s, then reset and run 0.15 s with vertical priority.
11	Explicit MPC against PD *Experiment*	a limit in the plan	Run under the explicit MPC through the ELM at 50 ms and the 19.7 mm kick at 0.15 s, then reset, select the loop before 50 ms and run the same case to 0.25 s.
12	Challenge at $\kappa = 1.95$ *Challenge*	graded run	Keep the plasma for 1.2 s, with $\lvert z - z_{SP}\rvert \le 4$ mm and $\lvert R - R_{SP}\rvert \le 1$ mm over the last 0.4 s, on the case as given.[a]

[a] As given, the plasma is lost at 661.0 ms. Priority allocation alone passes at the default seed, with 3.94 mm vertically and 0.80 mm radially over the last 0.4 s. The suggested settings ($K_p = 2\,\mathrm{V/mm}$, observer at $K_d = 1.5$ V per m/s, priority) pass with 3.23 mm and 0.18 mm. Before-Run and Reset settings give the same run. The verdict checks each 20 µs physics step and each kick instant.

TRY IT IN THE APP **2 · Let Go (MANUAL)**

Load scenario 2 and press Run. The loop is in manual with the coil voltage held at zero, and the plasma starts 0.1 mm above the midplane. Watch the vertical position grow faster and faster until the plasma reaches the wall at 20.0 ms. Press Reset, set the elongation to 1.5 before Run, and the same fall takes 42.3 ms.

The loop has a floor. At a position gain of 0.3 V/mm with no velocity term the plasma drifts away and is lost at 166.6 ms, a slow failure that looks like drift rather than instability. The floor at $\kappa = 1.8$, the gain at which a noise-free kick neither grows nor decays with the 100 µs latency, is about 0.45 V/mm. At 0.4 the plasma survives the first 0.2 s within 0.5 mm but is lost at 482 ms, and at 0.5 it holds to 0.023 mm rms over 0.2 s. The floor moves with the plasma. When the elongation is raised from 1.8 to 2.0 during a run at 0.5 V/mm, the growth rate triples to $1070\,\mathrm{s}^{-1}$, the floor rises to about 0.70 V/mm, and the plasma is lost 56 ms after the change, while 1.0 V/mm holds it to 0.023 mm rms. Halving the wall time raises the growth rate from 356 to $680\,\mathrm{s}^{-1}$ and the floor from about 0.45 to about 1.03 V/mm, so 0.6 V/mm loses the plasma at 158 ms and 2 V/mm holds it to 0.021 mm rms.

Pause. Try Exercise 12.1 before reading on.

Table 12.4 Measured Results

Case and condition	Result
1, noise 0.05 mm, 0.2 s	true z 0.023 mm rms
2, from 0.1 mm, κ 1.5 / 1.8 / 2.0	wall at 42.3 / 20.0 / 6.7 ms
3, $K_p = 0.3$	lost at 166.6 ms; floor $\approx$ 0.45 V/mm
3, $K_p = 0.4$	lost at 482 ms
3, $K_p = 0.5$, 0.2 s	0.023 mm rms
4, κ 1.8 to 2.0, $K_p = 0.5$	lost 56 ms later; floor $\approx$ 0.70 V/mm
4, $K_p = 1.0$, 0.2 s	0.023 mm rms
5, $\delta = 300\,\mu$s, K_p 10 / 5 / 2	1.45 / 0.063 / 0.022 mm rms
5, $K_p = 10$, saturation	64 % of the time
5, $\delta = 100\,\mu$s, $K_p = 10$	0.059 mm rms
6, noise 0.3 mm, 0.35 s, differentiator	0.216 mm, 33.9 J, 31,000 V/s
6, observer	0.116 mm, 9.1 J, 9,800 V/s
6, no velocity term	0.104 mm, 6.1 J
1, 3 mm ELM, K_d 0 / observer 1	overshoot 0.52 / 0.10 mm
7, $\kappa = 1.9$, 20 mm kick	margin 7 %, default gains lost at 67.9 ms
7, $K_p = 10$	lost at 58.9 ms
7, $\kappa = 1.8$	margin 28 %, held, peak 23.1 mm
7, $V_{max} = 15$ V	margin 38 %, held, peak 21.4 mm
8, τ_w halved, $K_p = 0.6$	γ 356 to 680 s^{-1}, lost at 158 ms
8, floor at $\tau_w = 2.5$ ms	$\approx$ 1.03 V/mm
8, $K_p = 2$, 0.2 s	0.021 mm rms
9, SP 10 mm, no integral	settles 0.51 mm high, peak 12.3 mm
9, $\tau_I = 20$ ms	within 0.1 mm in 41 ms (34 ms logged)
10, clip	lost at 72.7 ms
10, priority	held, peak 10.7 mm

12.5.2 The Gain Ceiling and the Velocity Estimate

With 300 µs of latency, a position gain of 10 V/mm settles into a limit cycle of 1.45 mm rms with the supply saturated 64 % of the time. The plasma is still held, with no reserve. At 5 V/mm it holds to 0.063 mm and at 2 V/mm to 0.022 mm, and cutting the latency to 100 µs lets 10 V/mm hold to 0.059 mm. The window between the floor and the ceiling, 0.45 V/mm to somewhere between 5 and 10 V/mm at this latency, is what the loop is designed within. For the reduced mode, Equation (1.11) shows the same fall of the ceiling as the latency grows.

🕮 **Pause.** Try Exercise 12.2 before reading on.

With the noise raised to 0.3 mm, a velocity gain of 2 V per m/s on the filtered differentiator leaves the position at 0.216 mm rms over a 0.35 s run, moves the voltage about 31 000 V/s summed over the control scans and delivers 33.9 J to the coils. The model observer gives 0.116 mm, about 9800 V/s and 9.1 J. With no velocity term at all the position is 0.104 mm rms for 6.1 J, so on noise alone the velocity term does not pay. It pays on kicks. With the noise at zero, a 3 mm ELM overshoots by 0.52 mm and settles within 0.15 mm in 5.3 ms without a velocity term, and overshoots by 0.10 mm and settles in 1.5 ms with the observer at a velocity gain of 1.

12.5.3 The Recoverable Region

At $\kappa = 1.9$ a 20 mm kick leaves a recoverable margin of 7 %. The voltage saturates and the plasma is lost at 67.9 ms, and raising the gain to 10 V/mm only loses it sooner, at 58.9 ms. Lowering the elongation to 1.8 cuts the growth rate from 542 to 356 s^{-1}, the margin after the same kick is 28 %, and the plasma is held with a peak of 23.1 mm. Raising the supply to 15 V at $\kappa = 1.9$ enlarges the region instead, to a margin of 38 %, and also holds it, with a peak of 21.4 mm. The kick at $\kappa = 1.9$ lands inside the region, so the tuning still decides whether the plasma is kept: a stronger velocity term, which reverses the voltage earlier, also holds it on the 10 V supply, and so does a weaker position gain. Gain decides how the loop behaves inside the region, while the growth rate and the supply voltage decide how large the region is, which is the reasoning behind specifying a maximum controllable displacement for a machine [1, 11].

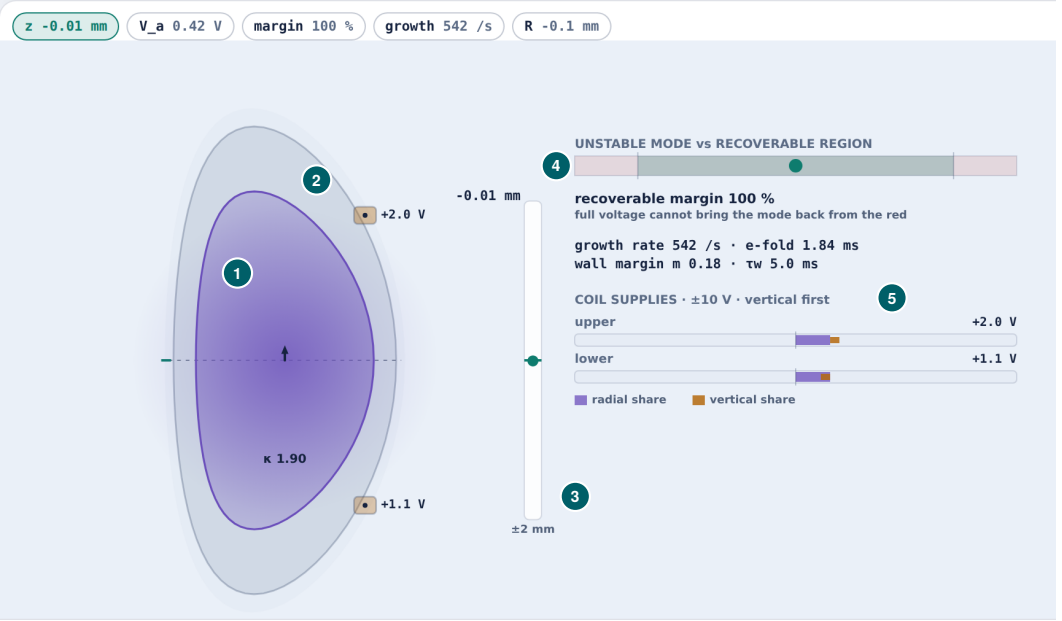

Figure 12.2 The vertical stabilization application running scenario 10 with priority allocation set before the run, at $t = 0.25$ s, 200 ms after a β drop and 195 ms after a 10 mm vertical kick at elongation 1.9, cropped to the plant view. The cross-section shows the vessel and the plasma at its elongation (1), the two in-vessel coils with their voltages (2) and a zoomed ±2 mm position gauge (3). Beside it the unstable mode is drawn against the recoverable region (4), above the growth rate of 542 s^{-1} and wall margin of 0.18, and the two coil supplies are split into radial and vertical shares (5). The plasma is held, back at the midplane.

> ✦ DID YOU KNOW? **More Gain, Lost Sooner**
>
> From the simulator: at elongation 1.9 a 20 mm kick leaves a recoverable margin of 7 %. The default gains lose the plasma at 67.9 ms, and a position gain of 10 V/mm, five times the default, loses it sooner, at 58.9 ms (Table 12.4).

12.5.4 Offset and Allocation

Moving the setpoint 10 mm up shows why vertical controllers carry a slow integral or a separate position loop. Holding the plasma off the midplane takes about −0.96 kA in the coils and −1.0 V across their resistance, so the proportional loop settles 0.51 mm above the setpoint after a peak of 12.3 mm. A 20 ms integral time brings the true position within 0.1 mm of the setpoint for good 41 ms after the step, while the trend logged at 2 kHz is inside that band from 34 ms, the noise carrying the plasma to the edge of the band in between.

When a β drop throws the plasma 15 mm inward and a 10 mm vertical kick follows 5 ms later at $\kappa = 1.9$, the radial loop asks for more than the supply has. With each coil clipped on its own, the plasma is lost at 72.7 ms, by which time the radial position has recovered only from 15 to 4.9 mm inward. With the vertical demand served first it is held with a peak of 10.7 mm (Table 12.4; Figure 12.2 shows the plasma held), and the radial position is back within 1 mm for good about 107 ms after the drop. The recoverable margin reads 53 % at the kick in both cases, because it assumes the vertical loop has every volt.

> **TRY IT IN THE APP** **10 · Two Jobs, One Supply**
>
> Load scenario 10 and press Run. At 50 ms a β drop throws the plasma inward and a 10 mm vertical kick follows. With each coil clipped on its own the plasma is lost at 72.7 ms. Press Reset, set the coil voltage allocation to Vertical first before Run, and the plasma is held with a peak of 10.7 mm.

12.5.5 The Graded Challenge

The challenge runs 1.2 s at $\kappa = 1.95$ with 0.2 mm of noise, two ELMs, a minor disruption, a β drop with a 10 mm kick and two more ELMs, starting from $K_p = 0.8\,\mathrm{V/mm}$, a differentiator velocity gain of 3 and clipped allocation. It is met when the plasma is kept and, over the last 0.4 s, the vertical position stays within 4 mm and the radial position within 1 mm. The verdict checks every 20 µs physics step and each instantaneous kick, so brief ELM excursions count even when they fall between the 2 kHz log samples. As loaded the plasma is lost at 661.0 ms, 11 ms after the β drop. Priority allocation alone keeps the plasma and meets the target at the default seed, with 3.94 mm vertically and 0.80 mm radially, although the differentiated noise saturates the vertical supply 22 % of the logged time. It passes at two of five additional seeds and fails at the other three, including a 5.02 mm vertical excursion at seed 3. Every tested combination of the observer, a velocity gain of 1.5 and $K_p = 2\,\mathrm{V/mm}$ that keeps clipped allocation loses the plasma within 0.3 ms of the original time. With $K_p = 2\,\mathrm{V/mm}$, the observer at a velocity gain of 1.5 and priority allocation, the last 0.4 s stay within 3.23 mm and 0.18 mm, including the ELM kicks. The supply saturates 0.1 % of the logged time and the position is 0.34 mm rms over the whole run. These suggested settings pass at all six seeds tested, with worst vertical errors of 3.17 mm to 3.33 mm.

12.5.6 Explicit MPC Against the Position and Velocity Loop

Case 11 holds the plasma of case 7 at $\kappa = 1.9$ under the explicit MPC, fires a 3 mm ELM at 50 ms and a 19.7 mm kick at 150 ms, a size chosen between the largest kicks the MPC and the loop recover. It is credited when the case has been run under the MPC and then, after Reset, under the loop selected before the ELM. After the ELM the MPC is back within 0.5 mm of the setpoint in 0.42 ms, against 1.02 ms for the loop with the differentiator and 0.62 ms with the observer, read at the control scans. The large kick leaves a recoverable margin of 7 % under each controller, and each holds the full −10 V for 7.5 to 7.8 ms. The MPC's command turns positive 7.82 ms after the kick with the plasma 3.8 mm from the midplane, the loop's at 7.72 ms with the plasma at 7.7 mm. Within the next 4 ms the MPC's unstable mode swings to 1.14 times ξ_{max} on the other side, past the edge of the recoverable region, and the plasma is lost at 168.0 ms, while the loop's mode reaches 0.93 times ξ_{max} and the plasma is held with a largest excursion of 34.2 mm, as shown in Figure 12.3.

Every row of Tables 12.5 to 12.7 uses one engine, noise seed, case script and duration, the loop at each case's own gains with each velocity estimate and both allocations, and the MPC at case 11's settings with nothing retuned per case. Two further rows isolate the MPC's ingredients. The loop retuned once by the MPC's own rule takes the LQR's gains on z and on the declared drift velocity at the elongation in force ($K_p = 8.55\,\mathrm{V/mm}$ and $K_d = 2.02\,\mathrm{V}$ per m/s at $\kappa = 1.9$), without the latency term, inside the MPC's box. The MPC with no box rows is the same law, estimator and target with the box removed from the program and the first move clipped afterwards, the strongest classical form of the MPC's own gain.

The MPC loses the headline comparison, the largest recoverable kick. At $\kappa = 1.9$ it recovers kicks up to 19.63 mm with clipping against 19.91 mm and 19.86 mm for the loop with each estimate, and at 1.8 up to 25.15 mm against 25.65 and 25.50 mm, as listed in Table 12.5 and drawn in Figure 12.4. The loop retuned to the LQR's gains recovers 19.63 mm and 24.81 mm with clipping (bisected to 0.05 mm). The plan sees 0.3 ms, and its terminal cost prices the motion after that as if any voltage were available, so the MPC holds −10 V until the plasma is nearer the midplane and the latency carries the swing past the region's far edge. Its knowledge of the box is worth 0.05 to 0.13 mm of recoverable kick over the same law without the rows. Held over a 2.1 ms horizon the same 27-region

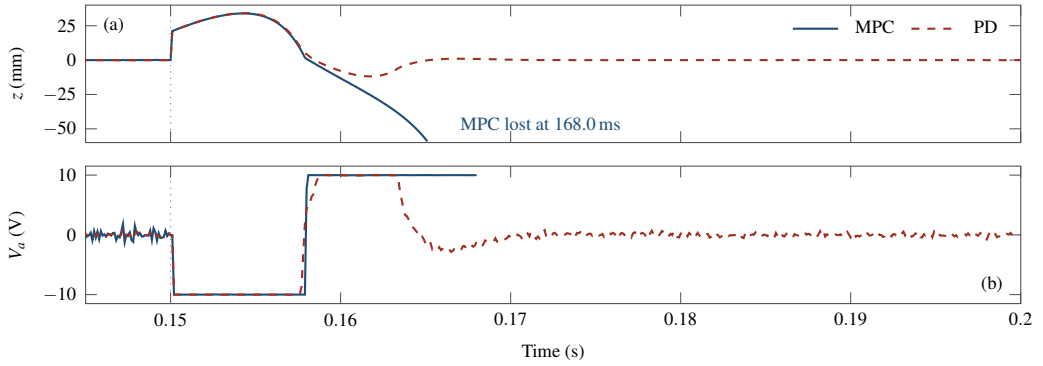

Figure 12.3 Case 11 to $t = 0.200$ s, run first with the explicit MPC and then, after Reset, with the loop and its differentiator. (a) Vertical position. (b) Antisymmetric voltage at the coils. The MPC is the solid line and the loop the dashed line. After the 19.7 mm kick at the dotted line both hold −10 V. The MPC reverses later and closer to the midplane, and its plasma crosses to the wall on the other side, while the loop's settles.

Table 12.5 Largest Recoverable Kick (mm) at 50 ms, Default Noise

κ, allocation	PD, diff.	PD, obs.	MPC	No rows	Held[a]
1.9, clip	19.91	19.86	19.63	19.58	20.19
1.9, priority	19.92	19.87	19.68	19.60	20.19
1.8, clip	25.65	25.50	25.15	25.03	26.27
1.8, priority	25.67	25.52	25.21	25.08	26.31

Bisection to 0.01 mm on case 7 without its scripted kick, 0.25 s runs. "No rows" is the MPC's law without its box rows, the first move clipped.
[a]Counterfactual: the same three moves held over 0.1, 0.4 and 1.6 ms, a 2.1 ms horizon, not offered in the application.

structure sees the braking, recovers 20.19 mm at $\kappa = 1.9$ and holds case 11, and over 5.1 ms it reaches 20.35 mm. No controller tested recovers more than about 20.4 mm at $\kappa = 1.9$ on a 10 V supply, so the supply's voltage limit decides how large a displacement can be recovered under every controller, and the look-ahead decides only how much of that region is used. Halving the integration step changes none of the rankings (19.91, 19.86 and 19.63 mm).

🕮 **Pause.** Try Exercise 12.3 before reading on.

The MPC is the fastest controller on every small kick, as listed in Table 12.6. After the ELM it is back within 0.5 mm in 0.52 ms against 1.32 and 1.02 ms, at 1.4 J of coil energy against 0.9 and 1.2 J, after the minor disruption in 0.52 ms against 1.52 and 0.82 ms, and at $\kappa = 1.95$ in 0.32 ms against 1.02 and 0.42 ms. The no-rows law gives the same numbers on all of these cases and on the noise and the challenge, so the difference from the loop is the LQR gain with the latency in its model, not the knowledge of the limit. The loop given the LQR's gains without that latency term limit-cycles at $\kappa = 1.95$, saturated for 76.2 ms of the 0.3 s run. The MPC loses the noise case. Under 0.3 mm of magnetics noise its position is 0.241 mm rms against 0.119 mm for the loop with the observer, for 87.9 J against 16.0 J and 14.6 ms saturated against none, because its position gain of 11.5 V/mm passes the noise to the coils. Taken at every integration step instead of the scans, the same runs give 0.212 against 0.117 mm and 81.5 against 16.0 J. In case 10 with clipping every controller is lost, the MPC 0.3 ms sooner, and with priority every controller holds a peak near 10.7 mm. In the challenge under priority allocation the MPC meets the target with 3.17 mm vertically against 3.23 mm for the suggested loop, at 27 % more coil energy (5082 J against 3991 J).

The worst-case table runs 53 cases under each controller and both allocations, as summarized in Table 12.7. The MPC loses 18 of 106 runs against 24 and 26 for the loop. It holds 8 runs a form of the loop loses, cases 3, 4 and 8 under both allocations, where the shipped low gains are the lesson

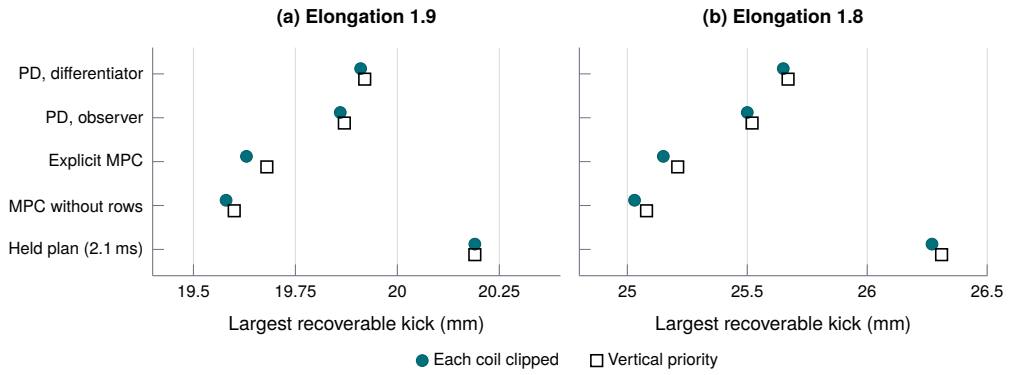

Figure 12.4 The largest kick at 50 ms that each controller recovers, from Table 12.5, at elongation 1.9 (a) and 1.8 (b). Filled circles: each coil clipped. Open squares: vertical priority. The last row is the counterfactual plan held over 2.1 ms, which the application does not offer. Notice that the explicit MPC recovers less than either form of the loop, that removing its box rows costs only a little more, and that the allocation barely matters here: the look-ahead, not the knowledge of the limit, decides how much of the region is used.

Table 12.6 Explicit MPC Against the Loop on the Shipped Cases

Case	Measure	PD, diff.	PD, obs.	MPC	PD, LQR	No rows
1, 3 mm ELM, no noise	back within 0.5 mm (ms)	1.32	1.02	0.52	1.02	0.52
	peak $\lvert V_a \rvert$ (V) / coil energy (J)	10.00 / 0.9	7.69 / 1.2	10.00 / 1.4	10.00 / 3.2	10.00 / 1.4
1, minor disruption, no noise	back (ms) / saturated (ms)	1.52 / 0.3	0.82 / 0.1	0.52 / 0.5	50.22 / 20.6	0.52 / 0.5
1 at $\kappa = 1.95$, ELM, no noise	back (ms) / saturated (ms)	1.02 / 0.2	0.42 / 0.0	0.32 / 0.3	199.92 / 76.2	0.32 / 0.3
6, noise 0.3 mm, 0.6 s, clip	rms (mm) / energy (J)	0.218 / 58.7	0.119 / 16.0	0.241 / 87.9	0.428 / 152.5	0.241 / 87.9
	saturated (ms)	0.5	0.0	14.6	68.4	14.6
10, clip	lost at (ms)	72.7	72.7	72.4	74.7	72.1
10, priority	back (ms) / peak (mm)	61.7 / 10.70	56.9 / 10.67	56.7 / 10.71	280.0 / 10.77	56.7 / 10.71
Challenge, priority[a]	last 0.4 s, z / R (mm)	3.94 / 0.80	3.23 / 0.18	3.17 / 0.20	4.39 / 0.62[b]	3.17 / 0.20
	coil energy (J) / saturated (ms)	11,415 / 270.9	3,991 / 2.2	5,082 / 2.8	11,380 / 288.3	5,082 / 2.8
11	ELM back (ms) / 19.7 mm kick	1.02 / held	0.62 / held	0.42 / lost 168.0	48.72 / lost 165.8	0.42 / lost 166.2

"Back" runs from the event until $\lvert z - z_{SP} \rvert$ stays within 0.5 mm, read at the 10 kHz control scans, "saturated" counts control periods with the vertical voltage at its limit or cut by the allocation (under the MPC also a plan on its box), and the rms is the application's running value over the run. Every MPC run decided every 100 µs with no rejected decision, and the online solve reproduces every MPC entry to the printed digit. "PD, LQR" is the loop with the LQR's gains and "No rows" the MPC's law without its box rows. [a]The differentiator column is the shipped tuning ($K_p = 0.8\,\mathrm{V/mm}$, $K_d = 3$) and the observer column the suggested one ($K_p = 2$, $K_d = 1.5$). With clipping every controller loses the plasma at 660.5 to 661.2 ms. [b]Fails the 1 mm radial target.

and the MPC's gain is fixed by its weights, and case 1 with a 1 ms wall, which the observer built on the declared 5 ms wall loses at 33.3 ms. In 12 runs every controller is lost and the MPC 0.2 to 2.4 ms sooner (case 7, case 10 and the challenge with clipping, and a learner's $g_w = 3$, $g_c = 0.1\,\mathrm{N/A}$, $L_c = 100\,\mu\mathrm{H}$ or $k_M = 0.8$). At the ends of the parameter ranges it is never lost where the loop holds. Between the ends, a sweep of 74 mismatched runs finds three where it is: case 9's 10 mm step with $g_c = 0.6\,\mathrm{N/A}$, $k_M = 0.5$ or $\tau_w = 2\,\mathrm{ms}$ is lost by the MPC at 66.1, 62.9 and 57.7 ms and by the loop with the observer, which holds the same declared model, while the loop with the differentiator holds it within 0.69 mm. The application states both the noise loss and these mismatch losses beside the controller.

In Node.js a lookup takes 0.08 µs at the median closed-loop decision and the online solve 7.2 µs.

Table 12.7 Worst Case Over 53 Cases and Both Allocations (106 Runs Each)

Output	PD, diff.	PD, obs.	MPC
runs lost	24	26	18
largest $\lvert z \rvert$ in a run held (mm)	24.66	23.07	24.28
largest coil current (kA)	6.00	6.00	6.00
longest at the coil current limit (ms)	22.4	22.4	20.7
longest at $\pm V_{max}$, run held (ms)	380.4	380.4	365.8
rejected MPC decisions	—	—	0

The ten scenarios and the challenge as shipped, the challenge with its suggested gains, each disturbance and a combined sequence, all 13 parameters at both ends of their ranges, elongations 1.5 and 2.02, 1 mm of noise, setpoint steps of 20 mm vertically and 30 mm radially, case 10 at $\kappa = 2.02$ and case 7 at 15 V.

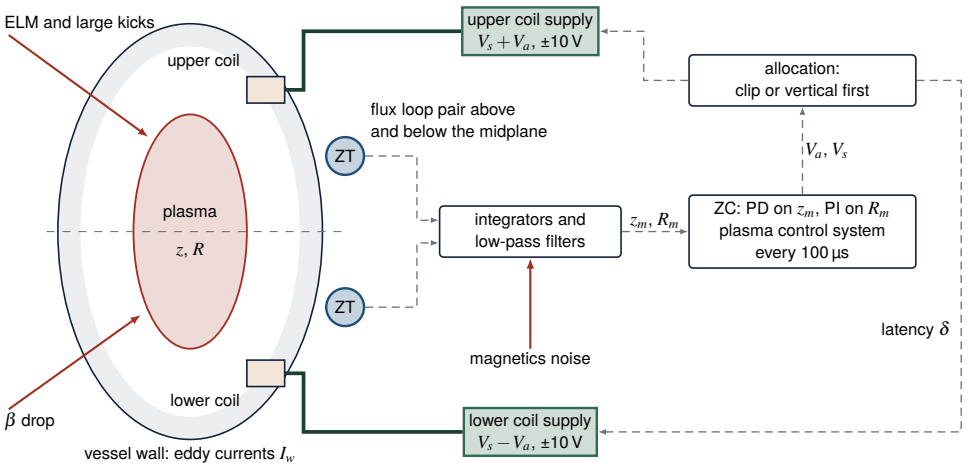

Figure 12.5 The loop of Figure 12.1 with the hardware of a tokamak. A pair of flux loops above and below the midplane (ZT) feeds integrators and filters, and the plasma control system ZC computes the position and velocity law and the radial PI every 100 µs. The allocation splits the two demands between the upper and lower coil supplies, fast converters limited to ±10 V, after the latency δ. Thick lines carry coil power, thin dashed arrows signals and solid arrows disturbances.

In the Chrome browser, on the application's own frame loop at 1×, the top speed of 10 000 decisions a second, the whole controller execution takes 1.1 to 1.3 % of one core under the lookup, 8.1 to 8.2 % under the online solve and 0.6 to 0.7 % under the loop, whose execution also runs the MPC's estimator. The host carried other computations while these times were measured, and the online figure lies within a factor of two of a tenth of a core.

12.6 Technology in Practice

The magnetics ZT, the controller ZC and the coil voltages of Figure 12.1 stand, on a tokamak, for magnetic sensors on the vessel, a real-time plasma control computer and a pair of coil power supplies. Figure 12.5 draws the loop with that hardware in place of the schematic's symbols.

What Measures It

No instrument touches the plasma. Its position is inferred from magnetic sensors and from the currents in the active coils. A learned controller on one research tokamak read 34 flux loops, 38 magnetic probes and 19 coil currents at 10 kHz, through a filtering stage modeled as a time delay with Gaussian noise [18]. On another machine the position from flux loop pairs fed the proportional term of a PID controller and their voltages the derivative term, and the integral term was never used

in its first operation [19], as in the chapter's default. The chapter's ZT is the true position plus white noise, and its latency δ stands for every delay from plasma to coil.

> **TECHNOLOGY** **Sensor: Flux Loops and Magnetic Probes**
>
> A tokamak never sees its plasma directly. Wire loops and small coils on the vessel pick up the changing magnetic field as an induced voltage. Pairs of flux loops above and below the midplane give the position: their flux difference is zero for a centered plasma and changes sign with the direction of a shift. The flux is the time integral of the measured voltage, and the voltage itself gives the velocity [19]. Integration brings drift. A repeatable linear drift is subtracted as a known slope and offset, while drift that does not repeat looks like slow plasma motion [20]. The coils' own fields must be corrected for, and one estimator run without those corrections erred by 1 cm or more in 38 % of its samples [19]. The chapter's ZT adds white noise only.

What Moves It

Coils whose radial field pushes the plasma up or down act together with eddy currents in the passive structure, and a 2010 design for ITER adds a pair of coils inside the vessel [12]. On one machine, adding two inboard coils to an outboard set roughly doubled the recoverable displacement. That displacement grows in proportion to the saturation voltage and falls roughly as the inverse of the growth rate, and a supply fast against the growth time makes its own dynamics unimportant [11]. Equation (12.5) carries the same two dependences for the chapter's two coils.

> **TECHNOLOGY** **Final Element: The Coil Power Supply**
>
> Fast power converters drive the coils. The thyristor supplies of one research tokamak were modeled for control design as a fixed bias and a fixed time delay identified from data. When a coil current changes polarity under too low a voltage request, such a supply can stick at zero output current for an extended time [18]. The supply planned for ITER's in-vessel vertical coils is a high-power chopper with energy storage and a 1 ms response. Two units are interleaved in series with the upper and lower coils, which are connected in anti-series, and every coil turn is brought out of the vessel so that a faulted turn can be left out of the circuit [12]. The chapter's supply is a ±10 V clamp behind a 100 µs delay line.

Where the Controller Runs

A plasma control system runs the loop on real-time computers. On the research tokamak above a compiled control policy ran at 10 kHz, the 100 µs period of the chapter's controller, after a conventional controller had formed the plasma [18]. ITER's plasma control system follows a pulse schedule, a supervisory layer changes the control in real time on plasma or plant events, and a central interlock system backs it up with simpler, robust algorithms. Every algorithm is first tested on a simulation platform [21]. The chapter's loop is one PID block in such a system, with the allocation rule between it and the supplies.

> **IN PRACTICE** **Size the Region Against the Noise**
>
> Vertical control is sized against noise. On operating tokamaks the standard deviation of the measured vertical position typically runs at 0.5 to 1 % of the minor radius, from power supplies, instruments, aliasing, crosstalk and plasma activity. Control turned marginal when the maximum controllable displacement was only 5 to 8 times that noise, and was lost near twice it [11]. From such statistics ITER took the empirical rule that the controllable displacement exceed a tenth of the minor radius, and operating machines keep it two to three times the largest expected perturbation [22].

What Is Modernizing It

Reinforcement learning has produced one controller that commands all of a tokamak's control coils. Trained on a simulator and run without further tuning, it held an elongation of 1.9 with a vertical growth rate of 1.4 kHz [18]. Data-driven estimators fit many signals at once: a linear model of 140

real-time magnetics signals reached a coefficient of determination of 0.906 on equilibria from 727 discharges [19]. And control systems are being designed to forecast approaching stability limits with physics-based and machine-learning predictors, so that a loss of vertical control can be mitigated [21]. The chapter's model observer is the simplest model-based estimate of this kind.

12.7 Limitations

The rigid massless plasma omits deformation, which matters at high elongation and for the definition of the margin [4].

One wall mode stands in for the vessel. A multi-mode vessel adds delay between the coil and the plasma and lowers the achievable gain ceiling beyond what latency alone produces [9].

> ⚠ **OPEN CHALLENGE** **Calculated Margins That Miss the Measurement**
>
> Models of the vessel and coils still predict the recoverable displacement poorly. On one tokamak the measured value was 0.15 to 0.24 m against about 0.40 m calculated, a gap far larger than the noise, attributed to complex non-axisymmetric passive structures and to the plasma striking the first wall. On two others the measurements agreed reasonably with the calculations but with significant variability, which reinforced the need for margin in ITER's design. Progress would be an understanding of non-axisymmetric structures and wall contact good enough to predict the measured value [11].

The margin law and the radial response are simple fits rather than equilibrium calculations, so absolute numbers carry less meaning than trends.

The magnetics noise is additive. On real machines the measurement also picks up the control coils, which couples the actuator into the sensor.

The disruption at the wall ends the simulation. Halo currents and the forces of the vertical displacement event are outside the model [5].

> ⚠ **OPEN CHALLENGE** **Seeing a Loss of Control Coming**
>
> ITER's design allows only 1 to 10 unmitigated vertical displacement events at the highest current and stored energy, a much stricter limit than the one on disruptions [22]. Its control system is to forecast approaching stability limits from many measurements and real-time calculations, and to call for impurity injection when it predicts a loss of vertical control [21]. The hard part is knowing how close the plasma is to the controllability boundary while it changes. Progress would be reliable, fast tools that assess that proximity in real time from reconstructed equilibria and other measurements, with warning long enough to reduce the elongation before the injection [22].

The explicit MPC's model holds the declared constants, so a learner's parameter change, the margin drop of a minor disruption, a faster wall, a latency other than 100 µs and the radial dynamics, which enter only through the box, are absent from its predictions, and the coil current limit stays a plant protection outside the program. Its plan covers 0.3 ms, the three single-period moves chosen before any comparison was run, and that horizon is why it recovers smaller large kicks than the loop. Its position gain makes it noisier than the loop with the observer.

12.8 Conclusion

An elongated tokamak plasma is an unstable plant held against a floor, a ceiling and a region, and the simulator of this chapter measures all three. Released at elongation 1.8 a 0.1 mm displacement reaches the wall in 20.0 ms, and at 2.0 in 6.7 ms. The position gain that holds the plasma must exceed about 0.45 V/mm at 1.8, about 0.70 V/mm at 2.0 and about 1.03 V/mm with the wall time halved, while with 300 µs of latency 10 V/mm settles into a 1.45 mm limit cycle. A model observer that reads the wall current from the force balance cuts the coil energy under 0.3 mm of noise from 33.9 J to 9.1 J against a filtered differentiator. A 20 mm kick at elongation 1.9 is lost at the default gains and

sooner at five times the position gain, and held with a 15 V supply or at elongation 1.8, and when the radial loop saturates the supply, priority allocation keeps a plasma that clipping loses 23 ms after the β drop. In the challenge that allocation alone meets the target at the default seed, while the suggested observer settings also pass at all five additional seeds tested. An explicit MPC with the voltage box in a 0.3 ms plan settles a 3 mm ELM in 0.52 ms against 1.32 ms but recovers kicks only up to 19.63 mm at elongation 1.9 against 19.91 mm for the loop, because it reverses the voltage late.

The model is rigid and massless with one wall mode and additive noise. A fuller model would add deformable plasma response, several wall modes, coil pickup in the magnetics and halo currents after loss. Within those limits the app lets a reader hold an unstable plant against hard limits and see which of them a tuning can move.

KEY POINTS

- An elongated plasma falls into the wall within milliseconds unless the coils push it back.
- The position gain must exceed a floor, about 0.45 V/mm at elongation 1.8, and the latency sets a ceiling above it.
- Gain decides how the loop behaves inside the recoverable region, while the growth rate and the supply voltage decide how large the region is.
- When the radial loop saturates the shared supply, priority allocation keeps a plasma that clipping loses.
- An explicit MPC with the voltage box in a 0.3 ms plan settles small kicks fastest but recovers smaller large kicks than the loop.

Where else this applies. An unstable plant held between a gain floor, a gain ceiling and a recoverable region is also the suspended ball of Volume 7, Chapter 6 and the balancing pole of Volume 7, Chapter 7.

References

[1] E. Schuster, M. Walker, D. Humphreys, and M. Krstić, “Plasma vertical stabilization with actuation constraints in the DIII-D tokamak,” *Automatica*, vol. 41, no. 7, pp. 1173–1179, 2005.

[2] J. W. Helton, K. J. McGown, and M. Walker, “Conditions for stabilization of the tokamak plasma vertical instability using only a massless plasma analysis,” *Automatica*, vol. 46, no. 11, pp. 1762–1772, 2010.

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

[4] A. Portone, “The stability margin of elongated plasmas,” *Nuclear Fusion*, vol. 45, no. 8, pp. 926–932, 2005.

[5] C. Clauser, S. Jardin, and N. Ferraro, “Vertical forces during vertical displacement events in an ITER plasma and the role of halo currents,” *Nuclear Fusion*, vol. 59, no. 12, p. 126037, 2019.

[6] M. Walker and D. Humphreys, “On feedback stabilization of the tokamak plasma vertical instability,” *Automatica*, vol. 45, no. 3, pp. 665–674, 2009.

[7] R. Stambaugh, L. Lao, and E. Lazarus, “Relation of vertical stability and aspect ratio in tokamaks,” *Nuclear Fusion*, vol. 32, no. 9, pp. 1642–1646, 1992.

[8] 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.

[9] R. Albanese and F. Villone, “The linearized CREATE-L plasma response model for the control of current, position and shape in tokamaks,” *Nuclear Fusion*, vol. 38, no. 5, pp. 723–738, 1998.

[10] G. De Tommasi, R. Albanese, G. Ambrosino, M. Ariola, P. J. Lomas, A. Pironti, F. Sartori, and J.-E. Contributors, “Current, Position, and Shape Control in Tokamaks,” *Fusion Science and Technology*, vol. 59, no. 3, pp. 486–498, 2011.

[11] D. Humphreys, T. Casper, N. Eidietis, M. Ferrara, D. Gates, I. Hutchinson, G. Jackson, E. Kolemen, J. Leuer, J. Lister, L. LoDestro, W. Meyer, L. Pearlstein, A. Portone, F. Sartori, M. Walker, A. Welander, and S. Wolfe, “Experimental vertical stability studies for ITER performance and design guidance,” *Nuclear Fusion*, vol. 49, no. 11, p. 115003, 2009.

[12] C. Neumeyer *et al.*, “Design of the ITER in-vessel coils,” Princeton Plasma Physics Laboratory, Princeton, NJ, Tech. Rep. PPPL-4576, Nov. 2010. [Online]. Available: https://www.osti.gov/biblio/1001665

[13] C. Neumeyer, “Design of the ITER In-Vessel Coils,” *Fusion Science and Technology*, 2011.

[14] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[15] S. Gerkšič and G. de Tommasi, "Vertical stabilization of ITER plasma using explicit model predictive control," *Fusion Engineering and Design*, vol. 88, no. 6–8, pp. 1082–1086, 2013.

[16] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[17] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[18] J. Degrave, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese *et al.*, "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.

[19] M. S. Parsons, A. H. Rasmussen, I. G. Stewart, C. Hansen, S. A. Sabbagh, S. Munaretto, J. W. Berkery, and S. P. Gerhardt, "A data-driven approach to real-time vertical position estimation for NSTX-U vertical stability control," *Plasma Physics and Controlled Fusion*, vol. 68, no. 2, p. 025008, 2026.

[20] J. D. King, E. J. Strait, R. L. Boivin, D. Taussig, M. G. Watkins, J. M. Hanson *et al.*, "An upgrade of the magnetic diagnostic system of the DIII-D tokamak for non-axisymmetric measurements," *Review of Scientific Instruments*, vol. 85, no. 8, p. 083503, 2014.

[21] J. A. Snipes, R. Albanese, G. Ambrosino, R. Ambrosino *et al.*, "Overview of the preliminary design of the ITER plasma control system," *Nuclear Fusion*, vol. 57, no. 12, p. 125001, 2017.

[22] D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson *et al.*, "Novel aspects of plasma control in ITER," *Physics of Plasmas*, vol. 22, no. 2, p. 021806, 2015.

12.9 Exercises

12.1 Concept check · 1 min

In scenario 3 the position loop runs at $K_p = 0.3\,\mathrm{V/mm}$ with no velocity term at elongation 1.8. The plasma wanders for about a sixth of a second and then runs into the wall at 166.6 ms. Which explanation is right?

(a) The gain is below the floor, so the coil pushes back more weakly than the plasma runs away

(b) The gain is above the ceiling, so the latency turns the loop into a slowly growing oscillation

(c) The magnetics noise winds up the integral until the supply saturates and the loop lets go

(d) The supply is too small, so the first burst of noise carries the plasma out of the region

12.2 Concept check · 2 min

In scenario 5 the latency is 300 µs and the loop runs at $K_p = 10\,\mathrm{V/mm}$ with no velocity term. The position settles into a limit cycle of 1.45 mm rms with the supply saturated 64 % of the time. With the latency set back to 100 µs the same gain holds the plasma to 0.059 mm. What sets this behavior?

(a) The filtered differentiator turns the magnetics noise into voltage that saturates the supply

(b) The latency removes phase near the crossover that the high gain reaches, so the loop is above its ceiling

(c) The longer latency lifts the gain floor to near 10 V/mm, so the plasma is only barely held

(d) The longer latency shrinks the recoverable region, so each noise sample leaves it for a moment

12.3 Concept check · 2 min

In scenario 11 the explicit MPC brings a 3 mm ELM back within 0.5 mm in 0.42 ms, against 1.02 ms for the loop with the differentiator. Yet it loses the 19.7 mm kick at 168.0 ms, while the loop holds it. Which explanation does the chapter support?

(a) The table has no region for so large a state, so the loop takes over too late to save it

(b) Its position gain of 11.5 V/mm passes the noise to the coils and wastes the supply

(c) Its 0.3 ms plan prices the later motion as if any voltage were available, so it reverses late

(d) Under clipping its box is $V_{max} - |V_s|/2$, so it has less voltage than the loop after the kick

12.4 Theory and calculation · 6 min

The wall time constant is $\tau_w = 5.0\,\text{ms}$, and the margin follows $m(\kappa) = \max[0.02,\, 1.2(2.05 - \kappa)]$, as given in Section 12.3.

(a) Find the margin and the wall-only growth rate of Equation (12.4) at $\kappa = 1.8$ and at $\kappa = 2.0$. *Checkpoint:* $m = 0.30$ and $667\,\text{s}^{-1}$. $m = 0.06$ and $3333\,\text{s}^{-1}$.

(b) The app shows 356 and $1070\,\text{s}^{-1}$ at these elongations with the coil voltage held. Find the ratio of the rate at 2.0 to the rate at 1.8 for the wall alone and for the app. *Checkpoint:* 5.0 for the wall alone, 3.0 for the app.

(c) Treat the growth as a single exponential at the app's rate. Estimate the time for a 0.1 mm displacement to reach the wall 150 mm away at each elongation, and compare with case 2 in Table 12.4. *Checkpoint:* 20.5 ms and 6.83 ms, against 20.0 ms and 6.7 ms measured.

12.5 Theory and calculation · 5 min

Scenario 9 moves the vertical setpoint to 10 mm at $\kappa = 1.8$, where $m = 0.30$. Use $g_w = 1.0\,\text{N/A}$, $g_c = 0.8\,\text{N/A}$, $L_w = 10\,\mu\text{H}$ and $R_c = 1.0\,\text{m}\Omega$ from Table 12.2, and take the plasma at rest.

(a) Find k_d, and the steady coil current and voltage that hold the plasma at $z = 10\,\text{mm}$. *Checkpoint:* 7.69×10^4 N/m, –0.96 kA and –0.96 V.

(b) With the integral off and $K_p = 2.0\,\text{V/mm}$, find where the plasma settles. *Checkpoint:* 10.51 mm, 0.51 mm above the setpoint.

(c) Find the position gain that would bring this offset down to 0.1 mm with no integral action. *Checkpoint:* 9.71 V/mm.

12.6 Theory and calculation · 6 min

Each coil supply is limited to ±10 V, and the coils receive $V_u = V_s + V_a$ and $V_l = V_s - V_a$, so that $V_a = (V_u - V_l)/2$. After a β drop the radial loop demands $V_s = +10\,\text{V}$, its own limit, and the vertical loop demands $V_a = -6\,\text{V}$.

(a) Find the two coil voltages and the delivered V_a and V_s when each coil is clipped on its own (Equation (12.8)). *Checkpoint:* $V_u = 4\,\text{V}$, $V_l = 10\,\text{V}$, $V_a = -3\,\text{V}$, $V_s = 7\,\text{V}$.

(b) Repeat with vertical priority. *Checkpoint:* $V_a = -6\,\text{V}$, $V_s = 4\,\text{V}$, $V_u = -2\,\text{V}$, $V_l = 10\,\text{V}$.

(c) With $V_s = +10\,\text{V}$ and every command inside ±10 V, find the largest $|V_a|$ that clipping can deliver and the command that delivers it. Compare with the box the explicit MPC is given under clipping. *Checkpoint:* 5 V, from a command of –10 V. The box $V_{max} - |V_s|/2$ is also 5 V.

12.7 App challenge · 8 min

Load scenario 3, Too little gain, at the default seed. It runs at $K_p = 0.3\,\text{V/mm}$ with no velocity term at elongation 1.8. Read the Plasma tile after each run. Where a value at 0.2 s is asked for, export the CSV file and read the column `metric_max_abs_z_mm` or `metric_rms_z_mm` in the row at `time_s` 0.2.

(a) Run the case as loaded until the plasma is lost. *Checkpoint:* Lost at 166.6 ms, within 0.5 ms.

(b) Press the reload button, set Position gain to 0.40 before Run, and run to 0.6 s. Read the largest excursion at 0.2 s and the Plasma tile at the end. *Checkpoint:* 0.48 mm at 0.2 s, within 0.02 mm. Lost at 482.4 ms, within 1 ms.

(c) Reload, set Position gain to 0.50 before Run, and run to 0.2 s. Read the rms position at 0.2 s. *Checkpoint:* 0.023 mm, within 0.002 mm, held.

Explain why the run at 0.40 looks healthy at 0.2 s, using the floor of about 0.45 V/mm.

12.8 App challenge · 8 min

Load scenario 6, Velocity from a noisy position, at the default seed. The noise is 0.3 mm and the velocity gain 2 V per m/s on the filtered differentiator. Make three separate runs of 0.35 s, pressing the reload button before each, and export the CSV file of each. Read the columns `metric_rms_z_mm` and `metric_coil_energy_J` in the row at `time_s` 0.35.

(a) Run the case as loaded. *Checkpoint:* 0.216 mm and 33.9 J.

(b) Set Velocity estimate to Model observer before Run. *Checkpoint:* 0.116 mm and 9.1 J.

(c) Set Velocity gain to 0 before Run. *Checkpoint:* 0.104 mm and 6.1 J.

Give each value within 0.003 mm and 0.3 J. Which velocity estimate would you keep on this machine, and why not drop the velocity term altogether?

12.9 App challenge · 9 min

Load the graded challenge, listed last as Challenge, a bad afternoon at κ 1.95, at the default seed. It starts at $K_p = 0.8\,\mathrm{V/mm}$ with a differentiator velocity gain of 3 and clipped allocation. For each run read the Plasma tile and the challenge verdict, and press the reload button before the next run.

(a) Run the case as loaded for 1.2 s. *Checkpoint:* Lost at 661.0 ms, within 0.5 ms.

(b) Before Run set Position gain to 2, Velocity estimate to Model observer and Velocity gain to 1.5, keep Clip each coil, and run again. *Checkpoint:* Lost at 661.2 ms, within 0.5 ms.

(c) Before Run change only Coil voltage allocation to Vertical first, and run the full 1.2 s. *Checkpoint:* Held. Vertical 3.94 mm and radial 0.80 mm over the last 0.4 s, within 0.02 mm: a pass.

Say which of the three threats of the challenge decides it.

12.10 Application area · 4 min

In scenario 7 a 20 mm kick at $\kappa = 1.9$ leaves a recoverable margin of 7 % on the 10 V supply. The margin is $1 - |\xi|/\xi_{max}$ with ξ_{max} from Equation (12.5).

(a) Predict the margin after the same kick with a 15 V supply. *Checkpoint:* 38 %.

(b) Find the supply voltage that would leave 50 % margin after this kick, and name one other rating that a larger supply does not provide by itself. *Checkpoint:* 18.6 V. The coil current limit, among others.

12.11 Application area · 3 min

In scenario 10 a β drop saturates the radial loop and a 10 mm vertical kick follows 5 ms later at $\kappa = 1.9$. With each coil clipped on its own the plasma is lost at 72.7 ms, with the radial position recovered only to 4.9 mm inward. With vertical priority it is held, with a peak of 10.7 mm. Why do real machines serve the vertical loop first, and what does the recoverable margin, which reads 53 % at the kick in both runs, fail to tell the operator?

12.12 Review the agent · 5 min

An AI assistant reviews scenario 7 and writes: “At $\kappa = 1.9$ the 20 mm kick lands outside the recoverable region. The plasma is lost at 67.9 ms, and five times the position gain only loses it sooner, at 58.9 ms. No tuning can save it. The only fixes are a 15 V supply or a lower elongation, so the larger supply should be ordered.” Find the flaw and say how you would check it.

CHAPTER 13

Stellarator Heat Exhaust Simulator in Which One Impurity Both Protects and Poisons

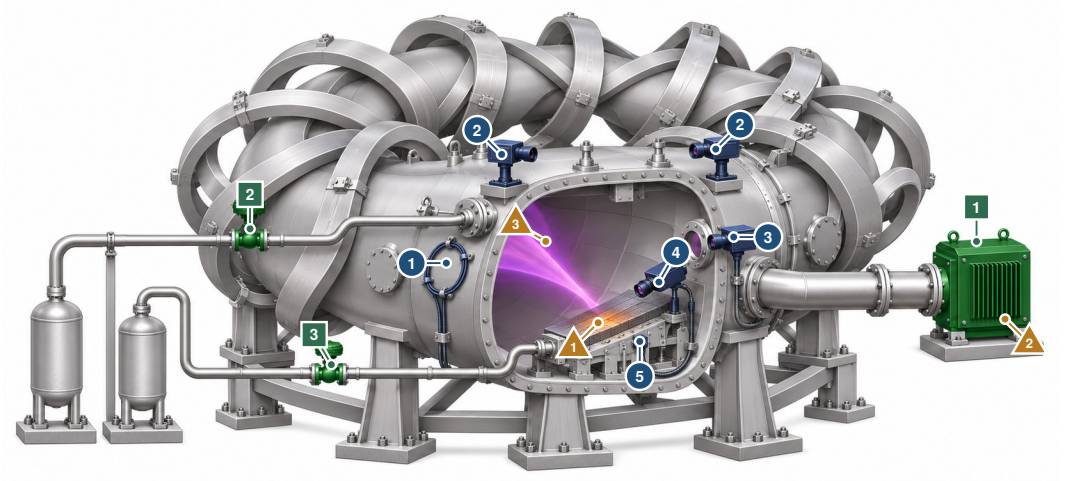

The controller moves 1 microwave heating 2 hydrogen fueling valve 3 neon seeding valve at the divertor

The controller reads 1 magnetic loop for the stored energy 2 interferometer for the density 3 bolometer for the radiated power 4 infrared camera on the target 5 thermocouple in the target

What disturbs it 1 loss of target cooling 2 heating sources trip 3 unannounced drop in confinement

Notice that the neon line ends beside the divertor target, so the gas that protects the warm strip enters at the edge, far from the core it can poison. Illustration, not to scale.

IN PLAIN TERMS **One Gas That Protects and Poisons**

The main idea. A stellarator is a fusion machine that holds very hot gas in a twisted magnetic cage for minutes. Heat leaking out lands on a few target plates, so a little neon is fed into the edge, where it gives most of that heat away as light.

The challenge. If too much neon reaches the center, the gas loses more heat than the heaters supply, and the run ends. Neon builds up slowly, so a controller that waits for the stored heat to drop acts too late.

Why it matters. The heat the machine can safely hold is set by what the plates and the neon can carry away.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 13.2, the fact boxes of Sections 13.1 and 13.6, the try-it boxes of Sections 13.4 and 13.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 13.6, Section 13.7, Section 13.8 and the open challenges

The control question A stellarator holds its plasma for minutes, and neon seeded into its edge protects the divertor targets from the exhaust but poisons the core if too much arrives there. How should three slow actuators share the limits of a plant that is stable but constrained?

Learning objectives

- Compute what more stored energy costs in heating under the ISS04 scaling.
- Compute the time constants of the neon inventories and the target element, and explain why the heat flux answers seeding slowly.
- Explain how a stored-energy loop hides an overseeding collapse, and how a core-radiation override on a low selector prevents it.
- Recognize an infeasible target after a source trip, and show how anti-windup on the available capacity removes the overshoot.
- Compare protected PI with an offset-free MPC on the heat flux and the neon each one uses.

Prerequisites Sections 1.4.6 (heat exhaust, detachment, radiative collapse and confinement scaling), 1.6.2 (velocity form) and 1.6.4 (running a case). Section 1.6.3 is needed only for the MPC. Chapter 12 is the fast, unstable counterpart.

The app https://apmonitor.com/apps/stellarator-control/

See also Volume 4, Chapter 3 for limiters that vote on one fuel valve with every integrator tracking the fuel commanded. Volume 4, Chapter 2 for a protective override that rests on its clamp in healthy operation. Volume 4, Chapter 6 for a vapor limit, not the product, that bounds the line speed. Volume 9, Chapter 4 for an offset-free MPC whose estimator carries the bias of a fuel no controller measures.



13.1 Introduction

A plant that is stable, slow and amply actuated looks easy to control until its limits are counted. In chemical plants the usual answer is a set of selectors, each giving an actuator to the constraint nearest its limit, and their systematic design is still an active topic [1]. A selector creates losing controllers, and a losing controller with integral action winds up unless it is told what the actuator is actually doing, which is why anti-windup and override design are treated together [2]. A jet engine's fuel valve, shared by a speed governor and its limiters, poses the same problem (Volume 4, Chapter 3). Model predictive control handles the same limits as constraints of one optimization, and an output-disturbance model makes it offset-free against a model that is wrong [3], as a boiler with an unmeasured fuel heating value shows (Volume 9, Chapter 4). Fired heaters, compressors (Volume 4, Chapter 2), reactors with a temperature limit and boilers with a drum-pressure limit all share this shape, as does the exhaust of a long-pulse fusion plasma.

A stellarator holds its plasma in a twisted magnetic field made entirely by external coils [4]. With no large plasma current there is nothing to drive and nothing to disrupt, and a stellarator can in principle run without stopping. Wendelstein 7-X (W7-X), a large optimized stellarator, has run discharges that injected 1.8 GJ under fully detached divertor conditions [5]. Over pulses that long, every watt of heating that the plasma does not radiate reaches the divertor targets, which in W7-X are built for 10 MW/m^2 in steady state [6]. Seeding an impurity such as neon into the edge radiates most of that power before it arrives, and the divertor plasma cools until it detaches [7]. The same neon that protects the targets dilutes and cools the core if it gets there, and a series of neon puffs in W7-X ended in exactly such a radiative collapse [7].

✦ DID YOU KNOW? Built for Half an Hour

Wendelstein 7-X is designed for discharges of up to 30 min and 18 GJ of heating energy, carried by superconducting coils and water-cooled plasma-facing components. Its operators raise the allowed heating energy per discharge in steps, from 1.3 GJ in one campaign to 2 GJ in the two that followed [5].

The simulator corrects the belief that a fusion plasma's performance is limited by the heating that can be bought. In this model, 20 % more stored energy costs about 60 % more heating, because the confinement time falls as the heating rises (Section 1.4.6), and every added megawatt leaves again as exhaust that the radiation loop and the targets must absorb. The binding limits are the target heat flux, the target surface temperature and the core radiation, and the controller's real task is to hold a narrow window between an attached divertor that overheats and a radiating layer that consumes the plasma. This makes the problem the opposite of the tokamak of Chapter 12, which stabilizes a fast unstable mode with one saturated actuator on a millisecond scale. Here three slow actuators (heating, hydrogen fueling and neon seeding) interact on a stable plant over minutes, and the lesson is coordination under constraints, protection logic and an infeasible target.

However, the teaching material on override control is drawn from process plants, and the fusion literature on radiated-power control is written for experts on a particular machine [8, 9]. This chapter works with a stellarator exhaust simulator, run in a web browser, in which the heating, the fueling and the seeding of a stellarator plasma are held by plain PI, by protected PI and by a constrained model predictive controller (MPC) on identical measurements, and it measures what each structure does against the exhaust limits. The two-zone balance, the ISS04 scaling [10], the two-point closure [11], the velocity-form PI law, the selectors, the bias-corrected MPC and the dual active-set quadratic program solver [12] are all established methods. What the simulator adds for study is an eight-state model whose energy, hydrogen and neon ledgers close to round-off and whose radiative collapse emerges from the balances without any scripted termination, a measurement packet that is the only input any controller reads, and every equation, constant and default that a reader needs to build the simulator again. Ten cases and a graded challenge each come with a rule that decides when they are done. The model is inspired by W7-X, but it is not a W7-X model, and no simulated result in this chapter is a W7-X result.

Section 13.2 traces the stellarator from its invention to the island divertor and radiated-power feedback. Section 13.3 sets out the two-zone power balance, the radiation of each zone, the radiating layer and detachment, the target element, the particle inventories and what the controllers measure. Section 13.4 describes the stable design point, the operating window and the unscripted collapse. Section 13.5 gives the PI loops, protected PI and the MPC. Section 13.6 measures the ten cases, the challenge and the paired comparison of the three structures, and Section 13.7 sets the loops beside the hardware of a real stellarator. Section 13.8 states what the model leaves out, and Section 13.9 summarizes the results.

13.2 Stellarators, Island Divertors and Radiated Power

The stellarator was conceived at Princeton's Project Matterhorn and described in full in 1958. A strong solenoidal field around a toroidal tube, given a rotational transform by a figure-eight tube or by additional helical windings, makes each field line trace out a closed toroidal surface [4]. Early stellarators were limited in plasma pressure and suffered high neoclassical transport and poor confinement of fast particles [13], and most fusion programs turned to the tokamak, whose transform comes from a current in the plasma itself. The current is also what a tokamak must drive, sustain and protect against disruption, which is why its vertical position needs the fast feedback of Chapter 12.

Optimization brought the stellarator back. Numerical design of the magnetic configuration reduced the neoclassical transport, raised the attainable pressure and improved the confinement of fast particles, and modular non-planar coils made such a configuration buildable [13]. W7-X, with its island divertor, is the result of that program, and the first history box below tells how it came into operation.

> ⧖ HISTORY **An Optimized Stellarator Comes On Line (2015 to 2018)**
>
> Wendelstein 7-X made its first plasma on 10 December 2015, after more than 15 years of construction. Its field comes from 50 non-planar and 20 planar superconducting coils, whose shape was found with coupled numerical codes for equilibrium, stability, transport, particle orbits and divertor loads. Its first plasmas, with five inboard limiters and no divertor, were reported in 2017 [14]. Its first operation with the island divertor followed, on ten identical uncooled graphite test divertor units, with the heating raised from six to ten gyrotrons [15]. The island divertor uses a chain of magnetic islands at the plasma boundary to guide the exhaust onto ten target units, and its stated purpose is steady-state exhaust for discharges of up to thirty minutes [6].

Radiative exhaust by impurity seeding was developed on tokamaks first. Nitrogen, neon and argon seeded into the divertor radiate power that would otherwise land on the strike lines, and the seeding rate is adjusted by feedback on a measure of the divertor state. Its extrapolation from present devices through ITER to a demonstration reactor was reviewed in 2013 [8]. Reduced models of the divertor, the two-point model above all, relate the target temperature and density to the power and density upstream and locate detachment below about 10 eV at the target [11]. The second history box below tells what the first seeding of W7-X's island divertor did.

> ⧖ HISTORY **Four Neon Puffs and a Radiative Collapse (2019 to 2020)**
>
> Impurity seeding in an island divertor was first shown on Wendelstein 7-X. Neon and nitrogen raised the radiated fraction to 80 % and more and cut the divertor heat loads by two thirds, at a price in stored energy. Neon went on radiating for tens of seconds after each puff ended, because it recycles from the carbon wall. In one discharge neon was puffed for 200 ms at 5, 7, 9 and 11 s. Each puff moved the radiating zone further inward, and after the fourth, about 8×10^{19} atoms in all, the radiated power approached the heating and a radiative collapse ended the plasma [7]. Detachment by density alone was later held stable with efficient particle exhaust: the peak divertor heat flux fell from above 5 to below 0.5 MW/m^2 as the radiated fraction rose from 0.35 to above 0.8 [6].

Feedback control of the radiated power is the current answer. A controller designed from identified frequency responses of the bolometer signal to nitrogen and neon injection now holds chosen radiated-power levels in W7-X, attached or detached, and kept the divertor detached through discharges longer than 300 s [9]. The energy confinement against which all of this is measured follows the ISS04 scaling, fitted across the international stellarator database with a renormalization factor for each configuration [10]. The density at which a stellarator plasma collapses radiatively scales with the square root of the heating power [16]. The coronal cooling rates that set how much an impurity radiates come from atomic data, and the public Open-ADAS effective coefficients rest on a generalized collisional–radiative model [17].

13.3 The Plasma and Exhaust Model

The simulated device has the plasma volume and field of W7-X, 30 m^3 and 2.5 T [15], a major radius of 5.5 m and a minor radius of 0.53 m, with five field periods and ten island divertor units, as shown in Figure 13.1. Ten microwave sources of 0.8 MW heat a hydrogen plasma at the second harmonic. A gas valve fuels it, and a neon inlet beside one divertor unit seeds the edge. The magnetic configuration is a preset chosen before the run (Standard, High iota or High mirror), and the coils are not actuators. Inside the boundary are the confined plasma as two zones, a core inside $r/a = 0.8$ and an edge shell that includes the island scrape-off layer (SOL), the hydrogen and neon inventories, the radiating layer and divertor plasma of the SOL, the bulk temperature of the hottest target element, the actuators and the diagnostics. Plasma current, magnetohydrodynamics, neutral beams, pellets, profile shapes, turbulence and fusion reactions are outside it.

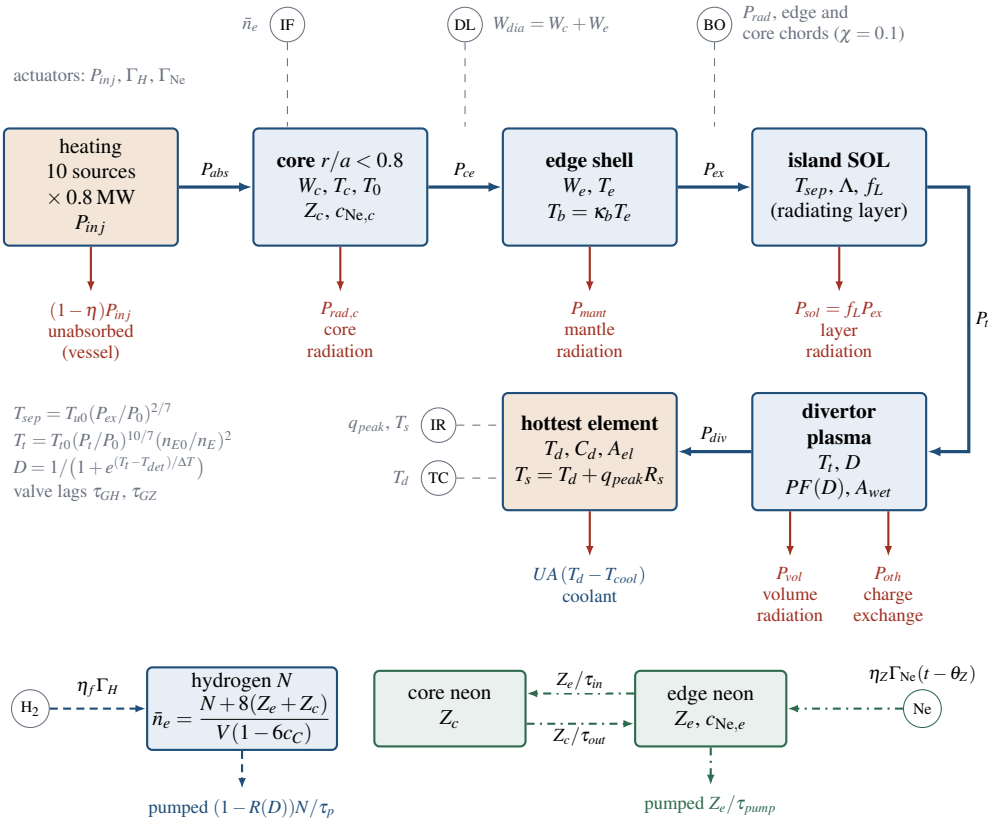

Figure 13.1 Power ledger, neon path and diagnostics of the simulated stellarator. Injected heating P_{inj} is absorbed as P_{abs} in the core, conducted as P_{ce} to the edge shell, leaves the shell as P_{ex} into the island SOL, crosses the radiating layer as P_t to the divertor plasma, and reaches the targets as P_{div} with peak heat flux q_{peak} on the hottest element. The short solid arrows below the zones (red) are the ledger's destinations, each nonnegative. Neon (dash-dotted arrows) enters near one divertor unit after the inlet delay θ_Z, exchanges between the edge and the core inventories and is pumped from the edge, and hydrogen (densely dashed arrows) is fueled and pumped. The diagnostics, joined by thin dashed taps, are the diamagnetic loop (DL), the interferometer (IF), the bolometers (BO), the infrared camera (IR) and a thermocouple (TC). After the app's Exhaust view, not to scale.

Table 13.1 classifies the variables. Eight states are integrated: six physical (the hydrogen ion inventory N, the core and shell thermal energies W_c and W_e, the edge and core neon inventories Z_e and Z_c, and the hottest element's bulk temperature T_d) and two actuator lags (the fueling and seeding flows that reach the plasma, G_H and G_Z). The heating slew, the seeding inlet delay and the instrument filters are discrete states of the actuator and sensor models. The model works in SI units with plasma temperatures in eV.

13.3.1 Two Zones and the Power Balance

The two zone energies follow the power balance of Equations (13.1) and (13.2), in which the core absorbs the heating, radiates, and conducts the rest to the shell, while the shell radiates its own share

Table 13.1 Classification of the Variables and Their Case 1 Design Values

Variable	Role	Design	Range
W	CV, stored energy	0.600 MJ	SP 0.2–0.9 MJ
$\bar{n}_e$	CV, line-averaged density	6.00×10^{19} m^{-3}	SP 2–12 $\times 10^{19}$
P_{rad} or f_{rad}	CV, radiation	2.41 MW, 48 %	SP 0–7.5 MW
q_{peak}	CV (cascade), constraint	4.50 MW/m^2	limit 10
T_s	constraint, target surface	480 °C	limit 1100
P_{inj}	MV, heating	5.03 MW	0–8, 4 MW/s
Γ_H	MV, hydrogen fueling	2.17×10^{21} s^{-1}	0–5 $\times 10^{21}$
Γ_{Ne}	MV, neon seeding	0.034×10^{18} s^{-1}	0–2 $\times 10^{19}$
f_{ren}	DV (hidden), confinement	0.65	$\times$0.5–1.2
UA	DV, target cooling	1000 W/K	20–100 %
n_{src}	DV, sources available	10	trip of 2
W_c, W_e	states, zone energies	0.529, 0.071 MJ	—
N	state, hydrogen ions	1.63×10^{21}	—
Z_e, Z_c	states, neon	8.6×10^{17}, 6.4×10^{17}	≥ 0
T_d	state, element bulk	255 °C	—
T_t, D	algebraic, divertor	21.4 eV, 0.000	$D \in [0,1]$

and passes the remainder into the island SOL,

$$\frac{\mathrm{d}W_c}{\mathrm{d}t} = P_{abs} - P_{ce} - P_{rad,c}, \tag{13.1}$$

$$\frac{\mathrm{d}W_e}{\mathrm{d}t} = P_{ce} - P_{mant} - P_{ex}, \tag{13.2}$$

where $P_{abs} = \eta(\bar{n}_e)P_{inj}$ is the absorbed heating, P_{ce} the power from core to shell, $P_{rad,c}$ the core radiation, P_{mant} the radiation of the shell (the mantle) and P_{ex} the exhaust into the SOL. The unabsorbed power $(1-\eta)P_{inj}$ goes to the vessel and is booked as its own destination. Second-harmonic X-mode heating at 140 GHz is absorbed only below its cutoff density $n_{cut} = 1.2 \times 10^{20}\,\mathrm{m}^{-3}$, and the absorption is the smooth step of Equation (13.3),

$$\eta(\bar{n}_e) = 0.10 + \frac{0.85}{1 + \exp\left[(\bar{n}_e/n_{cut} - 0.93)/0.025\right]}, \tag{13.3}$$

which gives 0.950 at the design density and 0.149 at the cutoff.

The temperatures follow from the energies, $T_c = W_c/(3\bar{n}_e V_c e)$ and $T_e = W_e/(3n_E V_e e)$, with the core volume $V_c = 0.64V$, the shell volume $V_e = 0.36V$, the edge density $n_E = 0.7\bar{n}_e$ and the elementary charge e. The temperature falls linearly across each zone, in the core from $T_0 = \max(2T_c - T_b, T_c)$ on the axis to $T_b = \kappa_b T_e$ at $r/a = 0.8$, and in the shell from T_b to the separatrix temperature T_{sep}, with the profile factor $\kappa_b = 1.6$. The core rests on its boundary. It loses $P_{ce} = (W_c - W_b)/\tau_c$ with $W_b = 3e\bar{n}_e V_c T_b$, so a shell that cools pulls the core down with it. The shell passes $P_{ex} = W_e/\tau_e$ into the SOL.

Confinement follows ISS04 times the configuration factor f_{ren}, as shown in Equation (13.4),

$$\tau_E = f_{ren}\, 0.134\, a^{2.28} R^{0.64} P^{-0.61} \bar{n}_{19}^{0.54} B^{0.84} \iota_{2/3}^{0.41}, \tag{13.4}$$

with P in MW, $\bar{n}_{19}$ in $10^{19}\,\mathrm{m}^{-3}$, a and R in m, B in T and the rotational transform $\iota_{2/3} = 0.9$ [10]. The power degradation is read from the stored energy, not from the heating. Inverting $W = \tau_E(P)P = C(\bar{n}_e)P^{0.39}$, the scaling of Equation (1.12), gives the loss power $P_W = (W/C)^{1/0.39}$ that the present energy implies, and $\tau_E = W/P_W$. A heating step therefore changes the losses only as the energy changes. The split is $\tau_c = c_1\tau_E$ and $\tau_e = c_2\tau_E$ with $c_2 = w_E = 0.12$ and $c_1 = 1 - w_E(1+k)$, $k = \kappa_b V_c/(0.7V_e) = 4.06$, so $c_1 = 0.392$. With no shell radiation, this split makes the steady stored energy exactly $\tau_E(P_{abs} - P_{rad,c})$. The factor f_{ren} is 0.65 for Standard, 0.75 for High iota and 0.58 for High mirror, teaching values, the Standard one chosen so that 5 MW holds 0.60 MJ at $6 \times 10^{19}\,\mathrm{m}^{-3}$. The

presets also change the wetted area and the peaking factors of Section 13.3.3, to 1.2 m^2 and 3.5→1.7 for High iota and 1.9 m^2 and 2.6→1.4 for High mirror.

13.3.2 Radiation in Each Zone

Each zone radiates from intrinsic carbon at a fixed fraction $c_C = 0.015$ of the electron density and from seeded neon, with the cooling rate averaged over the zone's linear temperature range, as shown in Equations (13.5) and (13.6),

$$P_{rad,c} = \bar{n}_e^2 V_c \Big[c_C \langle L_C \rangle + c_{\mathrm{Ne},c} \langle L_{\mathrm{Ne}} \rangle \Big]_{T_b}^{T_0} + P_{br}, \tag{13.5}$$

$$P_{mant} = n_E^2 V_e \Big[c_C \langle L_C \rangle + c_{\mathrm{Ne},e} \langle L_{\mathrm{Ne}} \rangle \Big]_{T_{sep}}^{T_b}, \tag{13.6}$$

where $\langle L \rangle_a^b = (b-a)^{-1} \int_a^b L dT$, $c_{\mathrm{Ne},c} = Z_c/(\bar{n}_e V_c)$ and $c_{\mathrm{Ne},e} = Z_e/(n_E V_e)$ are the neon fractions, and the hydrogen bremsstrahlung is $P_{br} = 5.35 \times 10^{-37}\, \bar{n}_e n_H \sqrt{T_c/1000}\, V_c$ in W with $n_H = \bar{n}_e(1 - 6c_C - 8c_{\mathrm{Ne},c})$. The cooling rates $L_{\mathrm{Ne}}(T)$ and $L_C(T)$, in W m^3, are two-bump analytic forms fitted to coronal rates computed from the Open-ADAS adf11 files at 10^{20} m^{-3} (https://open.adas.ac.uk/) [17]. The fits stay within a factor of 1.36 of the table for neon between 5 eV and 3 keV and within 1.61 for carbon between 2 eV and 3 keV. Neon radiates most near 30 eV (5.8×10^{-32} W m^3) and about fifty times less between 200 eV and 1 keV, which is why it radiates usefully in the cool edge and harmfully only once the radiating range reaches into the confined plasma.

13.3.3 Radiating Layer, Two-Point Closure and Detachment

The SOL receives P_{ex} at the separatrix temperature $T_{sep} = T_{u0}(P_{ex}/P_0)^{2/7}$ of the two-point model, with $T_{u0} = 110$ eV and $P_0 = 3$ MW [11]. A conduction-limited radiating layer takes the fraction f_L of it, as shown in Equation (13.7),

$$\Lambda = 10^{31} K_L \frac{n_{E,19}^2 T_{sep}^2 \big[c_C I_C + c_{\mathrm{Ne},e} I_{\mathrm{Ne}} \big]}{P_{ex,MW}^2},$$

$$f_L = f_{max}\big(1 - e^{-\Lambda/2 f_{max}}\big), \tag{13.7}$$

where $I(T_{sep}) = \int_0^{T_{sep}} \sqrt{T'} L dT'$ in W m^3 eV$^{3/2}$, $n_{E,19}$ is the edge density in 10^{19} m^{-3}, $P_{ex,MW}$ the exhaust in MW, $K_L = 3 \times 10^{-4}$ and $f_{max} = 0.85$. For small Λ the layer radiates $\Lambda/2$ of the power that enters it, the slope of the conduction-limited layer, and it saturates smoothly instead of reaching 100 % with infinite slope. The layer radiates $P_{sol} = f_L P_{ex}$ and passes $P_t = P_{ex} - P_{sol}$ to the divertor plasma, whose temperature and detachment degree follow Equation (13.8),

$$T_t = T_{t0}\left(\frac{P_t}{P_0}\right)^{10/7} \left(\frac{n_{E0}}{n_E}\right)^2, \qquad D = \frac{1}{1 + e^{(T_t - T_{det})/\Delta T}}, \tag{13.8}$$

with $T_{t0} = 30$ eV, $n_{E0} = 4.2 \times 10^{19}$ m^{-3}, $T_{det} = 5$ eV and $\Delta T = 1.5$ eV. Detachment acts three ways. The divertor plasma radiates $P_{vol} = D f_{vol} P_t$ with $f_{vol} = 0.5$ and loses $P_{oth} = (0.05 + 0.15D) P_t$ to charge exchange and the main-chamber wall, so the targets receive $P_{div} = P_t - P_{vol} - P_{oth}$. The peaking factor falls from 3.0 to 1.5, $PF = 3.0 - 1.5D$, so the heat spreads. The recycling falls from 0.95 to 0.92, so the pumping improves. The peak heat flux on the wetted area $A_{wet} = 1.5$ m^2 is $q_{peak} = PF\, P_{div}/A_{wet}$. Every destination is nonnegative because $f_L < 1$ and $f_{vol} + 0.05 + 0.15 < 1$.

13.3.4 Target Element, Particles and Terminations

The hottest target element of area $A_{el} = 0.05$ m^2 receives $q_{peak} A_{el}$, which is part of P_{div}, and its bulk temperature follows Equation (13.9),

$$C_d \frac{\mathrm{d}T_d}{\mathrm{d}t} = q_{peak} A_{el} - UA\,(T_d - T_{cool}), \quad T_s = T_d + q_{peak} R_s, \tag{13.9}$$

with $C_d = 10^4\,\mathrm{J/K}$, $UA = 1000\,\mathrm{W/K}$ (a 10 s time constant), $T_{cool} = 30\,^\circ\mathrm{C}$ and the surface resistance R_s of 50 K per MW/m^2. The infrared camera sees the surface T_s, the thermocouple the bulk.

The hydrogen and neon inventories follow Equations (13.10) to (13.12),

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \eta_f G_H - \big(1 - R(D)\big)\frac{N}{\tau_p}, \tag{13.10}$$

$$\frac{\mathrm{d}Z_e}{\mathrm{d}t} = \eta_Z G_Z - \frac{Z_e}{\tau_{pump}} - \frac{Z_e}{\tau_{in}} + \frac{Z_c}{\tau_{out}}, \tag{13.11}$$

$$\frac{\mathrm{d}Z_c}{\mathrm{d}t} = \frac{Z_e}{\tau_{in}} - \frac{Z_c}{\tau_{out}}, \tag{13.12}$$

with the valve lags $\tau_{GH}\dot{G}_H = \Gamma_H - G_H$ and $\tau_{GZ}\dot{G}_Z = \Gamma_{\mathrm{Ne}}(t - \theta_Z) - G_Z$, the recycling $R = 0.95 - 0.03D$, $\tau_p = 0.15\,\mathrm{s}$ (an effective particle time of 3 s), $\eta_f = 0.25$, $\eta_Z = 1$, $\tau_{pump} = 25\,\mathrm{s}$, $\tau_{in} = 8\,\mathrm{s}$, $\tau_{out} = 6\,\mathrm{s}$, $\theta_Z = 0.3\,\mathrm{s}$ and both lags 0.2 s. Quasi-neutrality with a stated average neon charge of 8 and fully stripped carbon gives the density $\bar{n}_e = (N + 8(Z_e + Z_c))/(V(1 - 6c_C))$ and $Z_{eff} = 1 + 30c_C + 56c_{\mathrm{Ne},c}$. Inventories never go negative, and a clamp, if one acts, is booked in the ledger.

Four terminations latch once, each with a historian mark and a named cause. A radiative collapse is declared when the radiation inside the separatrix, $P_{rad,c} + P_{mant}$, has exceeded P_{abs} for 0.3 s and W has fallen below half its largest value of the last ten seconds, or when T_c falls below 60 eV. The same test with $\eta < 0.5$ is a heating cutoff. A divertor interlock trips every source when T_s stays above 1200 °C for 1 s, and the end of the programmed pulse is a success. These tests only detect a collapse that the balances have already produced. Nothing in the derivative scripts it.

13.3.5 Parameters, Initialization and Numerics

Table 13.2 lists the parameters and their provenance. Every one is hidden from the learner except through the disturbance sliders, which scale f_{ren} (*teaching truth*, never told to a controller) and UA, add a delay to the diamagnetic signal or the bolometers, and scale the noise. A run starts lined out, by damped Newton iteration on the steady balances, at the operating point the case's controls define: each loop's setpoint when it is in Automatic and its manual station when it is in Manual, or the heat-flux setpoint when the cascade is on. The plant steps every 5 ms with fourth-order Runge–Kutta, and every ledger flow is accumulated with the same four stage weights, so the energy, hydrogen and neon books close to round-off. Two guards keep the zone energies nonnegative through a collapse. Every outflow of a zone (its radiation, and P_{ce} or P_{ex} leaving it) is multiplied by the gate $g(W) = x(2 - x)$ with $x = W/10\,\mathrm{kJ}$, which equals 1 at and above 10 kJ and falls to 0 at zero energy, so an emptied zone stops radiating. The least shell energy of any case that does not collapse is 0.0495 MJ (case 9 under the MPC), almost five times the gate. A step that would carry W_c or W_e below zero, or move either by more than half of $\max(W, 10\,\mathrm{kJ})$, is halved recursively, at most 14 times, each half adding to the ledger, so every held case integrates exactly as plain Runge–Kutta does.

13.3.6 What the Controllers Measure

Every controller receives the same packet every 20 ms: the diamagnetic stored energy $W_{dia} = W_c + W_e$ (2 ms lag), the interferometer density (1 ms lag), the bolometer total and an edge and a core chord (20 ms filter and a 20 ms base delay), the infrared peak heat flux and surface temperature (30 ms delay, saturating at 20 MW/m^2 and 1500 °C), the thermocouple (1 s lag), and the readbacks of the injected heating, the available source count and the two valve flows. The edge chord reads $P_{rad,e} + 0.1P_{rad,c}$ and the core chord $P_{rad,c} + 0.1P_{rad,e}$, a stated crosstalk $\chi = 0.1$. Each channel carries seeded Gaussian noise of standard deviation 4 kJ, $0.04 \times 10^{19}\,\mathrm{m^{-3}}$, 0.03 MW, 0.08 MW/m^2, 6 K and 1 K at the default level. The radiated fraction is measured radiation over measured heating, and the inferred average temperature $W_{dia}/(3\bar{n}_e V e)$ is labeled as an inference. The zone energies apart, the neon inventories, the absorbed fraction, f_{ren}, UA, D and the case's event schedule never enter the packet.

Table 13.2 Parameters, Standard Configuration

Symbol	Meaning	Value	Source
R, a, V, B	geometry and field	5.5, 0.53, 30, 2.5	W7-X [15]
ISS04	form and exponents	Eq. (13.4)	fitted [10]
f_{ren}	configuration factor	0.65 (0.75, 0.58)	teaching, calibrated
c_C	carbon fraction	0.015	teaching ($Z_{eff} \approx 1.5$ [6])
L_{Ne}, L_C	cooling rates	fits	Open-ADAS, $\times 1.36$, $\times 1.61$
κ_b, w_E	profile, shell share	1.6, 0.12	teaching
K_L, f_{max}	layer scale, saturation	3×10^{-4}, 0.85	teaching, calibrated
T_{u0}, T_{t0}	two-point prefactors	110, 30 eV	teaching
$T_{det}, \Delta T$	detachment	5, 1.5 eV	teaching
A_{wet}, PF	wetted area, peaking	1.5 m^2, 3.0→1.5	teaching
q_{lim}	heat-flux limit	10 $\mathrm{MW/m}^2$	published [6]
$T_{s,lim}$, trip	surface limit, interlock	1100, 1200 °C	teaching
A_{el}, C_d, UA	hottest element	0.05, 10^4, 1000	teaching
τ_p, R, η_f	hydrogen	0.15 s, 0.95→0.92, 0.25	teaching
$\tau_{in}, \tau_{out}, \tau_{pump}$	neon transport	8, 6, 25 s	teaching
θ_Z, lags	seeding inlet	0.3, 0.2 s	teaching
n_{cut}	X2 cutoff	1.2×10^{20} m^{-3}	physical
sources	heating	10×0.8 MW, 4 MW/s	teaching
valves	full scale	5×10^{21}, 2×10^{19} s^{-1}	teaching

13.3.7 Assumptions

The two zones keep the power balance, the loss of stored energy when the edge cools and the inward-moving radiating layer, and they omit transport equations, profile shapes and transport barriers. ISS04 is a cross-machine fit with a scatter of tens of percent, and the presets' factors are teaching values. Coronal cooling rates understate radiation at the low edge temperatures where the plasma is far from coronal equilibrium, so K_L is calibrated rather than derived: it makes the unseeded design radiate 28.4 % of its heating with carbon alone, against 0.35 measured attached [6]. The detachment closure is smooth and single-branched by construction, while real detachment can be abrupt, hysteretic and three-dimensional. One target element stands for ten units, so strike-line motion and asymmetries are absent. Neon is one species with two transport times and no neoclassical accumulation, and the absorption depends on density only, while real second-harmonic absorption also falls at low electron temperature.

13.4 Dynamic Character

13.4.1 A Stable Design Point With Slow Neon

At the design point of case 1, 5.03 MW is injected at 6.00×10^{19} m^{-3}, and the plasma stores 0.600 MJ (0.529 MJ in the core at an average of 955 eV and 0.071 MJ in the shell at 327 eV). Of the heating, 2.41 MW (48 %) is radiated, the cascade holds q_{peak} at 4.50 $\mathrm{MW/m}^2$, and the divertor plasma sits at 21.4 eV, attached. The ISS04 estimate typed independently from Equation (13.4) gives $\tau_E = 130.3$ ms and $\tau_E P_L = 0.6032$ MJ with $P_L = P_{abs} - P_{rad,c} = 4.629$ MW, 0.53 % above the model's 0.600 MJ, the difference being the shell's own radiation. The linearized eight-state plant, with the three actuators held, has the eigenvalues −0.0215, −0.100, −0.310, −0.334, −5.0, −5.0, −20.5 and −152.3 s^{-1}, all real and negative. The slowest, 46.5 s, belongs to the neon inventories, and the fastest, 6.6 ms, to the shell. The plant is stable, and its difficulty lies elsewhere. The neon that protects the target answers on a time scale about 350 times the energy confinement time.

With seeding off, the stored energy follows ISS04 times f_{ren} within 5.4 % over 2–8 MW and 3–10×10^{19} m^{-3}, the largest departure at 2 MW and 10×10^{19} m^{-3}, where carbon radiation in the shell is largest relative to the heating. The plasma holds up to 11.2–11.7 $\times 10^{19}$ m^{-3} at 3–8 MW before the absorption cutoff ends it, 1.3 to 2.0 times the Sudo density limit [16], which is an order-of-magnitude check and not a fit.

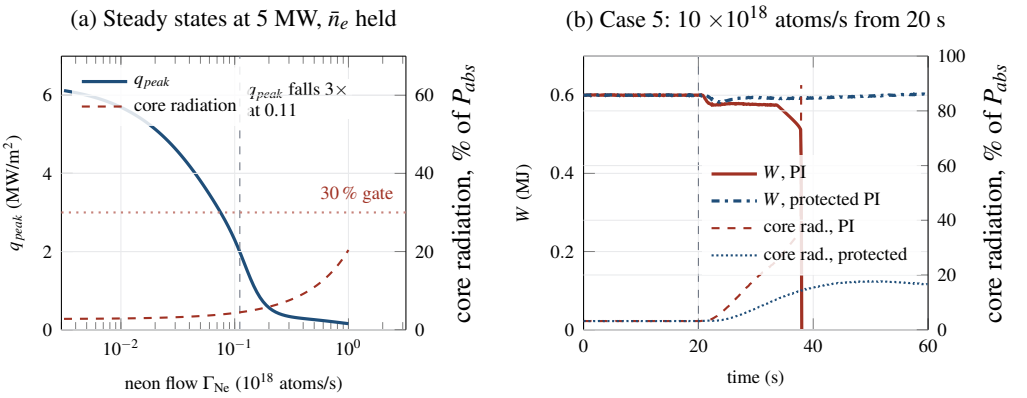

Figure 13.2 The operating window between an attached and a collapsing plasma. (a) Steady states of the model at 5 MW and 6×10^{19} m^{-3} against the neon flow: the peak heat flux (solid) falls smoothly and the core radiation (dashed) rises slowly. (b) Case 5 with the valve held at 10×10^{18} atoms/s from 20 s: under PI (stored energy solid, core radiation dashed) the heating rises to its ceiling and the plasma collapses 18.0 s after the valve opened, while protected PI's core-radiation override (stored energy dash-dotted, core radiation dotted) keeps the core under its limit and completes the pulse.

13.4.2 Detachment and the Operating Window

At fixed 5 MW and a held density, more neon lowers the peak heat flux smoothly and monotonically, as shown in Figure 13.2(a). Over 161 seeding rates from 0 to 10^{18} s^{-1}, q_{peak} falls from 6.33 MW/m^2 (unseeded, 28.4 % radiated, $T_t = 34.8$ eV) to 0.153 MW/m^2 (90.8 % radiated), and no step between neighboring rates moves it by more than 2.2 % of its unseeded value. It falls by three at 1.11×10^{17} s^{-1}, where the core radiates 4.4 % of the absorbed heating, well inside the app's design gate of 30 %. Integrated for 600 s from an attached and from a deeply detached start at 0.5, 1.2 and 3×10^{17} s^{-1}, the plant settles to the same q_{peak} within 9×10^{-5} MW/m^2, so the closure has one branch. The price is stored energy, 0.513 MJ at 10^{18} s^{-1} against 0.601 MJ unseeded, and core radiation, which reaches 20.4 % of P_{abs} at that flow. In the first seeding experiments on W7-X the stored energy fell by about 13 % at maximum radiation [7].

13.4.3 A Collapse Nobody Scripted

Holding the neon inlet at 10^{19} atoms/s from the design point, with the heating and fueling of case 1 held fixed, collapses the plasma without any termination code. The stored energy falls below 90 % after 3.1 s, and at 9.74 s the radiation inside the separatrix overtakes the absorbed heating and the shell's energy drops through the 10 kJ gate. Half the stored energy is gone at 9.80 s. By then the inlet has sent 9.8×10^{19} atoms, the core holds 1.99 % neon and Z_{eff} is 2.56. W7-X collapsed after about 8×10^{19} atoms with a core neon fraction of about 1.3 % inferred from a Z_{eff} of 2.2 [7]. Changing τ_{in} from 4 to 16 s moves the shell's emptying only from 9.1 to 10.2 s (and the loss of half the stored energy from 9.2 to 10.8 s), so the delay is of the order of τ_{in} but is set by the inventory the inlet must deliver, not by the transport time. This matters for control. A loop that waits for the stored energy to fall acts after the inventory is already in the plasma.

After that instant the model is a reduced account of a plasma that is ending. In case 5 under PI the gate holds both zone energies above 0.05 kJ, the energy left at the latch is 0.0024 MJ, and the collapse is declared at 37.985 s with 9.73 MW radiated inside the separatrix against 7.60 MW absorbed. Halving the step to 2.5 ms gives the same latch time.

TRY IT IN THE APP **5 · Seed Too Much**

Load scenario 5 and press Run. Plain PI holds the stored energy, and at 20 s the neon valve opens to 10×10^{18} atoms/s and stays there. Watch the heating climb to its 8 MW ceiling while the stored energy barely moves. Once no heating is left, the plasma collapses, 18.0 s after the valve opened. Switch Controller to Protected PI, press Reset, and the same pulse completes.

13.5 The Control System as Implemented

Three loops run every 20 ms on the packet of Section 13.3: the stored energy (WIC) on the heating, the density (NIC) on the fueling, and the radiation (RIC), either the total radiated power or the measured fraction, on the seeding. Each has its own Manual/Automatic switch. Every PI is the velocity-form PI of Section 1.6.2, with the tracking time equal to the integral time. The shipped tunings are $K_{c,W} = 5\,\mathrm{MW/MJ}$ with $\tau_{I,W} = 0.5\,\mathrm{s}$, $K_{c,n} = 0.5 \times 10^{21}\ \mathrm{s^{-1}}$ per $10^{19}\ \mathrm{m^{-3}}$ with $\tau_{I,n} = 3\,\mathrm{s}$, and $K_{c,R} = 0.3 \times 10^{18}\ \mathrm{s^{-1}}$ per MW with $\tau_{I,R} = 10\,\mathrm{s}$, all reverse acting ($a = +1$ in the convention of Section 1.3, since each measurement rises with its actuator) and without derivative action. Setpoint ramps of 0.05 MJ/s and $0.5 \times 10^{19}\ \mathrm{m^{-3}/s}$ act on the stored-energy and density references in every structure, and the panel shows them in every structure. The actuators then apply their own limits: the heating clamps to the available capacity and slews at 4 MW/s, the valves clamp to their full scale, and the seeding command reaches the plasma 0.3 s later through a 0.2 s lag.

13.5.1 Protected PI

Protected PI adds four elements to the same loops. A heat-flux cascade reads the infrared peak and sets the radiation setpoint, with a direct-acting PI ($a = -1$, since more radiation lowers the heat flux) of $K_{c,q} = 0.2\,\mathrm{MW}$ per $\mathrm{MW/m^2}$ and $\tau_{I,q} = 5\,\mathrm{s}$ whose output is bounded between 30 and 90 % of the injected heating. Two protectors and the WIC request the heating, and a low selector takes the smallest, as shown in Equation (13.13),

$$P_{inj} = \min(u_W, u_q, u_{T_s}), \qquad \Gamma_{\mathrm{Ne}} = \min(u_R, u_{core}), \tag{13.13}$$

where u_q holds the infrared peak at $8\,\mathrm{MW/m^2}$ ($K_c = 0.5\,\mathrm{MW}$ per $\mathrm{MW/m^2}$, $\tau_I = 1\,\mathrm{s}$), u_{T_s} holds the surface at 1050 °C ($K_c = 0.005\,\mathrm{MW/K}$, $\tau_I = 4\,\mathrm{s}$) and u_{core} holds the core chord at 20 % of the injected heating ($K_c = 20 \times 10^{18}\ \mathrm{s^{-1}}$ per unit fraction, $\tau_I = 5\,\mathrm{s}$). The density setpoint is capped at $10 \times 10^{19}\ \mathrm{m^{-3}}$, 0.83 of the cutoff. The three protectors are reverse acting like the loops they override. With anti-windup on, every losing selector input carries external reset, its integral set so that its output equals the selected value plus its own proportional part, so a protector far from its limit stays well above the selected value and noise alone cannot hand it the actuator. The heating PI's ceiling is then the available capacity from the source readback. With anti-windup off, the losers run free and the heating PI believes in the installed 8 MW. Every PI starts by tracking its own actuator at the lined-out point, so the first update is bumpless.

13.5.2 Model Predictive Control

The MPC moves the three actuators together every 0.2 s, a three-input case of the formulation in Section 1.6.3. Its model is the declared nominal plant, the equations of Section 13.3 with the preset's published values frozen at import, plus a 0.1 s heating lag. It never reads a running case's parameters. Each decision linearizes the model about the current estimate by central differences, discretized with an exact zero-order hold in scaled states (the raw states span 21 orders of magnitude) with its affine term, and the seeding delay is carried as two delay states (0.4 s against the 0.3 s inlet). The plan

covers 150 steps (30 s) in 12 graduated move blocks and minimizes

$$J = \sum_{k=1}^{N} \left[Q_W (W_k - W_{SP})^2 + Q_n (\bar{n}_k - \bar{n}_{SP})^2 + \ell \Gamma_{\mathrm{Ne},k} \right] + \sum_j \Delta u_j^{\mathsf{T}} S \Delta u_j + \sum_i (\rho_i s_i^2 + \sigma_i s_i), \quad (13.14)$$

with $Q_W = 300$ per MJ^2, $Q_n = 3$, $S = \mathrm{diag}(0.05, 0.5, 0.02)$ and the neon price $\ell = 0.02$ in the scaled units of MW, 10^{21} and 10^{18} s^{-1}. Four soft bands, each with one slack s_i judged every fifth step, keep $q_{peak} \le q_{SP}$, $T_s \le 1050\,^{\circ}$C, the core chord below 20 % of the heating and $\bar{n}_e \le 10 \times 10^{19}$ m^{-3}, priced (ρ, σ) of $(10^3, 30)$, $(10^2, 10)$, $(10^3, 30)$ and $(10^3, 10)$. The actuators keep hard boxes, with the heating ceiling from the source readback, and per-decision move limits of 0.8 MW, 10^{21} and 2×10^{18} s^{-1}. The linear neon price pulls the seeding down until the heat-flux band holds it up, so the MPC seeds what the heat flux needs. The program is a strictly convex quadratic program solved by the dual active-set solver that the apps share [12].

The estimate is the nominal model run forward every scan from the applied heating, the fueling command and the seeding the inlet is delivering, plus five output biases on W_{dia}, $\bar{n}_e$, q_{peak}, T_s and the core chord, each updated as $b \leftarrow b + \lambda (y_{meas} - y_{model} - b)$ with $\lambda = 1 - e^{-0.02/1}$, a one-second time constant. This output-disturbance form makes the plan offset-free [3], but it corrects only levels, not gains. A plan reaches the actuators only when the solver calls it optimal and all of its values are finite and within their boxes. Otherwise protected PI takes over, labeled as a fallback, and the PI integrals, which track the applied commands while the MPC acts, make the hand-over bumpless. The protectors of Equation (13.13) remain in force as a safety layer over the MPC's commands. With a loop in Manual the MPC stands by, which is not a fallback.

Three models are kept apart in this design. The simulation model is Section 13.3 with the case's true f_{ren}, UA and disturbances. The control model is the same equations with the nominal preset values, which a hidden confinement drop makes wrong. The PI loops use no model, and their tunings were set by measurement on the plant.

13.6 Ten Cases and a Challenge, Measured

The cases appear in Table 13.3. Every number in this section was measured without a display on the app's engine (Section 1.6.1), at the default seed, stepping exactly as the page does, on the shipped case unless the sentence names a change. The scorecard judges the true plant at every 5 ms step after the first 5 s. The operating envelope requires W within 5 % of its setpoint, $q_{peak} \le 10\,\mathrm{MW/m^2}$, $T_s \le 1100\,^{\circ}$C, the core radiation at most 25 % of P_{abs} and no termination. The stored-energy error is $\int |W - W_{SP}|/W_{SP}\,dt$ over the whole run, in seconds.

13.6.1 The Ledger and the Price of Stored Energy

Over the 300 s of case 1, 1510.8 MJ is injected. Of it, 675.0 MJ (44.7 %) lands on the targets and 647.8 MJ is radiated in the SOL layer, while the core radiates 46.2 MJ, the shell 30.8 MJ, other edge losses take 35.5 MJ and 75.5 MJ of microwaves is never absorbed. The energy, hydrogen and neon residuals are 7.5×10^{-15}, 8.9×10^{-15} and 2.7×10^{-15} of their throughputs, and the books close to 2.4×10^{-14} over the challenge and to 6.6×10^{-15} through the collapse of case 5, where the plasma's remaining energy is booked as lost at termination.

Raising W_{SP} to 0.70 MJ under plain PI with the radiated fraction held (case 2) takes the heating from 5.04 to 7.40 MW, and q_{peak} climbs from 4.50 to 7.01 MW/m^2 within 6.8 s of the step and stays at 6.64 MW/m^2 at 120 s. Holding a fraction lets the target take its share of every added megawatt. The same step in case 1, with the cascade in charge, peaks at 7.11 MW/m^2 5.8 s after the step and returns to 4.54 MW/m^2 at 120 s as the seeding rises and the heating settles at 7.51 MW. With the seeding loop in Manual the extra exhaust has nowhere to go, the heat-flux protector takes the heating 4.6 s

Table 13.3 The Ten Cases and the Challenge

#	Structure and events	Teaching point
1	Protected PI, 300 s, no event	every megawatt has a destination
2	PI, W_{SP} 0.60→0.70 MJ at 20 s	stored energy is paid in exhaust
3	PI, $\bar{n}_{e,SP}$ 6→8 at 20 s	density moves confinement, radiation and detachment
4	PI, seeding by hand from 0.03	detachment is slow
5	PI, seeding held at 10 from 20 s	delayed collapse, the core override
6	PI, anti-windup off, W_{SP} 0.68, two sources out 30–90 s	infeasible target, windup
7	Protected PI, $q_{SP} = 7$, cooling halved at 30 s	thermal inertia, surface override
8	PI, $K_{c,W} = 30$, W_{dia} 250 ms late at 20 s	delay under a high gain
9	MPC, W_{SP} 0.50, hidden $f_{ren} \times 0.75$ at 30 s	mismatch, offset-free estimate
10	Protected PI, W_{SP} ramp to 0.85 at 2 kJ/s	which limit binds first
•	Challenge, 360 s: W program, trip at 150, cooling 40 % at 200, sources back at 300 s	margin for the bad day

Units: W in MJ, $\bar{n}_e$ in 10^{19} m^{-3}, seeding in 10^{18} atoms/s, q_{SP} in MW/m^2.

after the step, and W settles at 0.675 MJ, short of its target. Performance was bought with exhaust. Raising the density to 8×10^{19} m^{-3} instead (case 3) lowers the heating from 5.03 to 3.58 MW through the ISS04 density dependence, cools the divertor plasma from 21.7 to 4.3 eV ($D = 0.61$) so the seeding loop closes its valve, cuts q_{peak} from 4.54 to 0.89 MW/m^2 and lifts the core radiation from 3.2 to 7.9 % of P_{abs}.

🕮 **Pause.** Try Exercise 13.2 before reading on.

13.6.2 Detachment by Hand and the Delayed Collapse

Raising the seeding by hand from 0.03 to 0.2×10^{18} atoms/s (case 4) cuts q_{peak} from 4.68 to a third in 52.8 s and to 0.85 MW/m^2 by 180 s, with core radiation at 5.3 %. At 0.4×10^{18} atoms/s the third arrives after 17.1 s. The heat flux follows the neon inventory, which builds over the 25 s edge removal time, not the valve.

With the valve held at 10×10^{18} atoms/s from 20 s (case 5), plain PI hides the problem. The WIC answers the falling stored energy by raising the heating to its 8 MW ceiling, the core neon reaches 3.92 %, and the plasma collapses at 37.985 s, 18.0 s after the valve opened, as shown in Figure 13.2(b). Protected PI's core-radiation override takes the seeding at 20.01 s, holds the core neon at or below 2.34 % and completes the 120 s pulse with the heating never above 7.19 MW. The override acts on a measured symptom, the core chord, before the stored energy falls. The screenshot of Figure 13.3 is case 5 under PI four seconds before the collapse.

🕮 **Pause.** Try Exercise 13.1 before reading on.

13.6.3 Windup, Cooling Loss and Delay

Two tripped sources in case 6 leave 6.40 MW, and no heating level lines the plant out at the 0.68 MJ target, so the target is infeasible and W sits at 0.662 MJ. With anti-windup off the heating integral winds to 8.00 MW, and when the sources return at 90 s the stored energy overshoots to 0.717 MJ (5.4 %) and q_{peak} to 9.53 MW/m^2. With anti-windup on the integral stops at the readback's 6.40 MW, the overshoot is 0.2 % and q_{peak} peaks at 7.85 MW/m^2.

> **TRY IT IN THE APP** **6 · Lose Two Heating Sources (PI)**
>
> Load scenario 6 and press Run. Two of the ten sources trip at 30 s and return at 90 s, with anti-windup off. Watch the stored energy overshoot to 0.717 MJ when the sources come back. Set Anti-windup to On, press Reset and run again: the overshoot falls to 0.2 %.

Halving the target cooling in case 7 doubles C_d/UA from 10 to 20 s. With the cascade at 7 MW/m^2 the bulk climbs for tens of seconds, the surface protector first takes the heating at 69.1 s, the surface

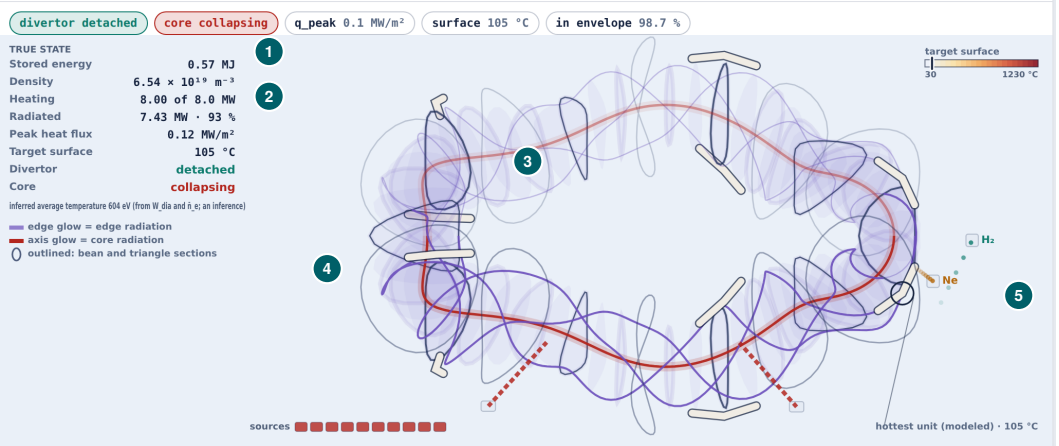

Figure 13.3 Case 5 under plain PI at $t = 34$ s, 14 s after the neon valve opened to 10×10^{18} atoms/s, on the Machine view. (1) The badges read the divertor detached and the core collapsing. (2) The true-state readout: 7.43 MW is radiated (93 %), $q_{peak} = 0.12\,\mathrm{MW/m^2}$ and the surface 105 °C. (3) The axis glow is red, core radiation. (4) The edge glows at full strength, edge radiation. (5) The neon inlet beside one divertor unit. The core radiates 26 % of the absorbed heating with 2.98 % neon, and the heating has climbed to its 8.00 MW ceiling while the stored energy has fallen to 0.571 MJ. The plasma collapses four seconds later.

peaks at 1061 °C, and W is 0.631 MJ at 240 s. Lowering q_{SP} to 5 MW/m^2 at 60 s keeps the protector out, the surface peaking at 1004 °C with W at 0.680 MJ. Radiation, not heating, takes the load.

A 250 ms delay on the diamagnetic signal (case 8) under the aggressive $K_{c,W} = 30\,\mathrm{MW/MJ}$ swings W between 0.562 and 0.638 MJ and the heating between 4.18 and 5.97 MW from 50 to 90 s. At the shipped 5 MW/MJ the same delay leaves W within 0.599–0.602 MJ. The same delay on the bolometers moves the radiation loop by less than 0.03 MW at any tested gain, because the 0.5 s seeding path and the slow neon inventory dominate, so case 8 delays the stored-energy signal and the bolometer delay is left as a disturbance slider.

13.6.4 Model Mismatch, Maximum Performance and the Challenge

In case 9 the confinement falls by a quarter at 30 s, relaxing over 0.5 s, and no controller is told. Every structure removes the offset: W returns within 0.01 MJ of 0.50 MJ and stays for 10 s from 36.4 s under the MPC, 38.6 s under protected PI and 39.8 s under PI, the MPC's stored-energy bias settling at −0.168 MJ. The difference is the exhaust. The MPC raises seeding with heating and holds q_{peak} at or below 4.53 MW/m^2 after the change, against 7.53 under protected PI and 7.32 MW/m^2 under PI.

🕮 **Pause.** Try Exercise 13.3 before reading on.

In case 10 the stored-energy setpoint ramps toward 0.85 MJ, more than the machine holds. Under protected PI the heating saturates at 82.4 s and W peaks at 0.720 MJ, with q_{peak} held at 4.50 MW/m^2 and 5.23 MW radiated, so actuator authority binds first. Lowering q_{SP} to 1 MW/m^2 costs 0.011 MJ (0.709 MJ). Plain PI reaches 0.725 MJ only by letting the target reach 9.87 MW/m^2 and 1018 °C. The MPC reaches 0.753 MJ, but by running the density at 6.62×10^{19} m^{-3}, 10 % over its setpoint, a trade its weights allow and the scorecard does not credit.

The challenge runs 360 s: the program raises W from 0.60 to 0.66 MJ from 30 s, two sources trip at 150 s, the cooling falls to 40 % at 200 s, and the sources return at 300 s. A pass needs 95 % of the judged time in the envelope, judged against the program, and a completed pulse. As shipped, protected PI with the cascade at 7 MW/m^2 spends 66.4 % in the envelope, losing 119 s to the stored

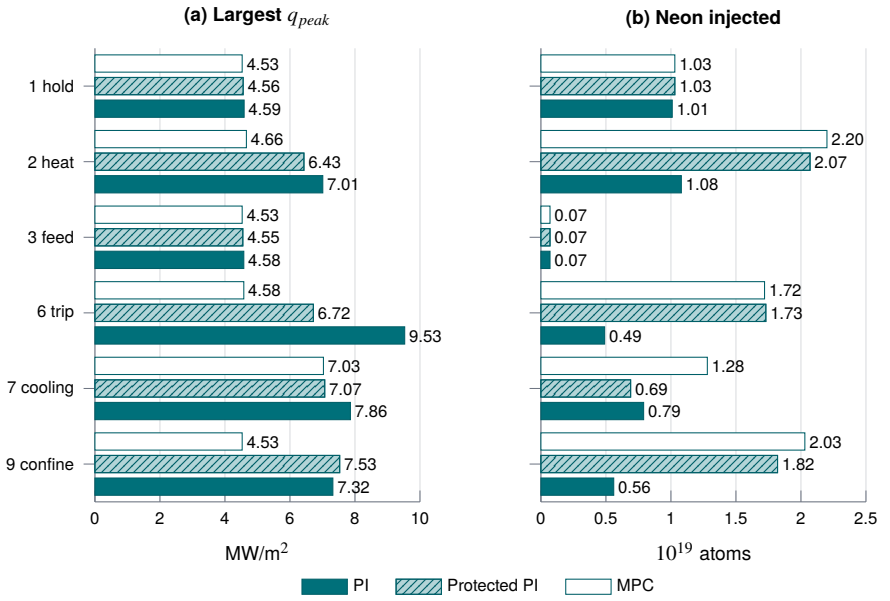

Figure 13.4 The paired comparison of Table 13.4 for the six cases whose rows report the heat flux and the neon: (a) the largest peak heat flux, against a limit of 10 MW/m^2 at the right edge of the axis, and (b) the neon injected. Solid: PI. Hatched: protected PI. White: MPC. Notice that the MPC keeps the largest heat flux near 4.5 MW/m^2 in cases 2, 6 and 9, where protected PI lets it reach 6.43 to 7.53 MW/m^2, and that it does so with more neon than protected PI in cases 7 and 9.

energy after the surface protector takes the heating at 221.2 s, with the surface peaking at 1074 °C. With the suggested q_{SP} of 4.5 MW/m^2 it passes at 100 %, the surface peaking at 822 °C. Plain PI trips the divertor interlock at 267.7 s, after 86.7 % of the time judged up to then, and the MPC passes at 100 % with the shipped settings while its soft surface band keeps the surface at or below 1060 °C.

13.6.5 The Paired Comparison

Table 13.4 lists the three structures on identical seeds, diagnostics, limits and event schedules, and Figure 13.4 draws its heat flux and neon columns. Protected PI uses less neon than the MPC in cases 7 and 9, and plain PI uses the least neon wherever it does not overheat the target. Cases 1 and 3 are ties in the envelope. In cases 4 and 5 a loop is in Manual, the MPC stands by and the runs are identical to protected PI. Case 8 ships a deliberately aggressive PI gain and is not a fair comparison. The MPC's advantage comes from coordination, raising seeding together with heating, which holds q_{peak} at 4.53–4.66 MW/m^2 in cases 2, 6 and 9 against 6.43–7.53 MW/m^2 for protected PI, and from its surface band in case 7 and the challenge. Its quadratic program solves in a median of 1.71 ms in Node.js (95th percentile 1.85 ms, largest 2.6 ms) over the 900 solves of case 9, and in a median of 1.7 ms in Chrome, where the 360 s challenge needs 18.0 s at the fastest playback, twenty times real time, under either structure. At that speed the five-minute case 1 takes 15.0 s of wall time.

+ DID YOU KNOW? Losing Cooling Costs the Most Neon

From the simulator: when the target cooling is halved in case 7, plain PI keeps the plant in the operating envelope only 22.3 % of the time and protected PI 77.2 %. The MPC stays in it throughout, at 100 %, and injects the most neon, 1.28×10^{19} atoms against 0.69 under protected PI and 0.79 under PI (Table 13.4).

Table 13.4 Paired Comparison: Time in the Envelope (%), Largest q_{peak} (MW/m^2) and Neon Injected (10^{19} atoms)

Case	PI	Protected PI	MPC
1 hold	100 / 4.59 / 1.01	100 / 4.56 / 1.03	100 / 4.53 / 1.03
2 heat	97.0 / 7.01 / 1.08	97.0 / 6.43 / 2.07	97.9 / 4.66 / 2.20
3 feed	100 / 4.58 / 0.07	100 / 4.55 / 0.07	100 / 4.53 / 0.07
6 trip[a]	99.8 / 9.53 / 0.49	100 / 6.72 / 1.73	100 / 4.58 / 1.72
7 cooling	22.3 / 7.86 / 0.79	77.2 / 7.07 / 0.69	100 / 7.03 / 1.28
9 confine	97.5 / 7.32 / 0.56	97.2 / 7.53 / 1.82	97.5 / 4.53 / 2.03
10 max W (MJ)	0.725	0.720	0.753[b]
Challenge, shipped	86.7, interlock	66.4, fail	100, pass
Challenge, $q_{SP} = 4.5$	—	100, pass	100, pass

[a] Case 6 ships with anti-windup off, which applies to the PI column. [b] With the density 10 % over its setpoint. Challenge figures are judged against the program.

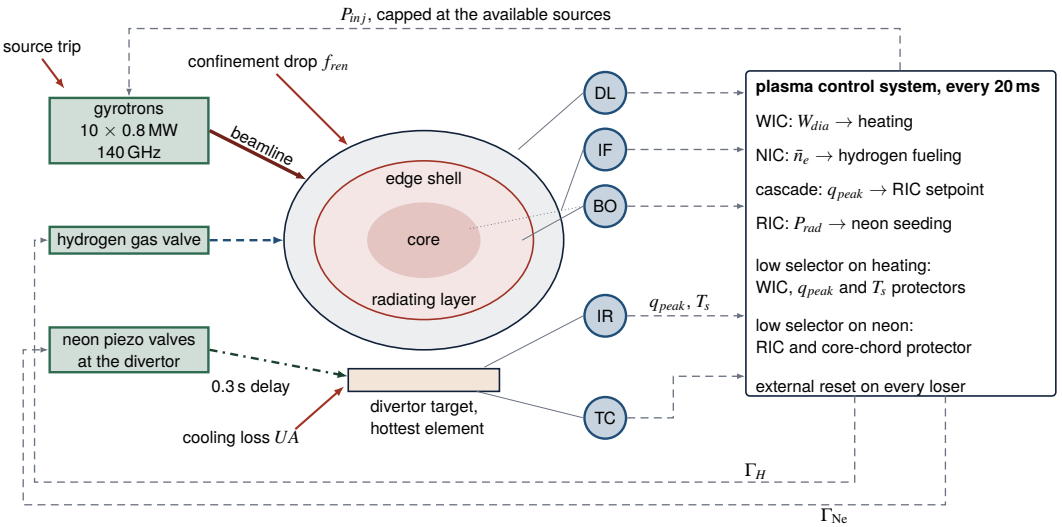

Figure 13.5 The loops of Figure 13.1 with the hardware of a stellarator. Diamagnetic loops (DL), an interferometer (IF), bolometer cameras with an edge and a core chord (BO), infrared cameras on the targets (IR) and a thermocouple (TC) feed the plasma control system, which runs the three PI loops, the heat-flux cascade and the two low selectors every 20 ms. Its commands go to the gyrotrons, the hydrogen valve and the neon valves at the divertor. Thick solid line: microwave power. Dash-dotted: neon. Densely dashed: hydrogen. Thin dashed arrows: signals. Solid arrows: disturbances.

13.7 Technology in Practice

The diagnostics DL, IF, BO, IR and TC and the three actuators of Figure 13.1 stand, on a real stellarator, for magnetic loops, an interferometer, bolometer cameras, infrared cameras on the targets, gyrotrons and gas valves. Figure 13.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The stored energy comes from diamagnetic loops, and the density from an interferometer as a line integral over a sight line [5, 15]. The target heat flux is not measured directly. Infrared cameras in the 8 to 10 µm or 3 to 5 µm band record the surface temperature, and the heat flux follows from the heat diffusion equation solved for the bulk of the tile. Ten divertor modules need ten camera systems,

and a detached target can fall below the system's detection limit of 0.4 MW/m^2 [15]. The chapter's infrared channel is a 30 ms delay that saturates at 20 MW/m^2, on one element that stands for all ten.

> **TECHNOLOGY** **Sensor: Resistive Bolometers**
>
> A bolometer measures radiation as heat. In the stellarator's system each detector is a blackened gold absorber 5 µm thick, with a platinum sensor and reference resistor behind it in a Wheatstone bridge. The bridge voltage and its rate of change give the absorbed power through a cooling time of about 0.1 s, calibrated before each discharge, sampled every 0.8 or 1.6 ms. Two pinhole cameras with 32 and 24 channels view one cross section. The foil also absorbs stray microwaves, which had to be suppressed about 300-fold, and thermal offset drift matters most in long pulses [18]. A plane far from the radiating zone reads low: radiation concentrated in the divertor may be substantially underestimated [7]. The chapter's BO is a 20 ms filter and delay with two chords that see each other's radiation.

What Moves It

The neon reaches the divertor through gas inlet boxes of five piezoelectric valves mounted behind a horizontal target plate, whose capillary nozzles end at the plate's surface and feed the gas straight into the divertor plasma. The valves open for as little as 2 ms, from reservoirs at 2 mbar to 60 bar [15]. In the chapter the seeding command reaches the plasma 0.3 s later through a 0.2 s valve lag.

> **TECHNOLOGY** **Final Element: Gyrotrons**
>
> The heating is microwaves at 140 GHz from gyrotrons, tubes in which an electron beam gives its energy to a wave in a resonator. The original tubes of the stellarator are rated at 1 MW. A 1.5 MW-class prototype was held to 1.3 MW by parasitic oscillations and run routinely at 1.1 MW. Single tubes fail without warning in long pulses, so a power stabilization system now raises the output of the others in real time and holds the total. In a 363 s discharge, cooling problems in the beamline caused growing power fluctuations from 270 s and arcing that ended the discharge [5]. The chapter's ten 0.8 MW sources slew at 4 MW/s, and its trip takes two of them out at once.

Where the Controller Runs

Feedback is not fundamentally required to control a stellarator plasma, yet the long pulses of Wendelstein 7-X now run on three real-time loops: radiated power from the bolometers on the seeding valve's voltage, a stabilization of the total heating power, and density control [5]. In the control system designed for ITER a pulse schedule drives nominal operation, a supervisory layer changes the control in real time on plasma or plant events, and multiple actuators call for shared actuator management [19]. The sources read give no cycle time for these loops. The chapter's PI loops run every 20 ms and its MPC every 0.2 s, and its two low selectors with external reset are that shared management in its simplest form.

> **IN PRACTICE** **Keep the Fueling Loop Out of the Seeding Test**
>
> Density feedback and impurity seeding disturb each other. In the first neon seeding of an island divertor, density changes caused by the seeding had always made the fueling feedback answer with unwanted hydrogen puffs that covered the effect of the neon, so its density target was lowered before the seeding began. Density lost to seeding can instead be made up by feedback-controlled fueling through the divertor [7]. A review of ITER's control makes the general point: core density and core radiation are tightly coupled goals, and where such goals are not joined in one algorithm, managing the actuators they share becomes an important challenge [20].

What Is Modernizing It

Radiated-power feedback has turned seeding from a programmed puff into a regulated variable. Neon led the radiated power through steps from 3.5 to 6.2 MW close to its target, and the same loop kept

a 363 s discharge detached [5]. A tokamak has gone a step further and given each limit its own impurity: argon at the outer midplane regulated the core radiation while nitrogen in the divertor held the peak target heat flux below 5 MW/m^2 with 23 MW of input power [20], the split this chapter's core-radiation override makes with one gas. Fueling is changing too. A screw-type pellet injector built for discharges of up to 30 min delivers pellets of 2 to 3 mm at 400 to 600 m/s and up to 5 Hz, and is to be joined to the real-time density control [5].

13.8 Limitations

The model is a teaching model assembled from public literature and calibrated to a few published measurements, not a W7-X model and not plant data, and its cases are not W7-X results. Every limit, setpoint and tuning is a teaching value. The 10 MW/m^2 heat-flux limit is the published rating of plasma-facing components [6], while the 1100 °C surface limit and the 1200 °C interlock stand in for values that come from the target material and a machine's protection system. Passing a case in simulation demonstrates no safety property. The two-zone profile, the coronal fits with a calibrated layer scale, the smooth single-branched detachment, one target element, a neon species with two transport times and a density-only absorption are the reductions of Section 13.3, and each one removes a behavior that real devices show.

> ⚠ **OPEN CHALLENGE** **A Divertor Regime Not Yet Settled**
>
> For ITER, a fully consistent way to run the divertor with a low enough target heat flux had not been found when its control was reviewed in 2015. The answer is expected to combine enough divertor radiation with partial detachment, but even the controlled variable is open: the local radiated power in core and divertor, the measured or inferred target heat flux, a local temperature or density in the divertor, or a line ratio that stands for the degree of detachment. Fueling and impurity gas or pellets are the actuators. Experiments on present machines have begun to test such schemes, and progress would be one that carries over to ITER [20].

The collapse is resolved only until the shell empties. The zone gate of Section 13.3 is a numerical guard rather than physics, and what follows the gate in a real device (a radiating front reaching the axis, a heating that is no longer absorbed at low temperature) is outside the model. The setpoint ramps act in every structure, so the steps of cases 2 and 3 reach the loops as ramps of 2 and 4 s.

The MPC's model is the nominal plant linearized every 0.2 s, and its biases correct levels but not gains, so a large hidden change is met by feedback rather than by prediction. Its horizon of 30 s is shorter than the 46.5 s neon mode, and the plan sees the far tail of an inventory change only through its last block. Nothing is guaranteed beyond the horizon, and the fallback is protected PI rather than a certified controller. The PI tunings were set by measurement on this plant and are not claimed optimal.

> ⚠ **OPEN CHALLENGE** **Measuring Radiation in a Twisted Plasma**
>
> In a stellarator the three-dimensional field makes an accurate total radiated power hard to derive from one cross section [15]. Tomography in one plane found an up-down asymmetry that three-dimensional edge modeling did not capture [18]. Imaging bolometers are one way forward. One installed in 2021 views the divertor region over 40 by 68 degrees with more than 500 channels, and its upgrade has more than 1400. Progress would be the radiated power measured in three dimensions, which four imaging bolometers combined have achieved on another helical device [21].

13.9 Conclusion

A long-pulse stellarator turns heat exhaust into a constrained control problem, and this simulator presents it with ledgers that close to 10^{-14} of the throughput. Its stored energy costs heating at the ISS04 rate, its divertor detaches smoothly as neon is added, q_{peak} falling by three with the core radiating 4.4 % of the absorbed heating, and a held neon flow ends the pulse in a collapse that no code

scripts and whose delay is set by the delivered inventory. Plain PI hides that collapse behind rising heating for 18.0 s, protected PI's core override prevents it, and anti-windup cuts the overshoot after a source trip from 5.4 to 0.2 %. The MPC holds the heat flux near 4.5 MW/m^2 where protected PI lets it reach 6.4–7.5 MW/m^2, and protected PI uses less neon in two of the cases. The model is reduced and calibrated rather than validated, and its collapse is followed only until the edge empties. A fuller model would add a radial transport equation, non-coronal edge radiation, a second impurity and a temperature-dependent absorption, and an identified rather than declared MPC model in the manner of current radiated-power control. The simulator shows that the highest sustainable performance of a long pulse is set by its exhaust and its actuators, not by the heating it can buy.

KEY POINTS

- Neon at the edge protects the divertor targets, but too much of it in the core ends the pulse in a radiative collapse.
- In this model 20 % more stored energy costs about 60 % more heating, and every added megawatt leaves again as exhaust.
- Plain PI hides an overseeding collapse behind rising heating for 18.0 s, and protected PI's core-radiation override prevents it.
- Anti-windup on the available capacity cuts the overshoot after a source trip from 5.4 to 0.2 %.
- The MPC holds the heat flux near 4.5 MW/m^2 by raising seeding with heating, where protected PI lets it reach 6.4–7.5 MW/m^2.

Where else this applies. A protective override that takes a shared actuator through a low selector while the losers track the winner also guards the fuel valve of Volume 4, Chapter 3 and the compressor of Volume 4, Chapter 2.

References

[1] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[2] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

[3] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[4] L. Spitzer, "The stellarator concept," *The Physics of Fluids*, vol. 1, no. 4, pp. 253–264, 1958.

[5] O. Grulke, G. Acton, J. Adamek, D. Aggelis, R.-M. Alamo-Calderon, C. Albert *et al.*, "Overview of Wendelstein 7-X high-performance operation," *Nuclear Fusion*, vol. 66, no. 11, p. 116003, 2026.

[6] O. Schmitz, Y. Feng, M. Jakubowski, R. König, M. Krychowiak, M. Otte *et al.*, "Stable heat and particle flux detachment with efficient particle exhaust in the island divertor of Wendelstein 7-X," *Nuclear Fusion*, 2020, published online 3 September 2020.

[7] F. Effenberg, S. Brezinsek, Y. Feng, R. König, M. Krychowiak, M. Jakubowski *et al.*, "First demonstration of radiative power exhaust with impurity seeding in the island divertor at Wendelstein 7-X," *Nuclear Fusion*, vol. 59, no. 10, p. 106020, 2019.

[8] A. Kallenbach, M. Bernert, R. Dux, L. Casali, T. Eich, L. Giannone *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.

[9] A. Tsikouras, F. Reimold, G. Partesotti, M. Krychowiak, T. Schröder, M. van Berkel, and the W7-X Team, "Improved feedback control of the radiated power in Wendelstein 7-X using system identification," *Plasma Physics and Controlled Fusion*, vol. 68, no. 8, p. 085004, 2026.

[10] H. Yamada, J. H. Harris, A. Dinklage, E. Ascasibar, F. Sano, S. Okamura *et al.*, "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.

[11] P. C. Stangeby, "Basic physical processes and reduced models for plasma detachment," *Plasma Physics and Controlled Fusion*, vol. 60, no. 4, p. 044022, 2018.

[12] 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.

[13] G. Grieger, W. Lotz, P. Merkel, J. Nührenberg, J. Sapper, E. Strumberger, H. Wobig *et al.*, "Physics optimization of stellarators," *Physics of Fluids B: Plasma Physics*, vol. 4, no. 7, pp. 2081–2091, 1992.

[14] 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.

[15] 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.

[16] S. Sudo, Y. Takeiri, H. Zushi, F. Sano, K. Itoh, K. Kondo, and A. Iiyoshi, "Scalings of energy confinement and density limit in stellarator/heliotron devices," *Nuclear Fusion*, vol. 30, no. 1, pp. 11–21, 1990.

[17] H. P. Summers, W. J. Dickson, M. G. O'Mullane, N. R. Badnell, A. D. Whiteford, D. H. Brooks *et al.*, "Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional–radiative model for light elements," *Plasma Physics and Controlled Fusion*, vol. 48, no. 2, pp. 263–293, 2006.

[18] D. Zhang, R. Burhenn, C. D. Beidler, Y. Feng, H. Thomsen, C. Brandt *et al.*, "Bolometer tomography on Wendelstein 7-X for study of radiation asymmetry," *Nuclear Fusion*, vol. 61, no. 11, p. 116043, 2021.

[19] J. A. Snipes, R. Albanese, G. Ambrosino, R. Ambrosino *et al.*, "Overview of the preliminary design of the ITER plasma control system," *Nuclear Fusion*, vol. 57, no. 12, p. 125001, 2017.

[20] D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson *et al.*, "Novel aspects of plasma control in ITER," *Physics of Plasmas*, vol. 22, no. 2, p. 021806, 2015.

[21] G. Partesotti, G. A. Wurden, F. Reimold, B. J. Peterson, F. Federici, K. Mukai *et al.*, "The infrared imaging video bolometer at Wendelstein 7-X," *Review of Scientific Instruments*, vol. 96, no. 6, p. 063503, 2025.

13.10 Exercises

13.1 ○ Concept check · 1 min

In case 5 under plain PI the neon valve is held at 10×10^{18} atoms/s from 20 s while the seeding loop is in Manual. The stored energy stays near 0.57 MJ for more than ten seconds and then collapses. What holds it there for so long?

(a) The neon needs the inlet delay and several transport times to reach the core, so the core radiation does not change until then

(b) The radiation loop closes the neon valve until the core chord reads twenty percent of the heating

(c) The fueling loop lowers the density, so each neon atom in the core radiates less than at the design point

(d) The stored-energy loop raises the heating toward its ceiling, which offsets the growing radiation until no heating is left

13.2 ○ Concept check · 2 min

In case 2 three plain PI loops hold the stored energy, the density and the measured radiated fraction at 48 %. The stored-energy setpoint steps from 0.60 to 0.70 MJ, and the peak heat flux settles near 6.6 MW/m^2 instead of the 4.5 MW/m^2 it had before. Why?

(a) The heating rises by nearly half, and the unradiated share of every added megawatt reaches the divertor

(b) The seeding valve saturates, so the radiated fraction falls below its setpoint and more power reaches the target

(c) The larger stored energy raises the peaking factor of the target toward its attached value of 3.0

(d) The density rises with the heating, and the denser edge heats the divertor plasma through the two-point closure

13.3 Concept check · 2 min

In case 9 the confinement falls by a quarter at 30 s and the MPC is not told. Its model keeps the nominal confinement factor. Which statement describes what happens?

(a) The MPC leaves a steady stored-energy offset, because the gain of its model is wrong after the change

(b) The output bias on the stored energy absorbs the steady error, so the offset goes, but a transient remains because a bias corrects levels and not gains

(c) The MPC estimates the new confinement factor from the diamagnetic signal, so its model is correct again within a second

(d) The MPC rejects its plans after the change and hands the heating to protected PI, whose integral removes the offset

13.4 Theory and calculation · 6 min

The neon inventories follow Equations (13.11) and (13.12) with $\eta_Z = 1$, $\tau_{pump} = 25\,\text{s}$, $\tau_{in} = 8\,\text{s}$ and $\tau_{out} = 6\,\text{s}$.

(a) Write the two equations as $\dot{z} = Az + bG_Z$ with $z = (Z_e, Z_c)$, and find the two eigenvalues of A and their time constants. *Checkpoint:* -0.0215 and $-0.310\,\text{s}^{-1}$, 46.5 s and 3.22 s.

(b) Find the steady inventories for a seeding flow of 0.034×10^{18} atoms/s, and compare them with Table 13.1. *Checkpoint:* $Z_e = 8.5 \times 10^{17}$ and $Z_c = 6.4 \times 10^{17}$, against 8.6×10^{17} and 6.4×10^{17} in the table.

(c) In case 4 the heat flux takes 52.8 s to fall to a third after the valve is raised. Which part of the seeding path sets that time: the 0.3 s inlet delay, the 0.2 s valve lag, or one of the two modes of part (a)? *Checkpoint:* The slow mode, 46.5 s.

13.5 Theory and calculation · 6 min

The hottest target element follows Equation (13.9) with $C_d = 10^4\,\text{J/K}$, $A_{el} = 0.05\,\text{m}^2$, $R_s = 50\,\text{K}$ per MW/m^2, $T_{cool} = 30\,^\circ\text{C}$ and a nominal $UA = 1000\,\text{W/K}$. The surface protector holds T_s at 1050 °C.

(a) Compute the steady bulk and surface temperatures at $q_{peak} = 4.5\,\text{MW/m}^2$ with nominal cooling, and compare them with Table 13.1. *Checkpoint:* 255 °C and 480 °C, as in the table.

(b) Show that the steady surface temperature is $T_{cool} + q_{peak}(A_{el}/UA + R_s)$. With the cooling at 40 %, as in the challenge after 200 s, find the bulk time constant and the largest steady q_{peak} that keeps the surface at or below the protector's setpoint. *Checkpoint:* 25 s, $5.83\,\text{MW/m}^2$.

(c) With the cooling at 40 %, compute the steady surface temperature at $7\,\text{MW/m}^2$ and at $4.5\,\text{MW/m}^2$. *Checkpoint:* 1255 °C and 817.5 °C.

13.6 Theory and calculation · 6 min

Under the ISS04 scaling with $\tau_E \propto P^{-0.61}\bar{n}^{0.54}$, the stored energy at a fixed density grows as $P^{0.39}$, where P is the loss power $P_{abs} - P_{rad,c}$.

(a) Find the heating factor that 20 % more stored energy needs. Then predict the heating after the case 2 step from 0.60 to 0.70 MJ, starting from 5.04 MW, if the absorbed and radiated shares stay fixed. *Checkpoint:* 1.60, 7.48 MW, against 7.40 MW measured.

(b) In case 3 the density steps from 6 to $8 \times 10^{19}\,\text{m}^{-3}$ at a held stored energy. Predict the loss power after the step from the design value of 4.629 MW, and then the injected heating, with $\eta = 0.950$ and the core radiating 7.9 % of the absorbed heating after the step. *Checkpoint:* 3.11 MW, 3.55 MW, against 3.58 MW measured.

(c) After the case 3 step the divertor plasma is at 4.3 eV. Compute the detachment degree from Equation (13.8), and compare it with the design point at 21.4 eV. *Checkpoint:* 0.61 against 0.000.

13.7 App challenge · 8 min

Load scenario 5, Seed too much, and change no setting. Plain PI runs with the seeding loop in Manual, and the neon valve is held at 10×10^{18} atoms/s from 20 s. Use the default seed and the 20× playback speed.

(a) Run until the pulse ends. Read the Discharge tile of the scorecard, and find in the exported CSV file the first row in which `P_inj_MW` reaches its ceiling. *Checkpoint:* Radiative collapse at 38.0 s. The heating reaches 8.00 MW at 33.7 s, within 0.2 s.

(b) Switch Controller to Protected PI, press Reset and run the full 120 s. Read the Discharge tile and the line "won" on the seeding card, and find in the exported file the largest `P_inj_MW` and the largest `cNe_core_pct`. *Checkpoint:* Pulse completed, seeding won by the core-radiation protector, 7.18 MW and 2.34 %, each within 0.02.

Explain from the two runs why the override acts on a measured symptom rather than on the stored energy.

13.8 App challenge · 8 min

Load scenario 6, Lose two heating sources (PI). Two of the ten sources trip at 30 s and return at 90 s, and Anti-windup ships Off. Use the default seed, the 20× playback speed and runs of 150 s.

(a) Run as shipped. In the exported CSV file read `P_req_MW` and `P_inj_MW` at 60 s, and find the largest `W_MJ` after 90 s. Read the Peak heat flux (max) tile. *Checkpoint:* 8.00 MW asked against 6.40 MW applied, 0.716 MJ, within 0.002 MJ, 9.53 MW/m^2 at 90 s.

(b) Set Anti-windup to On, press Reset and run again. Read the same three values. *Checkpoint:* 6.40 MW asked and applied, 0.681 MJ, within 0.002 MJ, 7.85 MW/m^2 at 4 s.

Express each stored-energy peak as an overshoot of the 0.68 MJ setpoint, in percent of the setpoint, and explain why the heat-flux tile of run (b) names a time before the trip.

13.9 App challenge · 9 min

Load the challenge, The long pulse. It runs 360 s: the stored-energy program rises to 0.66 MJ from 30 s, two sources trip at 150 s, the cooling falls to 40 % at 200 s and the sources return at 300 s. Protected PI ships with the Peak heat-flux setpoint at 7 MW/m^2. Use the default seed and the 20× playback speed, and read the Challenge target and Target surface (max) tiles at 360 s after each run.

(a) Run as shipped. *Checkpoint:* Not met, 66.4 % in the envelope, most lost to the stored energy (119 s), 1074 °C.

(b) Set Peak heat-flux setpoint to 4.5 MW/m^2, press Reset and run. Read the Neon injected tile as well, for both runs. *Checkpoint:* Met, 100.0 %, 822 °C, 3.37×10^{19} atoms against 0.65×10^{19}.

(c) Set the Peak heat-flux setpoint to the bound of Exercise 13.5(b), rounded down to the slider's step of 0.1 MW/m^2, press Reset and run. *Checkpoint:* 5.8 MW/m^2: met, 100.0 %, 1047 °C.

13.10 Application area · 3 min

The model uses three temperature levels on the target surface: the protector holds 1050 °C, the operating envelope ends at 1100 °C, and the interlock trips every source after 1 s above 1200 °C. Explain what each level is for, which of them a control engineer may tune, and two features of a real divertor target that the single-element model leaves out and that would move these numbers.

13.11 Application area · 4 min

In case 10 the stored-energy setpoint ramps toward 0.85 MJ. Under protected PI the heating saturates and the stored energy peaks at 0.720 MJ with the heat flux held at 4.5 MW/m^2. A planning engineer must choose one upgrade to raise the sustainable stored energy: more microwave sources, targets rated above 10 MW/m^2, or a larger allowance for core radiation. Which does the case support, and what will the upgrade cost in exhaust?

13.12 Review the agent · 5 min

A learner asks an AI assistant how to pass the challenge without spending more neon. The assistant replies: “Keep the Peak heat-flux setpoint at 7 MW/m^2 so the cascade seeds little neon. The run fails because the surface protector is too slow when the cooling falls, so raise the Heat-flux protector gain from 0.5 to 2.0. A faster protector will hold the surface and the stored energy together.” Find the flaw, and say how you would check it with Equation (13.9) and one run.

CHAPTER 14

Summary and Outlook

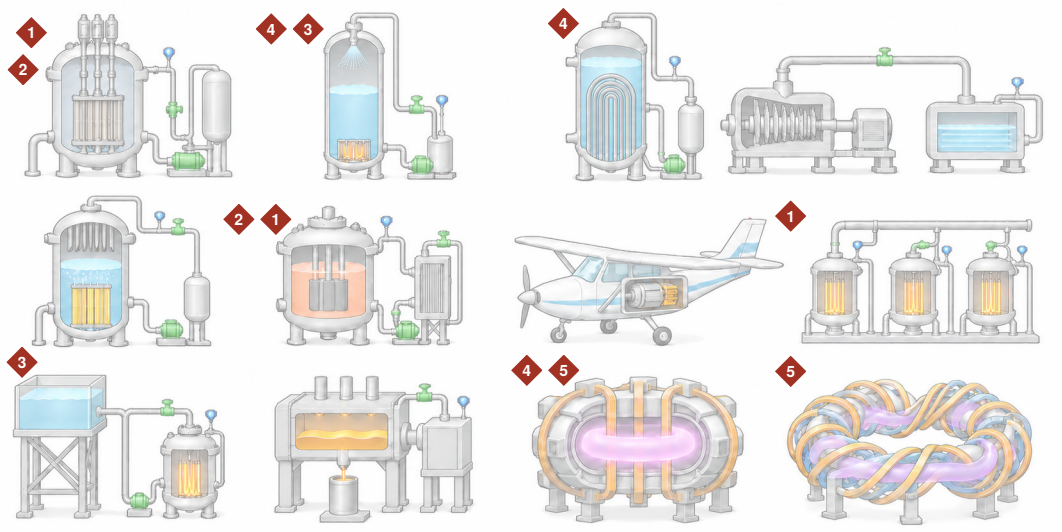

Open gaps 1 moving coefficients 2 unseen margins 3 hidden failures 4 layers that act 5 model error

Notice that the reactor with its rods, the salt reactor, the pressurizer and the tokamak each carry two gaps, and that no gap marks the steam generator, the boiling water reactor, the aircraft or the melter. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Nuclear and Fusion Plants Taught**

The main idea. Each chapter of this book ran one nuclear or fusion plant in a simulator, from the rods of a reactor to the hot edge of a fusion machine. In most of them the physics already steadied the plant, and the controller removed what was left.

The challenge. Some limits arrive faster than any reading can warn, and some margins never show on a gauge. A loop can hold its reading steady while the plant quietly spends its reserve.

Why it matters. The best plants measured the cause of a disturbance, let the protective rules act first, and tracked the margin left.

Twelve chapters controlled twelve plants, from a pressurized water reactor's control rods to the heat exhaust of a stellarator. Each plant was read in three parts in Section 1.4.7: physics that helps or hurts, a limit set by protection or equipment, and the control layers between them. Table 1.1 listed the three parts for every chapter. This chapter adds what each arrangement achieved and where it failed, draws a short guide for choosing a method, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

14.1 What the Chapters Showed

Tables 14.1 and 14.2 compare the chapters. The second column condenses Table 1.1. The other columns name a structure that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

The columns make five points, and Figure 14.1 draws one pair of results from each row.

Table 14.1 Physics, Limits and Layers, Chapters 2 to 7

Chapter	Physics and limit	What worked	What failed	Measured result
2, PWR rods	Feedback follows the load and leaves the offset to the rods. A thermowell, a deadband, an insertion limit.	Mismatch feedforward on a turbine ramp. A flux cascade on a dropped rod. A 0.8 °C deadband against noise.	$K_c = 8$, past the thermowell's margin. A 0.2 °C band with a minimum rod speed. Any tuning near the insertion limit.	Ramp IAE 41.1 against 10.1 °C min, rod travel 123 against 30.7 %. Dropped rod peak 5.93, 3.58 and 2.23 °C (Table 2.8). Bank at 26.8 % against a 30 % limit (Section 2.5.5).
3, pressurizer	An insurge compresses the bubble faster than full spray condenses. Relief valve and level limits.	Overrides first, then the split-range settings. The coolant's own compliance.	Tuning alone in the challenge. A high gain with no deadband, which cuts the peak at 8.9 times the travel.	Peak within 20 kPa for three architectures, time off level 1014 against 40 s (Table 3.4). Challenge worst error 1203 shipped against 27 kPa suggested (Table 3.5).
4, steam generator	Shrink and swell. The zero's crossing moves from 38 to 65 s as load falls. Two level trips.	Three elements with a fast inner flow loop. MPC forecasting the voids from the steam flow.	One element at a higher gain, which cycles at $K_c = 3.0$. The challenge as shipped, which trips high-high.	IAE 11486, 9094 and 6521 mm s for one, two and three elements (Table 4.6). Pickup peak 167.6 against 122.1 mm under MPC (Table 4.7).
5, steam dump	The shut-in secondary reaches its safety valve in 14.80 s against a 4 s RTD.	Turbine impulse pressure as feedforward. Two low-selected overrides and a reactor runback.	Feedback alone, at seven times the gain or with an instantaneous RTD.	Peak 83.37 feedback, 79.74 with feedforward, 76.39 bar coordinated (Table 5.4). Gain alone 81.73, instant RTD 82.87 bar (Table 5.5).
6, boiling water reactor	Voids are moderator, feedback and a level error at once. Level trips and a power-flow map.	Three-element feed. The six tunings and the feedforward gain restored together.	One element at the same gains. Gain scheduling switched on over the original gains.	Ramp worst level +287.4 against −63.0 mm (Table 6.6). Challenge worst level 410 as given, 47 corrected, 3097 mm scheduled, a scram (Table 6.8).
7, molten salt reactor	The pump sets the delayed neutron fraction, 419 pcm of 650 at design flow. A transport delay.	The fuel pump on outlet temperature. A lagged load feedforward. MPC with a margin limit.	The secondary flow pairing, with 16 s of dead time. A static feedforward, which raised the peak.	Peak 7.05 feedback, 7.42 static and 2.64 °C lagged (Table 7.7). Margin case 476.76 against 0 s below 415 pcm, IAE 299.24 against 2324.30 °C s (Table 7.10).

The physics does most of the work, and the controller removes what is left. With every rod parked, the reactor of Chapter 2 took its power from 100 to 89.9 % within eight minutes of a load cut and left the coolant 13.9 °C above program (Section 2.5.1). The aircraft reactor followed a takeoff with nobody driving, lagging by 90.9 points of rated power and then overshooting (Table 8.4). Natural circulation removed the decay heat of Chapter 10 with no pump. In each case the controller's task was the residue: an offset, an overshoot, an inventory. The same physics can turn. At a moderator coefficient of $+0.5\,\text{pcm}/^\circ\text{C}$ the rod loop's process gain rose by a factor of 2.8, and the shipped tuning swung 12.6 °C peak to peak (Section 2.5.6). A graphite coefficient of $+6\,\text{pcm}/^\circ\text{C}$ made the salt reactor unstable (Table 7.9). A design that rests on a physical property must keep measuring it.

The measurement or the transport, not the core, sets the gain. The rod loop's ultimate gain was 9.01 with a three-second thermowell and 37.3 without it (Table 2.3). The salt reactor's gain limit was set by its transport delay, and its secondary flow pairing waited 16 s before the outlet moved

Table 14.2 Physics, Limits and Layers, Chapters 8 to 13

Chapter	Physics and limit	What worked	What failed	Measured result
8, aircraft reactor	The core follows a takeoff unaided, then overshoots. Relative gain 0.86 at takeoff, 3.58 at idle.	Preview of the published mission. A coordinated predictive controller on the go-around.	The same controller with the preview off. Two independent loops on the go-around.	Takeoff IAE 6358 without preview, 729 % s with it (Table 8.4). Go-around peak outlet 859 two loops, 815 °C coordinated (Table 8.5).
9, three SMRs	The header holds 7.55 s of steam, the modules answer in 52.5 s. A coolant ceiling that moves.	Coordinated allocation with a 2 s bypass. A margin-proportional split. A turbine runback.	Master–slave on a large ramp. An equal megawatt split. A scram with no runback.	Ramp: master–slave 2.647 to 3.874 MPa and the customer lost, coordinated 3.166 to 3.449 MPa (Table 9.6). Margins 8.5, 4.9, 0.0 against 5.41, 5.66, 4.59 °C (Section 9.6).
10, decay heat	Natural circulation removes a power-law heat. The pool pays 165.2 kg per kelvin of cooldown avoided.	A margin override on a 2.0 MW auxiliary train with a setpoint pair.	A latched-open condenser. Throttling alone, with no auxiliary train.	Pool at 8 h 2.67, 7.61 and 31.76 t, tubes uncovered at 3.93, 6.26 h and never (Table 10.5). Challenge 69.9 and 120.8 minutes over the limit, overrides none (Table 10.7).
11, waste melter	A cold cap insulates the pool. Joule heating unstable at a fixed tap. Coverage and off-gas limits.	Power regulation. Feed under coverage control with overrides on a low selector. Feedforward to the exhauster.	A fixed voltage tap, which walks the pool. Independent loops, which overrun the coverage ceiling.	Glass 1968 against 2444 kg, 188.3 against 0 min over the ceiling, worst pressure 90.0 against 54.2 Pa with feedforward (Table 11.5).
12, tokamak plasma	An unstable mode at 356 s^{-1}. A gain floor, a latency ceiling and a recoverable region.	A gain between floor and ceiling. A model observer for velocity. Priority for the vertical loop.	$K_p = 0.3$, below the floor. 10 V/mm at 300 µs latency. Clipped allocation.	$K_p = 0.3$ lost at 166.6 ms. Limit cycle of 1.45 mm rms. Coil energy 33.9 against 9.1 J. Clipping lost at 72.7 ms, priority held (Table 12.4).
13, stellarator exhaust	Neon protects the targets and poisons the core. Stored energy costs heating.	A core-radiation override. MPC raising seeding with heating. A lower heat-flux setpoint.	Plain PI, which hides a collapse for 18.0 s. Protected PI at its shipped setpoint in the challenge.	Challenge in the envelope 86.7 PI with an interlock, 66.4 protected and 100 % MPC. Peak heat flux 4.53 to 4.66 against 6.43 to 7.53 MW/m^2 (Table 13.4).

at all (Table 7.4). A single-element drum loop stopped improving at $K_c = 2.0$ and cycled at 3.0, on a disturbance that needed a five-point valve move (Section 4.6). The tokamak's position loop settled into a 1.45 mm limit cycle at 10 V/mm with 300 µs of latency and held to 0.059 mm at 100 µs (Table 12.4). In every one of these plants the first number to find is the lag between the plant and the controller.

A disturbance faster than its measurement needs a measurement of the disturbance. The steam dump of Chapter 5 has fifteen seconds. Seven times the gain reached 81.73 bar and an instantaneous RTD 82.87 bar, while the turbine impulse pressure wired to the same loop reached 79.74 bar and lifted no safety valve (Tables 5.4 and 5.5). Steam flow in a three-element loop, the power mismatch on a turbine ramp, a lagged load signal in the salt reactor and the published mission of the aircraft did the same job. The point of entry decided the structure. On the turbine ramp the mismatch feedforward cut the rod travel from 123 to 30.7 %, and on the dropped rod the flux cascade gave the smallest peak, 2.23 against 3.58 °C, because the excore detector saw the rod first (Table 2.8).

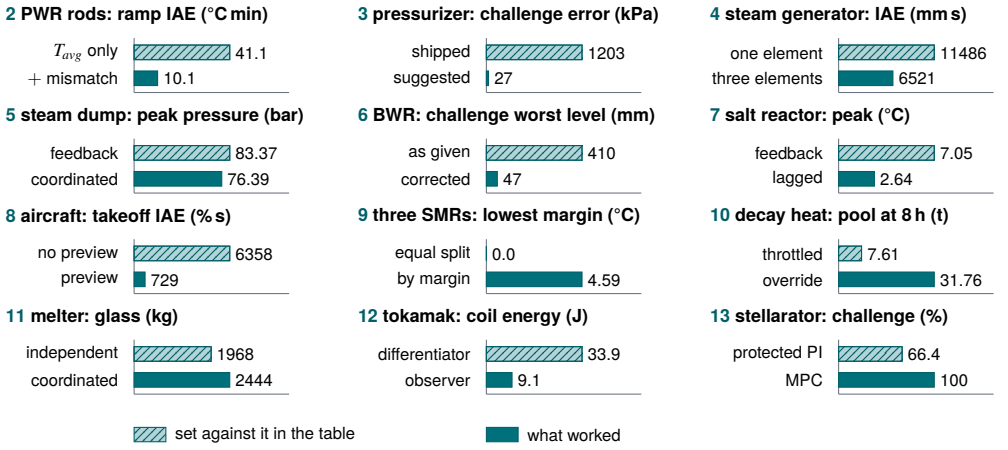

Figure 14.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 9, 10, 11 and 13, and a smaller one in the others. The three-module panel sets the lowest of the three coolant margins, and the equal split's 0.0 °C draws no bar. Notice that in most rows the better result came from a change of structure or measurement, not of tuning alone.

The limiting layer is structure, and tuning cannot replace it. Across the pressurizer's three architectures the insurge peak moved by 20 kPa. In the challenge the suggested split-range settings on the two-loop structure still ended 320 kPa low, and the same settings with the overrides passed (Section 3.7). The coordinated steam dump peaked at 76.39 bar, and its reactor runback was worth 2.71 bar of its 3.35 bar advantage (Section 5.6). On the steam header the turbine runback was the one change both necessary and sufficient in the challenge (Section 9.6). The tokamak kept a plasma under priority allocation that clipping lost (Table 12.4), and the stellarator's core-radiation override took the seeding at 20.01 s and completed a pulse that plain PI lost at 37.985 s (Section 13.6). A layer that switched on without its tuning made things worse: gain scheduling over the boiling water reactor's original gains scrammed the plant (Table 6.8).

The binding limit is often a margin or a resource that no setpoint shows. The rod bank of Chapter 2 passed through its insertion limit at 26.8 % while the temperature was on program, because nothing in the objective mentioned shutdown margin. The salt reactor's predictive controller kept the delayed neutron fraction above 415 pcm for the whole margin case at 7.8 times the temperature error (Table 7.10). The condenser pool of Chapter 10 is spent and does not come back, and two passing configurations one setpoint apart finished 2.97 t of water apart (Table 10.7). An equal megawatt split on the steam header left one module at 0.0 °C of margin with the header pressure back in its band (Section 9.6). A pressurizer with a spray valve stuck open held its pressure within 4 kPa by running every heater bank, and only two scorecard numbers, the fighting time and the heater energy, showed that the loop had spent its upward authority (Section 3.7). A controlled variable on its setpoint said little about the margin left.

14.2 Choosing a Method

The chapters suggest an order of questions for a nuclear or fusion plant that is new to the reader. Each names the chapter where its answer decided the result, and Figure 14.2 draws them as a chart.

1. *What does the plant do with no controller?* Run it with the actuators parked and size the residue, as the rod and aircraft chapters do. The residue is the controller's job.

2. *Does the helpful physics keep its sign over the fuel cycle, the load range and the fouling?* If not, schedule, retune or design across the range, and measure the range first (Chapter 2, Chapter 7).
3. *What lag sets the gain ceiling?* A thermowell, a transport delay, a moving zero or a latency each bounds the gain before the core does (Chapter 4, Chapter 12).
4. *How many seconds to the nearest protection limit, and which signals act inside them?* Count them with Equation (1.8). If no measurement of the effect is fast enough, measure the cause (Chapter 5).
5. *Where does each disturbance enter?* Feedforward from a disturbance that arrives through the heat path, a cascade on one that arrives as reactivity (Chapter 2).
6. *Which limit must win, and what enforces it?* Build the override, runback or priority allocation before tuning the loop beneath it (Chapter 3, Chapter 9, Chapter 12).
7. *Which margin or resource does no setpoint show?* Put bank margin, pool inventory, delayed neutron margin or actuator authority on the scorecard, and decide whether it becomes a constraint (Chapter 10, Chapter 7).

A predictive controller paid where it carried a forecast or coordinated actuators that the loops ran apart. Forecasting the voids from the measured steam flow, it held a pickup to 122.1 mm where three-element control reached 167.6 mm (Table 4.7). It cut the boiling water reactor's level dip on a pressure step from 114.9 to 17.7 mm, at 2.5 times the integrated pressure error (Table 6.10). Raising seeding with heating, it held the stellarator's peak heat flux near 4.5 MW/m^2 and passed the challenge as shipped (Table 13.4). It lost or tied elsewhere. On the pressurizer's insurge it used less spray and heat than the overrides for a 14.7 % higher IAE (Section 3.7.5). On the steam header it tied coordinated allocation, 71.61 against 72.56 MPa s (Table 9.7). In the melter it used less energy and made less glass on the loaded scrubber, and no controller won every constraint (Table 11.6). The tokamak's explicit MPC settled an ELM in 0.52 ms against 1.32 ms but recovered kicks only up to 19.63 mm against 19.91 mm, because it reversed the voltage late (Table 12.5). Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Nuclear plants that follow the grid. Nuclear plants are commonly run at full output whenever online, although they can ramp, follow load and provide reserves. A unit commitment and dispatch model that represents xenon transients and the change of core reactivity over the fuel cycle found, for a system with large shares of solar and wind power, that flexible nuclear operation lowered operating costs, raised the reactor owner's revenue and reduced the curtailment of renewables [1]. The rod chapter shows what such a schedule must respect. Its sixteen-hour maneuver drove the bank fully out for thirteen minutes on the return to full power, and a cut to 40 % crossed the insertion limit (Section 2.5.5). In the editor's view the xenon, boron and insertion limits will move from the rod controller into the dispatch that asks for the maneuver.

> ⧖ HISTORY **1979 to 2011: Instruments for the Worst Day**
>
> The second history box of Chapter 1 ended in 1979, with a commission that found instruments not designed to follow an accident. The accident at Fukushima Daiichi began on 11 March 2011. There the loss of instrumentation and control left the operators with little indication of actual plant conditions, and the accident guidelines did not cover the loss of the instruments that display the key parameters. Not knowing the state of the spent fuel pools led to an effort to add water to them. The report of the Director General of the International Atomic Energy Agency concluded that the instrumentation and control needed in accidents beyond the design basis must remain operable, and that strategies for manual control of vital equipment need to improve [2]. In the editor's reading the lesson of 1979 was learned again, for a plant without power.

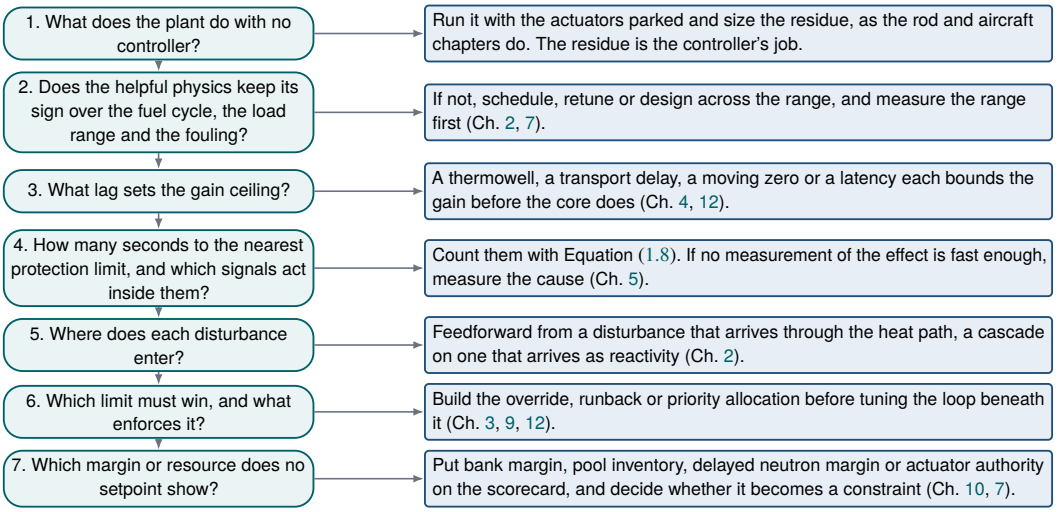

Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the plant with no controller and the lag come before any question about structure, and that the last question is about what no setpoint shows.

Rod control that adapts to the core. From 2008 to 2015 an adaptive reactor control system was installed in several pressurized water reactors in Germany. It receives a new set of reactivity coefficients for every core reload, determined for the fuel cycle as a function of the boron concentration, and balances reactivity, xenon included, to keep a margin to the bank limits that assure shutdown reactivity and the insertion limits. Its more precise control allows reduced dead bands near full power [3]. The same guide notes that unplanned maneuvering can move the power and burnup profile of a core away from its predictions, so the core physics calculations may need to be checked again after each maneuver. The rod loop of Chapter 2 runs on one fixed deadband, and nothing in its objective mentions shutdown margin. In the editor's view the next step is coefficients estimated from plant data between reloads, with their uncertainty stated.

Autonomy for multimodule plants. A framework for the autonomous control of advanced reactors starts from multiple units on one site and from remote small modular reactors, whose economics depend on lower operating and maintenance costs. It calls for diagnosis, planning, reconfiguration and self-validation inside the control system, and it describes an autonomous control system for nuclear power as a goal not yet realized [4]. The steam header of Chapter 9 is a small instance of that problem. Its allocator splits load against margins it estimates from measurements, and its runback sheds turbine load after a scram so that the industrial customer keeps its steam. In the editor's view an allocator of this kind will have to report the uncertainty of each module's estimated limit before an autonomous system can be trusted to act on it.

Instruments watched while they run. Online condition monitoring detects and diagnoses abnormalities from the process signals while the plant operates. Its intent is to mine the recorded data for changes in the relationships between measured parameters, which the indicators of a typical control system do not include. Shown that a sensor has not drifted since its last calibration, a plant can extend the calibration period [5]. The scorecards of this volume did the same on a small scale. The fighting time and the heater energy of Chapter 3 showed a stuck spray valve that the pressure trend hid, and the pool inventory of Chapter 10 separated two configurations that the temperature scorecard could

not. In the editor's view such measures of spent authority belong beside the drift checks of an online monitoring system.

Plasma control that must work on the first discharge. A review of plasma control for ITER finds its design driven by the nuclear mission, by aggressive commissioning constraints and by a limited number of discharges, with little time to tune algorithms on the machine. It names actuator sharing and fault handling among the novel problems [6]. Priority allocation in Chapter 12 is a two-loop case of actuator sharing, and the stellarator's low selectors hand the heating to a protector and the seeding to a core override as limits approach. In the editor's view fusion controllers will be designed and tuned almost entirely in simulation, which makes the model's stated recoverable region the quantity to verify.

Learned controllers and predictors. A controller trained by deep reinforcement learning has commanded the full set of control coils of a tokamak, producing and holding a diverse set of plasma shapes while meeting physical and operational constraints, with a notable reduction in design effort [7]. A deep learning method has forecast disruptions and made reliable predictions for machines other than the one it was trained on, which a large reactor needs because it cannot afford training disruptions [8]. The vertical loop of Chapter 12 gives the yardstick. Any controller, learned or written, is bounded by the same gain floor, latency ceiling and recoverable region. In the editor's view a learned controller will be accepted on the kicks and latencies it was not trained on.

Exhaust under feedback. Detachment is required to keep heat and particle fluxes on the divertor tolerable, and feedback control of the location of an impurity emission front has been demonstrated with real-time diagnostics, a dynamic characterization of the divertor plasma and an offline controller design [9]. The stellarator of Chapter 13 adds the core to that problem, since the same neon that detaches the divertor collapsed the plasma 18.0 s after a valve was held open. In the editor's view exhaust control will be designed as a constrained multivariable problem with core radiation as a hard limit.

Long pulses with the exhaust under feedback. Wendelstein 7-X is designed for discharges of up to 30 min and 18 GJ of heating energy. A new feedback loop holds its total radiation at a target by changing the voltage of the seeding valve from bolometer measurements, and a second raises the remaining microwave sources in real time when one of them trips. With both, a discharge ran 363 s with nearly constant heating and radiation and 1.8 GJ of heating energy, until arcing caused by cooling problems in a heating beamline ended it [10]. The same report notes that feedback is not fundamentally required for plasma control in a stellarator, whose magnetic configuration is steady by design. Chapter 13 works with the same two actuators, the seeding valve and the heating, and adds the core limit that a long pulse must respect. In the editor's view the pulses of a reactor-scale stellarator will depend on how the exhaust loops share those actuators.

Figure 14.3 sets these developments in three threads.

14.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [11]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

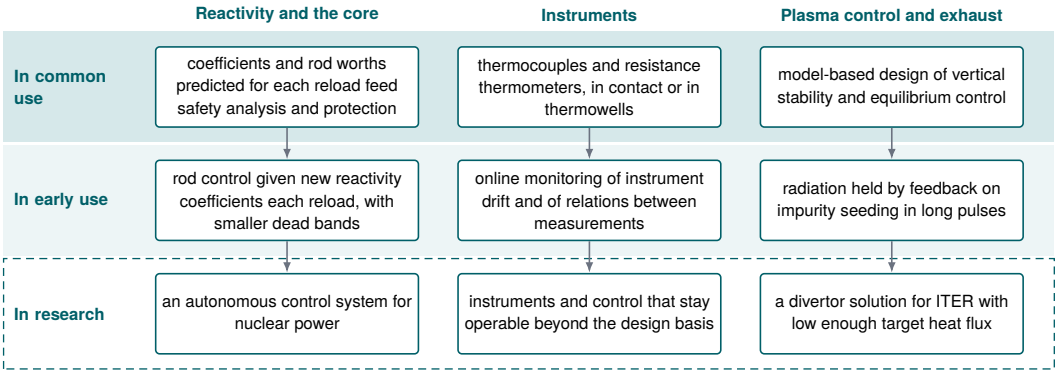

Figure 14.3 Developments in three threads. The top lane holds what the cited source describes as usual or routine practice, or as the most common, the middle lane what it reports in service at some plants or used operationally on one machine, and the bottom lane what it reports as a goal or a need not yet met, in the editor's reading. Sources by thread: reactivity and the core [3, 4], instruments [2, 5] and plasma control and exhaust [6, 10]. Notice that each thread moves from what is computed or measured before operation toward what must be known while the plant runs.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [12]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

14.4.1 How This Volume Was Made

Each of the twelve apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. Eight of the twelve referee reports call themselves self-reviews, and the other four say they were written against the draft and applied or describe a reviewer's stance. None names a reviewer independent of the writing. Seven carry a later pass on predictive-control material, two of them called independent. Eleven papers were measured again in a consistency round and an app-fix round, and the twelfth was written after both. The defect logs hold 125 entries with a severity, in eleven of the twelve logs, since the tokamak's log numbers its findings without one. Of these, 30 are wrong output or behavior, 71 documentation or content, 15 cosmetic or display and 9 something else, such as a completion rule or a design question. Some round notes list fixes without a severity, so the count is a floor. The editor set the scope of each round and what each agent could change. Four cases show what went wrong and what caught it. In each, a layer between the physics and the limit was not doing what its text said.

Overrides that never acted. The pressurizer's texts said the high-pressure override held the charging valve shut through the peak, and counted it among three things that win the challenge. Measured, the selector was active for 0.0 s in the challenge, and an override setpoint raised out of reach reproduced the run to the last digit. The insurge raises the level, so the level loop has already shut the charging valve before the pressure reaches the override. The steam dump's level override never selected in twelve scenarios, because on a load rejection the pressure only rises and the level cannot swell. Agents measuring the apps' claims found both. The first was confirmed by raising the override out of reach and seeing nothing move, the second by lowering its setpoint until it acted. The texts now say when each override acts.

A pressure loop that sprayed while the pressure was low. After an outsurge 1 MPa below setpoint, the pressurizer's back-calculation with a 30 s tracking time wound the integral to +594 % against a proportional and derivative sum of −825 %. As the error shrank the output crossed the deadband, and the spray was commanded for 589 s with the pressure more than 70 kPa low. A consistency round that measured every scenario again found it. Tracking at the integral time cut the figure to 12.2 s, and the same fix was then needed on the override, which had taken the charging valve with the pressure 836 kPa below its own setpoint.

Controllers that knew the plant. Six entries in four logs record a controller or an operator aid that read the run's hidden parameters or the plant's true state. The aircraft reactor's dynamic matrix controller and preview feedforward solved their models on the run's parameters and scaled the prompt jump by the true neutron power. With the instruments untouched, 16 of 44 parameters scaled by 1.25 moved its gains, and the fouling scheduled in the challenge reached the feedforward as a better model. Fixing it raised the takeoff power error from 634 to 729 % s, so the predictive controller now trails the two loops there. The steam header's allocator read each module's hidden thermal ceiling, and a hidden change moved its split by up to 57.310 MW. The tokamak's observer read the margin drop of a minor disruption that no instrument measures. Each entry records the same probe: move a hidden value with the measurements held and ask whether the controller moves.

Causes the runs did not support. The book round tested each chapter's causal sentences by run and corrected six in four chapters without adding a number. The boiling water reactor's runback dip was said to come from fewer voids plus a feedwater cut. At the dip the core void had risen from 40.3 to 43.2 % and the vessel had gained 1.5 t, so the dip is the steam held below the surface falling with steam production. The steam dump's long time off program was said to be overcooling, yet every one of its 114.8 seconds had the coolant above its reference. The tokamak's 20 mm kick was said to be lost at any gain, and a velocity gain of 2 V per m/s held it on the same supply. That last claim had reached the app's briefing and Chapter 1.

The book round. The twelve chapter records list 140 rows of numeric verification, each a stated value or group of values set beside its reproduction. Eight report no drift, and four name a printed value that a run does not reproduce exactly. On the steam header the margins of the margin-proportional split had moved after a fix to the allocator, and a later repair round measured them again. The writer of Chapter 1 read the apps' action flags and found five places in three chapters that named a loop's action against the convention of Section 1.3, and its own reactor example first settled 17 % low until it was checked against its closed form. Three statements of Chapter 1 were corrected from the chapter reports. Every chapter pass logged at least one defect in its app. The melter's pass found that after an exhauster trip the low selector and the pressure feedforward chattered, moving the feed command between about 630 and 1800 kg/h from one scan to the next. The editor approved a repair of the model and its results were regenerated. Eleven chapters of this volume were then synchronized with the repaired apps and papers, and the melter's exercises were run again. With the repair the loaded-scrubber ranking reversed, and coordinated control now makes more glass than the predictive controller.

Every error above was caught by a check that could fail: a layer switched off to see what moved, a scenario measured again, a hidden value moved while the measurements were held, a causal sentence tested against the variable it names. None was caught by rereading the text that contained it. Four working rules follow. Disable each protective layer once and confirm that it acts on the event it guards. Check what an integral does while its output is clamped. Give a controller only what the plant measures and the design declares. Test every causal sentence with the variable it names.

14.5 Gaps in the Control of Nuclear and Fusion Systems

Five gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 14.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 14.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *Coefficients that move.* The rod loop's process gain rose 2.8 times when the moderator coefficient went slightly positive, and the salt reactor became unstable at a graphite coefficient of $+6\,\mathrm{pcm}/^{\circ}\mathrm{C}$. The steam header now estimates each module's ceiling from measurements because the fouling moves it. A guide to flexible operation describes rod control that receives a new set of reactivity coefficients for each core reload, and notes that unplanned maneuvering can move the core away from the predictions that its physics calculations rest on [3]. The gap is hard because a coefficient shows only through the loop it changes, and a licensed plant cannot be stepped at will to measure it. Online estimates of reactivity coefficients and thermal limits that state their own uncertainty, and a licensed way to schedule a loop on them, are not routine.

2. *Margins that the controller cannot see.* The rod bank crossed its insertion limit with the temperature on program, and the salt reactor held its delayed neutron margin only by accepting 7.8 times the temperature error. The adaptive rod control of Section 14.3 balances reactivity, xenon included, to keep a margin to the bank limits that assure shutdown reactivity and the insertion

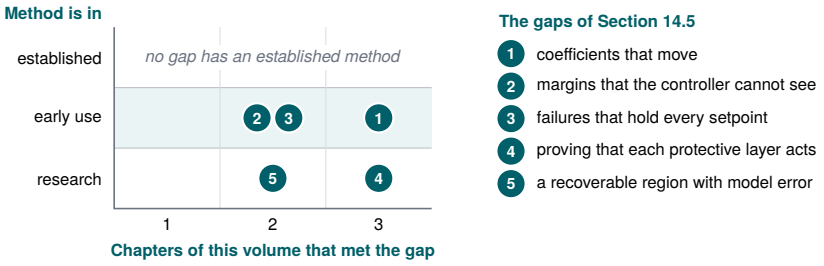

Figure 14.4 The five gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 2, 7 and 9), two for gap 2 (2 and 7), two for gap 3 (3 and 10), three for gap 4 (3, 5 and 12) and two for gap 5 (12 and 13). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the two gaps about proof, of a protective layer and of a recoverable region, sit lowest.

limits [3]. The gap is hard because the margins of a safety analysis are computed in advance for assumed conditions, while a controller acts on the present ones. Methods that put the margins of a safety analysis into a control layer's objective, while keeping that layer outside the safety system, are not settled.

3. *Failures that hold every setpoint.* Two lost heater banks changed nothing on the pressurizer's trends until an outsurge asked for heat. A quarter of the decay-heat pool was lost with no temperature showing it. Online monitoring mines recorded data for changes in the relationships between measured parameters, relationships that the indicators of a typical control system do not include [5]. The gap is hard because a loop that holds its setpoint hides the effort it spends, and that effort sets off no alarm. A standard set of authority measures, such as fighting time, actuator reserve and inventory used, computed beside the error, is not yet common practice.

4. *Proving that each protective layer acts.* The level cap on the steam dump never selects in any of its chapter's twelve scenarios, because a load rejection there can only shrink the level (Section 5.6). The pressurizer's high-pressure override takes the charging valve in an early proportional hand-over, before the pressure reaches its setpoint (Section 3.5.2). The tokamak's verdict checks every physics step, because brief excursions fall between the 2 kHz log samples. Test generation that forces every override, runback and interlock to act, and checks the record at the rate of the event, is not routine. The Open challenge box after this list takes the gap further.

5. *A recoverable region with model error.* The vertical loop's region was computed for a rigid, massless plasma with one wall mode, and the stellarator's model is calibrated, not validated. On one tokamak the largest controllable displacement, measured by switching vertical control off for varying intervals, was 0.15 to 0.24 m against about 0.40 m calculated, a difference far beyond the noise and more likely due to errors in modeling the non-axisymmetric structures and to the plasma striking the wall [13]. ITER permits only 1 to 10 unmitigated vertical displacement events at its highest current and stored energy [6]. The gap is hard because the measurement that settles the region is the event the design must avoid. A recoverable region stated with the model's uncertainty, which a controller designed in simulation could be held to on the first discharge, is missing.

OPEN CHALLENGE **Showing That a Protective Layer Acts**

Three chapters met gap 4, as many as any gap. A guide to nuclear instrumentation and control notes that digital systems may be too complex to test in all aspects, and that a common cause failure that leaves the outputs as they are is not revealed until a safety action is demanded [5]. A review of plasma control for ITER requires every control algorithm to be verified in simulation under the expected noise, disturbances and key off-normal events before use [6]. In the editor's view the gap is hard because a layer written for a rare event is exercised only by that event. Progress would be a test that shows each layer acting on its own event, run again after every change.

14.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [14]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [12]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [15]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [16]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [17]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that

has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- The physics does most of the work, and the controller removes what is left.
- In every one of these plants the first number to find is the lag between the plant and the controller.
- A disturbance faster than its measurement needs a measurement of the disturbance.
- The limiting layer is structure, and tuning cannot replace it.
- A controlled variable on its setpoint said little about the margin left.

References

[1] J. D. Jenkins, Z. Zhou, R. Ponciroli, R. B. Vilim, F. Ganda, F. de Sisternes, and A. Botterud, "The benefits of nuclear flexibility in power system operations with renewable energy," *Applied Energy*, vol. 222, pp. 872–884, 2018.

[2] International Atomic Energy Agency, *The Fukushima Daiichi Accident*. Vienna, Austria: International Atomic Energy Agency, 2015, sTI/PUB/1710, report by the Director General with technical volumes 1–5.

[3] ——, *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.

[4] R. T. Wood, B. R. Upadhyaya, and D. C. Floyd, "An autonomous control framework for advanced reactors," *Nuclear Engineering and Technology*, vol. 49, no. 5, pp. 896–904, 2017.

[5] 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.

[6] D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson, A. Kallenbach, E. Kolemen, J. Lister, D. Moreau, A. Pironti, G. Raupp, O. Sauter, E. Schuster, J. Snipes, W. Treutterer, M. Walker, A. Welander, A. Winter, and L. Zabeo, "Novel aspects of plasma control in ITER," *Physics of Plasmas*, vol. 22, no. 2, p. 021806, 2015.

[7] J. Degrave, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de las Casas, C. Donner, L. Fritz, C. Galperti, A. Huber, J. Keeling, M. Tsimpoukelli, J. Kay, A. Merle, J.-M. Moret, S. Noury, F. Pesamosca, D. Pfau, O. Sauter, C. Sommariva, S. Coda, B. Duval, A. Fasoli, P. Kohli, K. Kavukcuoglu, D. Hassabis, and M. Riedmiller, "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.

[8] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," *Nature*, vol. 568, no. 7753, pp. 526–531, 2019.

[9] T. Ravensbergen, M. van Berkel, A. Perek, C. Galperti, B. P. Duval, O. Février, R. J. R. van Kampen, F. Felici, J. T. Lammers, C. Theiler, J. Schoukens, B. Linehan, M. Komm, S. Henderson, D. Brida, and M. R. de Baar, "Real-time feedback control of the impurity emission front in tokamak divertor plasmas," *Nature Communications*, vol. 12, no. 1, p. 1105, 2021.

[10] O. Grulke, G. Acton, J. Adamek, D. Aggelis, R.-M. Alamo-Calderon, C. Albert *et al.*, "Overview of Wendelstein 7-X high-performance operation," *Nuclear Fusion*, vol. 66, no. 11, p. 116003, 2026.

[11] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[12] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[13] D. Humphreys, T. Casper, N. Eidietis, M. Ferrara, D. Gates, I. Hutchinson, G. Jackson, E. Kolemen, J. Leuer, J. Lister, L. LoDestro, W. Meyer, L. Pearlstein, A. Portone, F. Sartori, M. Walker, A. Welander, and S. Wolfe, "Experimental vertical stability studies for ITER performance and design guidance," *Nuclear Fusion*, vol. 49, no. 11, p. 115003, 2009.

[14] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[15] International Electrotechnical Commission, "Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels," IEC 62443-3-3:2013, Edition 1.0, 2013.

[16] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[17] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

14.6 Exercises

14.1 Discussion · 8 min

An agent reviewing Chapter 5 proposes to replace the coordinated structure with the feedback-only dump loop at a seven times higher gain and a faster temperature sensor, "because it has fewer parts to maintain". Using Tables 5.4 and 5.5, say what the proposal misses and what you would keep.

14.2 Discussion · 8 min

The rod loop of Chapter 2 and the pump loop of Chapter 7 both rest on reactivity feedback that helps. Using the moderator coefficient case of Chapter 2 and the graphite coefficient sweep of Table 7.9, discuss how you would keep these loops tuned over a fuel cycle, and what you would measure.

14.3 Discussion · 8 min

Section 14.4.1 reports two overrides in this volume that never acted on the event they were written for. Write a test plan, in five to seven steps, that an agent could run on any app or plant model to show that each override, runback and interlock acts when it should and stays out when it should.

14.4 Discussion · 8 min

In Chapter 10, Chapter 7 and Chapter 9 the binding limit is a margin or a resource rather than a setpoint. For each, say what quantity belongs on the scorecard, and argue whether it should become a constraint in a predictive controller or stay a number for the operator.

14.5 Discussion · 8 min

A vendor offers a learned controller for the vertical position loop of Chapter 12, trained in simulation. Using the floor, the ceiling and the recoverable region of that chapter and of Section 1.4.6, list what you would ask the vendor to show before the controller is tried on a plasma.

14.6 Discussion · 10 min

Section 14.4.1 reports six entries in which a controller or operator aid read the run's hidden parameters or the plant's true state. An agent is asked to commission the module allocator of Chapter 9 on a new plant. Write a procedure, in six to eight steps, that keeps the allocator honest, states what the agent may read and write, and says what a person must approve.

Answers to Selected Exercises

Short answers are given for the odd-numbered exercises. Numeric answers from an app are for its default seed and settings unless the exercise says otherwise.

Chapter 1. Introduction

1.1 (a) 1.25 and 50 s. (b) 3.85 ms and 4.15 times.

1.3 (b). A velocity actuator supplies the integrator, so any steady error would keep the bank moving.

1.5 (a) 0.857 and 0.571, 386 pcm and 1157 pcm. (b) It rises at 0.090 per hour.

1.7 (d). Feed follows steam, so the inventory stays balanced, while the voids still move the level.

1.9 (a) 2.0 V/mm and 6.0 mm. (b) 4.0 V/mm and 3.0 mm, and 7.5 mm with 15 V.

Chapter 2. PWR Rod Control and Inherent Feedback

2.1 (d). Bank position is the integral of the commanded speed, so the open loop holds an integrator and a proportional controller leaves no offset.

2.3 (b). Removing the thermowell lag raises the ultimate gain from 9.01 to 37.3, a factor of 4.14.

2.5 (a) 308.17 °C, $\Delta T_c = -1.83\,^{\circ}\mathrm{C}$. (b) $\Delta T_f \approx -31.7\,^{\circ}\mathrm{C}$. (c) About +102.5 pcm, 5.7 % of travel.

2.7 (a) 89.9 % and +13.9 °C, about 346 MW shed. (b) +10.8 °C with xenon at 101.1 %, which supplies part of the negative reactivity that the hot coolant had been supplying.

2.9 (a) 8.10 °C, 49.0 %, 71 %, missing the error and travel targets. (b) 1.74 °C, 49.6 %, 52 %, a pass. (c) 2.41 °C, 49.2 %, 41 %, missing the error target. The inner flux loop of the cascade sees the dropped rod first.

2.11 (a) $7/2.5 = 2.8$. At $K_c = 3$ the loop swings about 12.6 °C peak to peak in the last ten minutes of a forty-minute run. (b) Two of gain scheduling, periodic retuning and a design robust across the range, all of which need the measured range of α_M over the cycle.

Chapter 3. Pressurizer Split-Range Control

3.1 (c). The surge flow falls by W_{sp}, and W_{sc} cancels when the two mass balances are added.

3.3 (a). Full spray condenses 7.26 kg/s while the surge line peaks at 34.72 kg/s, and the peak moves only 20 kPa across three architectures.

3.5 (a) 11 236 kg, 19.01 m^3, 0.373 % per percent against the programmed 0.370 %. (b) 1.257 m^3 and 129 kg of steam. (c) 1013 kg of pool water, about 840 kg of it charged and the rest condensed steam.

3.7 (a) +1.243 kPa/s and −11.450 kPa/s, a ratio of 9.21. (b) 1388.5 kW and 0.155 kg/s, about 10.7 % as steam. The chapter's demand acts half a second later.

3.9 (a) 159 kPa, no PORV lift, lowest level 33.2 %, not passed. (b) 35 kPa, no lift, 33.2 %, passed, with 1000 kg of spray. The structure removes most of the error and the gain removes the rest.

3.11 (a) About 40 s. (b) About 5.1 K. (c) The relief valves always reseat, and the property fits are clamped to 11 to 19 MPa, so a loop that goes two-phase cannot be followed.

Chapter 4. Three-Element Drum Level and Inverse Response

4.1 (c). Each kilogram of feedwater condenses 0.238 kg of steam, and the void term of Equation (4.1) falls faster than the water term rises.

4.3 (a). With feedwater equal to steam the inventory holds, and the level falls by the full void shrinkage that the lower load requires.

4.5 (a) 12.7, 16.2 and 18.1 sec. (b) −38, −181 and −244 mm. (c) 38 mm against 63 mm.

4.7 (a) −3.95 mm at 19.4 sec after the step. (b) +7.16 mm and −11.11 mm. (c) 38.1 sec, against 65 sec at a quarter load.

4.9 (a) 10.4 kg/s with the valve at 69 %; −149 mm at 340 sec; 86.51 m s. (b) −7.6 mm and 1.459 m s, a ratio of 59.

4.11 The valve-position feedforward is 14.5 points wrong at 40 % flow, and a fixed load change is 2.2 times worse at 40 % load. The model has no circulation instability.

Chapter 5. Nuclear Steam Dump on a Load Rejection

5.1 (b). It balances the heat the tube bundles deliver against the latent heat of the steam leaving and the sensible heat the feedwater absorbs.

5.3 (d). Disabling the runback raises the peak to 79.10 bar, so it is worth 2.71 bar of the 3.35 bar gained.

5.5 (a) −0.0513 °C/%. (b) 19.2 %/°C, and 0.31 at the shipped gain. (c) 110 % and 37 % of valve, against 5.2 °C for the whole dump.

5.7 (a) 83.4 bar and 1529 kg, half open 12.1 s after the breaker. (b) 81.70 bar with nothing lifted. (c) The reference follows excore power, so no input moves when the breaker opens.

5.9 (a) 83.0 bar, 1223 kg, 123 mbar, derate down to 0.78. (b) 80.0 bar, nothing lifted, 110 mbar, 3519 kg through the reliefs, no derate. (c) The override holds the back pressure below the interlock band, and the reliefs and the runback take the balance.

5.11 The window between a normal day and a reportable lift is the two seconds between 8 and 10 s, and no training program makes it reliable. A feedback only loop is half open only 12.1 s after the breaker, as late as the operator, so the automation needs the turbine signal.

Chapter 6. Boiling Water Reactor Coupled Loops

6.1 (a). Pressure is a reactivity term held by the fastest handle, feedwater is the only inventory handle, flow moves power fastest, and the rods relieve the pumps.

6.3 (c). It is the first line of Equation (6.12).

6.5 (a) 295.4 kg/s per unit valve per MPa. (b) 2770.7 kJ/kg. (c) −4.19 kPa/s per percent and −16.8 kPa/s for the park, steeper in the run because the power falls.

6.7 (a) −40.9 mm at 33.4 s with the collapsed level at +11.8 mm. (b) −1.68 m^3, 2.1 % of the sub-surface steam. (c) 7.64 mm/s, about 26 mm by 3.4 s at the full rate.

6.9 (a) 77 mm, 0.166 MPa, 0.45. The level now holds, and the pressure band fails. (b) A low-level scram near 104 s, then Level 8 near 146 s. (c) The schedule doubled an already excessive power-loop gain.

6.11 (i) High power on low flow raises the decay ratio of the density-wave mode, and the plant is left by adding flow. (ii) One point-kinetics node with no regional mode, and a calibrated oscillator whose amplitude is illustrative.

Chapter 7. Molten Salt Reactor: Pump Sets the Margin

7.1 (d). The core rise $P_0 n/W_f$ follows the pump at once, and the same speed sets the transit times of Equation (7.3).

7.3 (b). The colder inlet raises the power and widens the core rise, and the cold leg is a full loop transit away.

7.5 (a) 38.3 % and 96.4 %. (b) 7.5 s and 13.3 s, then 246 pcm, 37.8 %, 0.247 dollars. (c) 1.22.

7.7 (a) 78.2 %, then 615.8, 655.0 and 576.6 °C. (b) +0.8, −10.0 and +11.6 °C, with a rise of 78.4 °C. (c) Rod and pump fixed, so the temperature terms must return to their starting sum. $P_0 n/W_f = 78.2\,^{\circ}\mathrm{C}$.

7.9 (a) 2.64, 7.42 and 7.04 °C, then 188, 329 and 416 °C s. (b) 63 % and 55 %. (c) The correction arrives a loop transit before its disturbance and has to be withdrawn.

7.11 (a) At +5 pcm/°C, where the sum is zero. At +4 the outlet settles 128.6 °C low at 14.3 % power. (b) The sum $\alpha_f + \alpha_g$, measured in service. IAE goes from 269 to 914 and travel from 25 to 170 % between -1 and +6.

Chapter 8. Aircraft Molten Salt Reactor Load Following

8.1 (d). Only $\beta_{\text{eff}} = 390$ pcm of the static 650 pcm decays in the core, and the core must make up the 260 pcm that decays outside it.

8.3 (b). The event is commanded and not scheduled, so the profile the feedforward reads ahead does not contain it.

8.5 (a) 0.893 MW/°C, 140 °C, 7.39 °C. (b) 18.5 °C and 5.91 °C. (c) About 6.3 °C.

8.7 (a) 149.5 % at 157 s. (b) 90.8 points at 104.5 s. (c) 48 °C. The reactivity arrives through a chain of delays and is still present when the ramp ends.

8.9 (a) –9 °C (859 °C). (b) 35 °C (815 °C). (c) 146.5 s against 100.25 s: the predictive controller moves the pump when the go-around is commanded, the two loops only after the outlet has climbed.

8.11 (a) The controller is unconstrained and then clipped, with the outlet in a weighted zone, and the thermal model is not validated. (b) A constrained quadratic program with a hard outlet limit. The model would still be notional.

Chapter 9. Multi-Module Steam Header Load Allocation

9.1 (c). Header pressure is the measured CV that every module pushes up and every user pulls down.

9.3 (a). The bypass answers the error in seconds, the valve position controller walks the reactors in behind it, and the allocator divides the demand by available duty.

9.5 (a) 12.8, 13.0, 13.1 °C. (b) 16.1, 13.0, 9.3 °C, spread 6.9 °C. (c) 612 MW against 654 MW.

9.7 (a) 70.5 MW at 10 minutes and 91.7 MW at 20 minutes. (b) 3.385 to 3.401 MPa, inside the ±0.08 MPa band. (c) 2.693 to 3.917 MPa.

9.9 (a) 1.936 MPa, the user lost, the header back on 3.400 MPa. (b) 2.943 MPa, 41.7 MW shed, the user kept. (c) 0.042 MPa, kept, no relief: the challenge passes.

9.11 Per-module margin and margin spread, thermal reserve and the industrial user's status. A real header makes it worse, because transmitters at opposite ends disagree by more than a calibration bias.

Chapter 10. Decay-Heat Removal on a Depleting Heat Sink

10.1 (b). The plant has a limit, a level and a coping time to meet, not a setpoint, so the task is to spend a fixed inventory over a horizon.

10.3 (d). On a controlled loop the disturbance appears in the manipulated variable.

10.5 (a) 372.76 MJ/K and 165.2 kg/K. (b) 32.84 GJ, about 14.55 t. (c) 3.19 t/h and 25.1 t.

10.7 (a) 15.72 t, 211.9 °C and 29.92 t, 299.8 °C. (b) 3.93 h and 6.27 h. (c) 14.53 t predicted, 14.20 t measured.

10.9 (a) 342 °C, 69.9 min, 4.81 t, a fail. (b) 313 °C, 0.0 min, 11.56 t, 28.3 GJ, a pass. (c) 319 °C, 41.1 min, 11.14 t, a fail.

10.11 (a) 22.0 K/h. (b) It does not hold: after the first ten minutes the peak rate is 32.2 K/h. The argument for throttling is water, the 14.2 t it saves by four hours.

Chapter 11. Vitrification Melter Across Four Time Scales

11.1 (a). At steady state every kilogram of dry solids that arrives leaves as glass, so production is $\dot{F}(1-w)y_{ox}$ whatever the power.

11.3 (c). At a fixed tap the delivered power falls by 7.52 kW/K, more than the 6.84 kW/K the pool sheds, so the design point is open-loop unstable.

11.5 (a) 1.82. (b) 7.52 kW/K, net 0.68 kW/K of positive feedback, about 13 h. (c) 1605 kW and 1178 kW, set by the transformer voltage.

11.7 (a) 1129.3 °C, 1104.1 °C, 869 kW. (b) 1145.3 °C. (c) On a fixed tap the power falls with the pool by 7.52 kW/K, more than the 6.84 kW/K the pool sheds.

11.9 (a) 0.6 min, about +35 Pa. (b) 1236 kg, 1052 kg/h. (c) 785 kg, about 36.5 % less, with the feed fixed at 1800 kg/h.

11.11 (a) 3.5 t per day, 29 % of the design production. (b) Each degree closer to the ceiling is throughput bought with melter life. The corrosion rate of the refractory against temperature is outside the model.

Chapter 12. Tokamak Vertical Stabilization

12.1 (a). The floor at $\kappa = 1.8$ is about 0.45 V/mm, and below it the coil cannot overcome the destabilizing push.

12.3 (c). The MPC turns its voltage positive with the plasma 3.8 mm from the midplane, the loop at 7.7 mm, and the late reversal carries the plasma past the far edge of the region.

12.5 (a) 7.69×10^4 N/m, –0.96 kA, –0.96 V. (b) 10.51 mm. (c) 9.71 V/mm, near the top of the slider and close to the gain at which scenario 5 limit-cycles.

12.7 (a) Lost at 166.6 ms. (b) 0.48 mm at 0.2 s, then lost at 482.4 ms. (c) 0.023 mm. Just below the floor the plasma diverges slowly, so a short run hides the loss.

12.9 (a) Lost at 661.0 ms. (b) Lost at 661.2 ms. (c) Held, 3.94 mm and 0.80 mm, a pass. The allocation decides: the radial loop takes the supply when the kick follows the β drop.

12.11 Losing the vertical position is a vertical displacement event and ends the discharge, while a radial error can be corrected later. The margin assumes the vertical loop has every volt, so under clipping it overstates what is left.

Chapter 13. Stellarator Heat Exhaust Control

13.1 (d). The heating climbs from about 5 to 8 MW and then has nothing left to give, so the stored energy falls away.

13.3 (b). The bias settles at –0.168 MJ and the stored energy is back within 0.01 MJ of 0.50 MJ from 36.4 s.

13.5 (a) 255 °C and 480 °C. (b) 25 s and 5.83 MW/m^2. (c) 1255 °C and 817.5 °C.

13.7 (a) Radiative collapse at 38.0 s. Heating at 8.00 MW from 33.7 s. (b) Completed, core-radiation protector, 7.18 MW, 2.34 %.

13.9 (a) Not met, 66.4 %, 1074 °C. (b) Met, 100.0 %, 822 °C, 3.37×10^{19} against 0.65×10^{19} neon atoms. (c) At 5.8 MW/m^2: met, 100.0 %, 1047 °C.

13.11 More heating. Each 20 % of stored energy needs about 1.6 times the heating, and the cascade must radiate the extra exhaust with neon.

Chapter 14. Summary and Outlook

14.1 A good answer computes the clock: the shut-in secondary reaches the first safety valve in 14.80 s against a 3 s dump stroke and a 4 s RTD. It reads the tuning table: seven times the gain reaches 81.73 bar, an instantaneous RTD 82.87 bar, and both together 81.30 bar, while the turbine impulse pressure added at the shipped gain reaches 79.74 bar and lifts no safety valve. It concludes that the gap is closed by a measurement of the disturbance, not by gain or sensor speed. It keeps the overrides and the runback, which take the peak to 76.39 bar and are worth 2.71 bar of the 3.35 bar the coordinated structure gains.

14.3 A good answer lists the protective layers and the event each guards, runs each event with the layer in place and with it disabled or its setpoint out of reach, and compares the runs, so that an identical run flags an inert layer. It forces each layer to act with a setpoint moved into range, records how long it is selected, checks the record at the rate of the event, and tests the hand-back after the event. It reports inert layers to a person rather than deleting them, since a layer may be inert on one event and needed on another.

14.5 A good answer asks for the gain floor and latency ceiling of the learned law, or tests that bracket them, such as holding at elongation 2.0 where the floor rises to about 0.70 V/mm and at 300 µs of latency where 10 V/mm cycles at 1.45 mm rms. It asks for the largest recoverable kick against the loop's 19.91 mm at elongation 1.9, for behavior under clipped and priority allocation, for coil energy under 0.3 mm of noise against the observer's 9.1 J, and for results on disturbances and parameters outside the training set, with the model's omissions stated.

Glossary

Each term is given as this volume uses it, in one plain sentence, followed by the chapter where it is first used in earnest. The two indexes give every page.

Aircraft Reactor Experiment A small test reactor of 1954 whose molten fluoride fuel was pumped through its core, built for aircraft propulsion. *Chapter 1*.

average power range monitor (APRM) The average of many detectors inside a boiling water reactor core, read as the reactor's power. *Chapter 6*.

back-calculation A way to stop windup: when an output is limited, the controller's integral is pulled toward the output actually applied. *Chapter 2*.

boiling water reactor (BWR) A reactor in which the cooling water boils in the core and the steam goes straight to the turbine. *Chapter 1*.

bolometer A detector that measures the power a plasma radiates from the heating of a thin absorbing foil. *Chapter 13*.

boron A neutron absorber dissolved in the coolant of a pressurized water reactor and changed slowly to carry large changes of reactivity. *Chapter 2*.

borosilicate glass The glass into which high-level waste is melted, so that the waste is locked into a durable solid. *Chapter 11*.

bubbler A tube that blows gas into the glass pool of a melter to stir it and speed the melting. *Chapter 11*.

circulating fuel Fuel dissolved in a salt that a pump carries through the core and around an outside loop. *Chapter 1*.

cold cap The layer of unmelted feed that floats on the glass pool of a melter. *Chapter 11*.

control allocation Dividing one control demand among several actuators by a fixed rule or an order of priority. *Chapter 12*.

control layers The levels of control, from protection and regulation up to supervision, each acting on its own time scale. *Chapter 1*.

control rod A rod of neutron-absorbing material moved into or out of the core to change its reactivity. *Chapter 1*.

control rod drive The mechanism on the reactor vessel head that moves a control rod on automatic or manual signals. *Chapter 1*.

coping time The hours for which a passive heat sink can take the decay heat before its water is spent. *Chapter 1*.

deadband A band around a setpoint inside which a device makes no correction. *Chapter 2*.

decay heat Heat that the fission products go on making after a reactor stops, a few percent of full power at first. *Chapter 1*.

delayed neutron fraction The share of all fission neutrons that are delayed, about 650 pcm for uranium-235. *Chapter 1*.

delayed neutrons The small share of fission neutrons released seconds to a minute late by decaying fission products; they set how fast power can change. *Chapter 1*.

density-wave oscillation An oscillation of power and flow in a boiling core, carried by waves of steam bubbles that travel up the channels with a delay. *Chapter 6*.

detachment A state in which the plasma edge is cooled so much by radiation that it no longer reaches the divertor targets at full heat. *Chapter 1*.

differential pressure level transmitter An instrument that infers a level from the difference in pressure between a tap low on a vessel and one high on it. *Chapter 1*.

divertor The part of a fusion device where exhaust heat and particles are guided onto cooled targets. *Chapter 1*.

dollar Reactivity measured in units of the delayed neutron fraction; one dollar makes a reactor prompt critical. *Chapter 1*.

Doppler coefficient The change in reactivity per degree of fuel temperature; it is negative, so hotter fuel slows the chain reaction. *Chapter 1*.

dynamic matrix control (DMC) A form of model predictive control that predicts the plant from its measured step responses. *Chapter 8*.

eddy currents Currents induced in a conducting wall by a changing magnetic field; they slow the vertical motion of a plasma. *Chapter 1*.

elongation How much taller than wide a tokamak plasma is drawn; more elongation makes its vertical instability faster. *Chapter 12*.

external reset Feeding the applied output back to each controller's integral, so that controllers not in command do not wind up. *Chapter 11*.

fuel salt A molten salt with the reactor fuel dissolved in it. *Chapter 7*.

gain ceiling The highest gain that a loop on an unstable plant tolerates, set by its delay, before it oscillates. *Chapter 1*.

gain floor The lowest gain at which a loop can hold an unstable plant at all. *Chapter 1*.

growth rate How fast an unstable motion grows; its inverse is the time in which the motion multiplies by about 2.7. *Chapter 1*.

gyrotron A microwave tube that heats a plasma by beaming power at its electrons. *Chapter 13*.

heat exhaust The heating power that leaves a plasma and must be taken away by the divertor and the walls. *Chapter 1*.

impurity seeding Puffing a gas such as neon into the plasma edge so that it radiates heat before the heat reaches the targets. *Chapter 1*.

insurge Water flowing into the pressurizer from the primary circuit when the coolant expands. *Chapter 3*.

inverse response A first move of a measurement in the direction opposite to where it finally settles. *Chapter 1*.

ion chamber A gas-filled detector whose current measures the neutron flux, and so the reactor's power. *Chapter 1*.

island divertor A stellarator divertor that uses a chain of magnetic islands at the plasma edge to guide exhaust onto its targets. *Chapter 13*.

isolation condenser A heat exchanger in a water pool that condenses reactor steam and returns the water by natural circulation. *Chapter 10*.

joule heating Heating by passing an electric current through a material, here the molten glass between electrodes. *Chapter 11*.

latency The delay between a measurement and the action it causes. *Chapter 1*.

load following Changing a plant's power to follow the demand for electricity. *Chapter 2*.

load rejection A sudden loss of electrical load, so the steam the turbine was taking must go somewhere else. *Chapter 3*.

loop pairing Choosing which manipulated variable controls which controlled variable in a plant with several loops. *Chapter 6*.

mid-ranging Using a slow actuator to bring a fast one back toward the middle of its range. *Chapter 9*.

model predictive control (MPC) A controller that predicts the plant over a horizon and chooses its moves by solving an optimization at every step. *Chapter 1*.

moderator A material that slows neutrons so they cause fission more readily; water in this volume's water reactors. *Chapter 1*.

moderator temperature coefficient The change in reactivity per degree of coolant temperature; negative in the water reactors of this volume. *Chapter 1*.

molten salt reactor (MSR) A reactor whose fuel is dissolved in a molten salt that a pump carries through the core. *Chapter 1*.

natural circulation Flow driven by the difference in density between hot and cold fluid, with no pump. *Chapter 1*.

neutron detector An instrument that counts or measures neutrons, inside or outside the reactor vessel, as a measure of power. *Chapter 1*.

non-condensable gas Gas such as air that does not condense and blankets the tubes of a condenser. *Chapter 10*.

off-gas train The coolers, scrubbers and fan that draw and clean the gases leaving a melter. *Chapter 11*.

outsurge Water flowing out of the pressurizer when the primary coolant shrinks. *Chapter 3*.

override control A structure in which a second controller takes a final element from the first when a limit comes near. *Chapter 1*.

passive safety system A safety system that works by gravity, natural circulation or stored energy, without pumps or operator action. *Chapter 10*.

pcm A unit of reactivity equal to one hundred-thousandth, from the Latin per cent mille. *Chapter 1*.

peaking factor The ratio of the highest heat flux on a target to the average over the area it wets. *Chapter 1*.

point kinetics A model that treats a reactor core as one neutron population plus groups of delayed neutron precursors. *Chapter 1*.

power-flow map The chart of reactor power against core flow on which a boiling water reactor must stay. *Chapter 6*.

power-operated relief valve (PORV) A pressurizer valve that opens on high pressure to release steam and must close again when the pressure falls. *Chapter 3*.

pressurized water reactor (PWR) A reactor whose cooling water is kept liquid under high pressure and passes its heat to steam generators. *Chapter 1*.

pressurizer The vessel whose single steam bubble sets the pressure of the whole primary circuit of a pressurized water reactor. *Chapter 1*.

prompt critical A reactor critical on prompt neutrons alone, at a reactivity of one dollar, where the delayed neutrons no longer set the pace. *Chapter 1*.

prompt jump The almost instant change of power after a step of reactivity, before the delayed

neutrons follow. *Chapter 1*.

radiated fraction The share of a plasma's heating that it radiates instead of sending to the divertor. *Chapter 1*.

radiative collapse The end of a plasma whose core radiates more power than it absorbs. *Chapter 1*.

reactivity The fractional excess of neutron production over loss; power rises while it is positive and falls while it is negative. *Chapter 1*.

reactivity ledger The sum of the rod, fuel, coolant and xenon contributions to a reactor's reactivity. *Chapter 1*.

reactor trip The fast insertion of all control rods that stops the chain reaction, also called a scram. *Chapter 1*.

recirculation pump A pump that drives water through the core of a boiling water reactor; its speed changes the reactor's power. *Chapter 6*.

recoverable region The largest displacement of an unstable plant from which full actuator effort can still bring it back. *Chapter 1*.

resistance temperature detector (RTD) A thermometer whose metal winding changes its electrical resistance with temperature. *Chapter 1*.

right-half-plane zero A zero of a transfer function that makes a response start the wrong way and limits how fast any loop can be. *Chapter 1*.

runback A fast automatic cut of reactor power when a limit comes near. *Chapter 1*.

safety valve A valve that lifts by itself at a set pressure to protect a vessel from overpressure. *Chapter 1*.

self-regulation A plant returning toward balance after a disturbance with no controller acting. *Chapter 1*.

shrink and swell The level of a boiling vessel first falling when the feed rises (shrink) or rising when the steam demand rises (swell). *Chapter 1*.

small modular reactor A reactor of modest power, often planned with several units sharing one site. *Chapter 9*.

soft sensor A model that estimates a variable that cannot be measured directly from measurements that can. *Chapter 11*.

split-range control One controller output that drives two actuators, each over its own part of the range. *Chapter 3*.

spray valve The valve that sprays cooler water into the pressurizer steam to lower the pressure. *Chapter 3*.

steam dump Valves that send steam past the turbine to the condenser or the air when the turbine cannot take it. *Chapter 5*.

steam generator A heat exchanger in which the heat of the primary coolant boils the water of a second circuit. *Chapter 1*.

stellarator A fusion device that makes its twisted magnetic field with outside coils alone, so it can run for long pulses. *Chapter 1*.

temperature program A setpoint for the average coolant temperature that changes with the load. *Chapter 2*.

thermocouple A temperature sensor made from a junction of two different metals that gives a small voltage. *Chapter 1*.

thermowell A closed tube in a pipe or vessel that holds a thermometer, protects it and slows its response. *Chapter 1*.

three-element control Level control that adds the steam flow and the feed flow to the level, so the level loop only trims. *Chapter 1*.

tokamak A ring-shaped fusion device that drives a current in its plasma to help confine it. *Chapter 1*.

turbine impulse pressure The steam pressure just inside the turbine, a fast measure of how much steam the turbine takes. *Chapter 5*.

turbine trip The sudden closing of a turbine's steam valves, as when its generator leaves the grid. *Chapter 5*.

velocity actuator An actuator that takes a speed rather than a position, so it adds up the controller's output over time. *Chapter 2*.

vertical displacement event A loss of vertical control in which the plasma drifts into the wall. *Chapter 12*.

vertical instability The tendency of an elongated plasma to move up or down ever faster unless coils push it back. *Chapter 1*.

vitrification Melting waste together with glass-forming materials into a glass. *Chapter 11*.

void feedback The change in reactivity as steam bubbles form in a boiling core; more bubbles slow the chain reaction. *Chapter 6*.

xenon poisoning The absorption of neutrons by xenon-135, which builds up for hours after a cut in power. *Chapter 1*.

Index of Control Topics

actuator asymmetry, 63
actuator saturation, 321
  loss of process gain, 221
AI agents
  causal claims, 374
  engineering practice, 371
  open problems, 376
  reproduction, 374
  truth audit, 373
  verification record, 373
  while studying, 18

back-calculation, 35, 115, 213, 302, 325
bubbler level measurement, 312
bumpless transfer, 16, 93, 156, 242, 352

cascade control, 35, 92, 174, 352
  inner loop speed, 99
constraint control, 261, 303
  margin, 368
consumable resource, 261
control allocation, 326
  priority, 326
control layers, 12
control rod drive, 12
controller action, 93, 157, 213, 302, 325

data export, 17
dead time, 302, 355
  gain limit, 189
deadband, 34
decay heat, 10, 118
defense in depth, 12
delayed neutrons, 7
  circulating fuel, 174, 208
  effective fraction, 179
density-wave instability, 145
  regional mode, 168
detachment control
  open problems, 359
disturbance model, 95
drum level control, 81
dynamic matrix control, 204, 214
dynamic testing
  in service, 223

effective delayed neutron fraction, 209
energy ledger, 349
external reset, 303, 352

feedforward
  fast disturbance, 367
  preview, 214
feedforward control, 92, 112, 157
  dynamic compensation, 174
  lead-lag, 125
  model inversion, 303
  power mismatch, 35
final control element, 12
flow measurement
  molten salt, 199
flow scheme, 4
fundamental limitations, 113

gain
  nonlinear actuator, 63
gain ceiling, 321, 330
gain floor, 321, 329
gain margin
  measurement lag, 366
gain scheduling, 33, 158, 185, 272
gap controller, 34
growth rate, 324

heat exhaust, 11, 342
hidden failure, 70

impurity seeding, 11, 343
infeasible setpoint, 354
inferential measurement, 293
instrument calibration
  channel check, 254
integral action
  on an integrating actuator, 37
integral windup, 354
  back-calculation, 373
integrating process, 298
  in the actuator, 36
integrator windup, 65
interlock, 65, 93, 120
inverse response, 8, 81, 143, 184

Kalman filter, 17, 244
  disturbance states, 66
  steady state, 95, 159

steady-state, 304

latency, 325
level measurement, 12
limit cycle, 330
deadband and minimum step, 40
load allocation, 229
margin-proportional, 241
spill pass, 243
load following
molten salt reactor, 198
loop action
direct acting, 185
loop interaction, 53
fueling and seeding, 358
shared inventory, 66
loop pairing, 142, 156

maximum controllable displacement
model error, 337
noise, 336
measurement lag, 31, 121
measurement noise, 41, 190, 330
measurement packet, 349
mid-ranging, 242
model mismatch, 353
model predictive control, 94, 158, 243, 304, 343
against simpler control, 369
explicit, 326
fallback, 353
in the apps, 16
margin constraint, 187
two priced actuators, 66
move blocking, 353
moving horizon estimation, 17, 314
multi-loop control, 290
multivariable coupling, 142

neutron detection
high temperature, 223
neutron detector, 12
non-equilibrium model, 57
nuclear plants
control design procedure, 368
open problems, 374

observer, 326
offset
position actuator, 186
offset-free control, 328, 343
override control, 9, 47, 65, 112, 274, 290, 342
nuclear plants, 368
published cap, 125
testing, 373

pairing, 212
manipulated variable selection, 174
passive system
reliability, 284
performance measures
resource remaining, 284
phase margin, 37
PID control
anti-windup, 16, 65, 93, 156, 186, 241, 274, 343
derivative filter, 41
velocity form, 16, 64, 92, 125, 156, 241, 274, 302, 325
PID tuning
SIMC rule, 309
plantwide control, 229
one integrator per controlled variable, 242
point kinetics, 5, 148
positive feedback, 293
preview control, 204
process gain, 33, 124
variation with valve position, 272
process monitoring
harsh environment, 314
prompt jump approximation, 5, 29, 117, 179
protection limits, 9

reactivity feedback, 5, 153
coefficient sum, 193
reactivity ledger, 30
recoverable region, 10, 320
recoverable margin, 324
reinforcement learning
plasma control, 371
relative gain array, 67, 204, 212
right-half-plane zero, 83, 143
runback, 243

scenarios, 17, 305
completion rules, 17, 328
graded challenge, 69, 133, 159, 308, 332, 355
laboratory sequence, 67, 127
scorecard, 17, 43
selective control, 53
low selector, 65, 115

selector, 343
  low, 352
  min/max, 292
selector control, 274
self-regulation, 24, 37, 144, 188, 216, 267, 366
sensor bias, 191
setpoint program, 213
setpoint programming, 55, 124
shrink and swell, 8, 86, 121, 154
  load pickup, 91
simulation engine, 15
simulator framework, 15
single-element control, 84, 145
soft constraint, 353
soft sensor, 298
  cold-cap coverage, 314
split-range control, 53
  asymmetric actuators, 53
  deadband, 64
step test
  open loop, 90
stored heat, 272
surveillance testing, 70
testing
  protective layers, 376
three-element control, 9, 84, 145
  state feedback, 85
throughput manipulator, 301
time-scale separation, 290, 350
transport delay, 174
two-element control, 84

ultimate gain, 36
unstable plant, 10, 320
unstable steady state, 300

valve characteristic
  equal percentage, 59
valve position control, 156, 242, 308
velocity actuator, 24, 126
velocity estimation, 326
void feedback, 142

xenon poisoning, 8, 42, 183

zone control, 215

Index of Applications

air heater bypass, 211
aircraft nuclear propulsion, 204
Aircraft Reactor Experiment, 176, 205
    circulating fuel, 207
    operation, 176
atmospheric relief valve, 116
average coolant temperature limit, 240

β drop, 324
boiler master, 231
boiling water reactor, 142
    early experiments, 144
bolometer, 358
    three-dimensional radiation, 359
boron, 31
    end of cycle, 47
borosilicate glass, 292
bubbler, 292
    installation, 296

chemical and volume control system, 56
circulating fuel, 174, 208
cold cap, 291
    plenum temperature, 313
compressor discharge temperature, 210
condenser
    back pressure, 120
condenser bypass, 236
condenser pool, 265
control bank, 28
control rod
    regulating rod, 222
control rod drive, 24
    magnetic jack, 46
    wear, 46
coping time, 269
coronal cooling rate, 348

decay heat, 235
    Way-Wigner correlation, 265
decay-heat removal, 10, 261
Defense Waste Processing Facility, 292
    melter life, 293
delayed neutron precursors, 179
density limit, 345
detachment, 343

differential-pressure level transmitter, 84, 145
    reference leg, 166
    wet leg, 105
divertor, 11, 343
    detachment control, 371
    target element, 349
drift-flux correlation, 147
drum boiler, 81

eddy currents, 324
edge localized mode, 324
elongation, 321
equal-percentage valve, 87
exclusion region, 145

feedwater valve, 87
    positioner, 106
flashing, 58
flexible operation
    level control, 104
    steam bypass, 232
    turbine bypass, 137
foaming, 293
force balance, 323
freeze plug, 196
fuel pump
    variable-speed drive, 198
fuel salt, 177
Fukushima Daiichi, 264
    instrumentation, 369
    isolation condenser, 265
fusion, 1

gauge glass, 106
go-around, 219
graphite moderator, 177
gyrotron, 358

heat exchanger
    counterflow effectiveness, 210
    fouling, 192
Heat Transfer Reactor Experiment, 207
heater banks
    contactor hysteresis, 59
heater uncovery, 61
house load operation, 136

immersion heater, 56
impurity seeding
  first island divertor experiments, 345
in-vessel coils, 324
industrial steam user, 236
inherent load following, 216
instrument, 12
instrumentation and control
  digital upgrade, 47
insurge, 68
integral pressurized water reactor, 232
ion chamber
  outside the vessel, 221
island divertor, 344
isolation condenser, 265
  standby, 283
  tube rupture, 264
ISS04 confinement scaling, 345
ITER, 323
  in-vessel coils, 322
  plasma control, 371

jet pump, 145
joule heating
  glass resistivity, 297
  power supply, 313
joule-heated melter, 290

LaSalle Unit 2, 145
level trip, 90
load following, 37, 144, 174, 204
  history, 27
  nuclear, 369
load rejection, 68, 112

magnetic diagnostics, 336
main steam safety valve, 116, 235
maximum controllable displacement, 323
minor disruption, 324
moderator temperature coefficient, 43
molten salt reactor, 7, 174
  aircraft, 204
Molten Salt Reactor Experiment, 177
  outlet temperature control, 184
motor-operated valve, 283

natural circulation, 85, 267
non-condensable gas, 268
nuclear power, 1
  milestones, 3
nuclear waste
  high-level, 290

off-gas train, 293
once-through steam generator, 231
online monitoring
  instrument drift, 370
outsurge, 68

P-4 permissive, 115
parallel boilers, 230
passive safety system, 261
peaking factor, 348
power supply
  vertical stabilization, 336
power-flow map, 145
power-operated relief valve, 55
  stuck open, 62
pressure transmitter, 73
  calibration bias, 237
  differential capacitance, 253
pressurized water reactor, 4, 112
  pressurizer, 53
  rod control, 24
pressurizer, 53
  fatigue monitoring, 74
  level program, 67
  steam bubble, 55
pump trip, 191

radiating layer, 348
radiative collapse, 343
reactivity coefficients
  adaptive rod control, 370
reactor runback, 126
reactor trip, 9, 121, 158
  decay heat after, 261
recirculation flow control, 144
  flexible operation, 167
recirculation pump
  variable-frequency drive, 167
recirculation pump trip
  1988 event, 145
resistance temperature detector, 45
  thermowell lag, 133
rigid massless plasma model, 323
rod insertion limit, 42
rod worth, 30

saturated steam properties, 88
scram, 235
scrubber loading, 297
self-regulation
  history, 3
shutdown margin, 42

small modular reactor, 229
  autonomous control, 370
  multi-unit dispatch, 255
spray valve, 56
  split range, 74
steam drum, 8, 85
steam dump, 112
steam dump valve
  positioner, 135
steam generator, 28, 81, 116
  low-power level control, 107
steam header, 229
stellarator, 11, 342
  long pulses, 371
surge line, 56
  thermal stratification, 75

temperature program, 31, 124
thermal energy storage, 236
thermocouple, 283
  in molten salt, 196
thermowell, 31
Three Mile Island, 56
  emergency feedwater, 82
  instrumentation, 4
  relief valve, 56
tokamak, 10, 320
turbine bypass, 114
turbine bypass valve
  positioner, 254
turbine impulse pressure, 124
turbine trip
  P-9 permissive, 115
turbojet, 210
two-point model, 345

vertical displacement event, 321
  prediction, 337
vertical instability
  growth rate measurement, 322
vertical stability, 10, 320
vitrification, 290

wall stability margin, 324
Wendelstein 7-X, 343
  history, 345

xenon transient, 8, 31

About the Editor

John D. Hedengren is a professor in the Department of Chemical and Biological Engineering at Brigham Young University, where he leads the Process Research and Intelligent Systems Modeling (PRISM) group. He received the B.S. and M.S. degrees in chemical engineering from Brigham Young University and the Ph.D. degree from the University of Texas at Austin in 2005. He developed advanced process control applications for polymer production at ExxonMobil Chemical Company and commissioned nonlinear model predictive controllers with PAS, Inc. before joining the faculty at Brigham Young University in 2011.

His teaching covers process dynamics and control, engineering computing, data science, machine learning, and dynamic optimization, with course material published online for use inside and outside the classroom. He developed the Temperature Control Lab, a portable heat transfer device now used by many universities for process control and data science instruction, and he maintains the APMonitor and GEKKO optimization software used in those courses. He serves as Chair of Control Education for the IEEE Control Systems Society and as a trustee of the CACHE Corporation. He received the AIChE David Himmelblau Award in 2014, the AIChE Computing Practice Award in 2018, and the John R. Ragazzini Award in 2024, and he was inducted into the Process Automation Hall of Fame in 2025.

Volumes in This Series

Interactive Control Engineering: Principles and Applications has 20 volumes. Nine are organized by control technology and eleven by industry or application area. Every chapter in every volume has a simulator at https://apmonitor.com/apps/.

1 **Process Dynamics, Modeling, and Identification.** How plants respond, and how to find the model.
2 **Practical PID Control.** Tuning, anti-windup and implementation that works on real loops.
3 **Feedforward, Cascade, and Ratio Control.** Choosing a structure from where the disturbance enters.
4 **Override, Split-Range, and Constraint Control.** When limits decide what the controller may do.
5 **Delays, Slow Measurements, and State Estimation.** Controlling what arrives late or cannot be measured.
6 **Multivariable and Plantwide Control.** Pairing loops, managing inventories, closing the recycle.
7 **Nonlinear Dynamics, Oscillations, and Stabilization.** Plants that will not stay put without control.
8 **Batch, Run-to-Run, and Hybrid Control.** Recipes, endpoints and learning from the last run.
9 **Model Predictive Control and Economic Optimization.** Planning moves ahead, and choosing where to operate.
10 **Control in Upstream and Midstream Oil and Gas.** From the wellbore to the terminal.
11 **Control in Petroleum Refining.** Throughput, quality and limits from crude unit to blending.
12 **Control in Petrochemicals and Polymers.** Severity, selectivity and grade from cracker to pellet.
13 **Control in Chemical Plants.** Recycle, equilibrium and energy in continuous chemical processes.
14 **Control in Semiconductor Manufacturing.** Run-to-run control where little can be measured in the chamber.
15 **Control in Mining and Mineral Processing.** Variable ore, late assays and equipment at its limit.
16 **Control in Iron, Steel, and Metals Production.** Slow furnaces, fast strip, and measurements that arrive once.
17 **Control of Energy Conversion and Storage.** Fast regulation and slow dispatch on a changing grid.
18 **Control of Nuclear and Fusion Systems.** Reactor physics, protection limits and coordinated control.
19 **Control of Biological Systems and Human Performance.** Variable subjects, sparse samples and limits that cannot be crossed.
20 **Motion Control, Robotics, and Autonomous Systems.** Tracking, navigation, contact and allocation.