INTERACTIVE CONTROL ENGINEERING

Multivariable and Plantwide Control

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 6

Multivariable and Plantwide Control

Pairing loops, managing inventories, closing the recycle

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 6: Multivariable and Plantwide Control
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/v06-multivariable-plantwide/

How to cite this volume. John D. Hedengren, Ed., *Multivariable and Plantwide Control*, vol. 6 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 loops that interact and units that share inventories. Chapters 2 to 6 work inside one unit: which measurement to pair with which valve, when to decouple, and what a weak direction of the plant does to measurement noise. Chapters 7 to 11 look across a plant: where to set throughput, how recycle amplifies a disturbance, how a published plantwide benchmark is held together, and how inventory decisions pass along a supply chain.

Volumes 2 to 4 cover the single loops that are combined here. Volume 13 applies the plantwide ideas to chemical plants with recycle.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Pairing, Interaction, and Plantwide Structure 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 2
- 1.3 Notation and Conventions . . . 3
- 1.4 Structure Theory on One Plant . . . 3
- 1.5 Instruments and Final Elements in This Volume . . . 12
- 1.6 The Simulator Framework . . . 14
- 1.7 Working With AI Agents While Studying . . . 17
- 1.8 Chapter Roadmap . . . 18
- 1.9 Exercises . . . 20

2 Seawater Reverse Osmosis Virtual Laboratory to Teach Loop Pairing and Decoupling 22
- 2.1 Introduction . . . 23
- 2.2 Background . . . 24
- 2.3 Process Model . . . 25
- 2.4 Gain Structure and Loop Pairing . . . 28
- 2.5 Controller Design . . . 29
- 2.6 Results . . . 33
- 2.7 Technology in Practice . . . 40
- 2.8 Limitations . . . 42
- 2.9 Discussion . . . 43
- 2.10 Conclusion . . . 44
- 2.11 Exercises . . . 46

3 Paper Machine Dryer Simulator for Control Through a Scanning Quality Gauge 49
- 3.1 Introduction . . . 50
- 3.2 Quality Control on Paper Machines . . . 51
- 3.3 Process Model . . . 52
- 3.4 Dynamic Analysis . . . 59
- 3.5 Controller Design . . . 61
- 3.6 Results . . . 64
- 3.7 Technology in Practice . . . 72
- 3.8 Limitations . . . 74
- 3.9 Discussion . . . 75
- 3.10 Conclusion . . . 75
- 3.11 Exercises . . . 77

4 Vapor-Chain Concentration Control in a Three-Effect Milk Evaporator Simulator 81
- 4.1 Introduction . . . 82
- 4.2 From the Sugar House to the Dairy Train . . . 83
- 4.3 The Evaporator as Modeled . . . 84
- 4.4 The Vapor Chain and the Temperature Budget . . . 89
- 4.5 Implemented Controller . . . 91
- 4.6 Cases and Their Measured Outcomes . . . 94

4.7 Technology in Practice . . . 101
4.8 Limitations . . . 103
4.9 Conclusion . . . 103
4.10 Exercises . . . 105

5 Roll-to-Roll Coating Simulator for Web Tension and Coat Weight During Speed Changes 109
5.1 Introduction . . . 110
5.2 Tension and Coat Weight on Converting Lines . . . 111
5.3 Model of the Coating Line . . . 112
5.4 Implemented Controllers . . . 117
5.5 Three Reduced Models . . . 121
5.6 Scenarios and Measured Results . . . 123
5.7 Technology in Practice . . . 130
5.8 Limitations . . . 132
5.9 Conclusion . . . 133
5.10 Exercises . . . 135

6 Orbit Feedback for a Storage Ring Whose Weak Modes Amplify Noise 138
6.1 Introduction . . . 139
6.2 From Strong Focusing to Feedback in Mode Space . . . 140
6.3 A Declared Ring of 96 Monitors and 120 Correctors . . . 141
6.4 Strong and Weak Directions of the Response . . . 145
6.5 Feedback in Mode Space as Implemented . . . 146
6.6 Ten Cases, Measured . . . 148
6.7 The Ring on the Multivariable Ladder . . . 153
6.8 Technology in Practice . . . 154
6.9 Limitations . . . 156
6.10 Conclusion . . . 156
6.11 Exercises . . . 158

7 Switchable Control Structures and the Snowball Effect in a Reactor–Column Recycle Plant 161
7.1 Introduction . . . 162
7.2 Recycle, Inventory and the Throughput Handle Since 1964 . . . 163
7.3 The Reactor, the Column and the Line as Modeled . . . 164
7.4 One Steady State and Three Handles . . . 168
7.5 The Loops as Coded . . . 171
7.6 Nine Cases on the Running Plant . . . 172
7.7 Technology in Practice . . . 176
7.8 Limitations . . . 178
7.9 Conclusion . . . 178
7.10 Exercises . . . 180

8 A Verified Browser Port of the Tennessee Eastman Plant and Its Shutdown 183
8.1 Introduction . . . 184
8.2 A Challenge Problem From Kingsport . . . 185
8.3 The Public Model and Its Port . . . 186
8.4 How Closely the Port Follows the Code . . . 191
8.5 Control Structures as Implemented . . . 193
8.6 Ten Cases, Measured . . . 195
8.7 Technology in Practice . . . 200

8.8 Limitations . . . 202
8.9 Conclusion . . . 202
8.10 Exercises . . . 204

9 GLP-1 Peptide Purification Where a Chromatography Cut Is Judged at the Filter 207
9.1 Introduction . . . 208
9.2 Peptides on Resin, Peptides on Silica . . . 209
9.3 The Route as Modeled . . . 210
9.4 The Recipe, the Loops and the Three Cuts . . . 215
9.5 Code Verification, and What the Sources Can Check . . . 216
9.6 Nine Cases, a Challenge and Three Cut Policies . . . 217
9.7 Technology in Practice . . . 221
9.8 Limitations . . . 223
9.9 Conclusion . . . 224
9.10 Exercises . . . 225

10 Sealed Orbiting Cabin Simulator Where Added Oxygen Shuts the Nitrogen Valve 229
10.1 Introduction . . . 230
10.2 From Expendable Canisters to Regenerable Beds . . . 231
10.3 The Cabin as Modeled . . . 232
10.4 The Sum, the Saw-Tooth and the Shadow . . . 237
10.5 The Control Layers as Implemented . . . 238
10.6 Ten Cases and a Week, Measured . . . 240
10.7 Technology in Practice . . . 245
10.8 Limitations . . . 247
10.9 Conclusion . . . 247
10.10 Exercises . . . 249

11 Supply-Line Feedback and the Bullwhip Effect in a Perishable Dairy Chain 253
11.1 Introduction . . . 254
11.2 From Servomechanisms to the Beer Game and Fresh Chains . . . 255
11.3 The Chain and Its Ledgers . . . 257
11.4 Ordering Rules on a Decision Calendar . . . 261
11.5 The Rule's Frequency Response and the Floor . . . 264
11.6 Nine Cases on One Chain . . . 266
11.7 Technology in Practice . . . 270
11.8 Limitations . . . 272
11.9 Conclusion . . . 273
11.10 Exercises . . . 274

12 Summary and Outlook 278
12.1 What the Chapters Showed . . . 278
12.2 Choosing a Structure . . . 281
12.3 Where Practice Is Heading . . . 282
12.4 Engineering With AI Agents . . . 284
12.5 Gaps in Multivariable and Plantwide Control . . . 287
12.6 Exercises . . . 291

Answers to Selected Exercises 292

Glossary 296

Index of Control Topics 299

Index of Applications **301**

About the Editor **303**

Volumes in This Series **304**

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: Pairing, Interaction, and Plantwide Structure

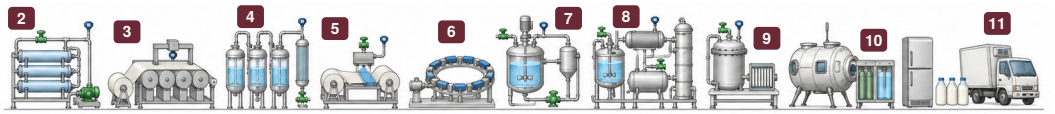

The chapters 2 reverse osmosis train 3 paper machine 4 milk evaporator 5 coating line 6 storage ring 7 recycle plant 8 Tennessee Eastman plant 9 peptide purification 10 orbiting cabin 11 dairy chain

Notice that the recycle plant and the Tennessee Eastman plant each pipe material back to their reactor, while the coating line, the purification bench and the dairy chain pass every change on to the next unit downstream. Illustration, not to scale.

IN PLAIN TERMS **Which Handle Holds Which Target**

The main idea. A plant has several handles, such as valves, pumps and motor speeds, and several things worth holding steady. Before any loop is tuned, someone must decide which handle holds which target, and where the rate of production is set.

The challenge. Each handle moves several targets at once, so two loops can fight or undo each other's work. Streams that loop back, and tanks in a row, can also make a small change grow as it travels.

Why it matters. A wrong choice cannot be fixed by tuning, and the same few tests serve a seawater plant, a particle accelerator and a chain of milk depots.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	7 min	This page, Section 1.2 with its timeline and two History boxes, the Did you know? boxes in Sections 1.4.2, 1.4.6 and 1.4.8, Figure 1.11, and Key points in Section 1.8. No equations.
🕮	**Study**	40 min	The whole chapter in order, with each number of the recycle example worked, then Exercises 1 to 10.
△	**Practice**	5 min	This page, Section 1.4.9 with Table 1.2, Section 1.5 with Figures 1.7 and 1.9, and Section 1.8 with Figure 1.11.

1.1 The Question This Volume Answers

A single loop has one decision before tuning, the sign of its action. A plant with several actuators and several variables worth holding has a larger decision, and it is made before any loop is tuned. Which variables receive setpoints, which actuator holds each one, which flow sets the production rate, and which vessels absorb the upsets? A wrong answer cannot be repaired by tuning. A pairing can be unstable for every tuning, and a recycle can turn a small change of feed into a large change of load on the equipment inside it. This volume asks how such a structure is chosen, and which measures computed from a model predict whether it will work.

Ten chapters answer it, each on a plant with an app that runs in a browser. Chapter 2 and Chapter 3 pair two actuators with two product properties and add a static decoupler. Chapter 4 and Chapter 5 hold a quality along a chain of units that share one supply of heat or one line speed. Chapter 6 inverts a response of 96 monitors and 96 slow correctors whose weak directions amplify noise. Chapter 7 and Chapter 8 are recycle plants, where inventory control and the choice of production handle come first. Chapter 9 judges an upstream decision by its effect two units downstream, Chapter 10 holds two gas loops that share one total pressure, and Chapter 11 passes orders along a supply chain whose inventories amplify demand.

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

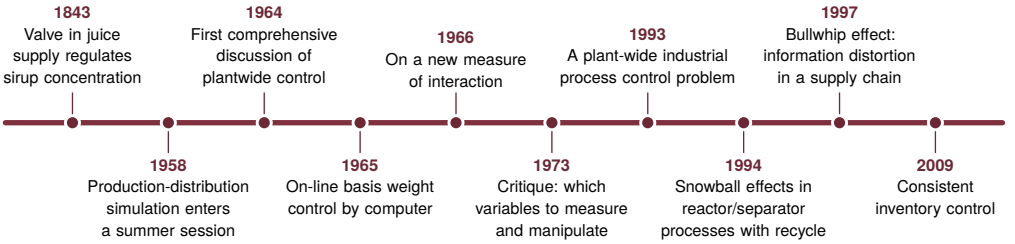

Figure 1.1 Milestones of pairing, interaction and plantwide structure in the systems of this volume. The dates are from [4] (1843), [5] (1958), [2] (1964 and 1973), [6] (1965), [7] (1966), [8] (1993), [9] (1994), [10] (1997) and [11] (2009). Notice that a sugar works, a paper mill and a classroom game sit beside the theory: each met a structural problem before a measure for it existed. Events are evenly spaced, not to a time scale.

- Count the degrees of freedom, and say which of them must stabilize inventories and which remain for production and quality.
- Place the production handle and arrange the inventory loops so that every vessel is held.
- Predict the snowball of a recycle and the handle that avoids it.
- Compute the relative gain array and the Niederlinski index, choose a pairing and reject one that no tuning can save.
- Design a static decoupler and say what it costs.
- Read the singular values of a gain matrix and decide how many directions a controller should try to correct.
- Say where a disturbance goes and whether a chain of inventories amplifies it.

This chapter develops each skill on one small plant and names the chapters that use it. PID tuning is not repeated here. Volume 2 develops it, and every loop below is a PI controller tuned by the rules of that volume. Standard texts treat the multivariable theory [1] and plantwide design [2, 3] in fuller detail.

1.2 A Short History

The milestones of Figure 1.1 follow the structural decisions of this volume through very different systems. A pairing was made long before it had a name. A patent of 1843 regulated the concentration of the sirup leaving the last of a train of sugar evaporators with a differential thermometer that turned a throttle valve in the juice supply, so the concentration was held with the feed [4], the pairing that Chapter 4 sets against the steam.

Interaction was met in practice before it was measured. On the first paper machine under computer control, whose on-line basis weight control first ran in 1965, the first identification experiment showed that the thick stock flow moved both the moisture and the dry basis weight, while the steam moved only the moisture [6]. A measure of interaction for multivariable process control, the relative gain, was proposed in 1966 [7], and the Niederlinski index followed in 1971 [12].

The plantwide questions grew with the plants. In the late 1970s and early 1980s one chemical company moved the production handle of numerous plants away from the process feeds, against a mindset that the feeds must stay constant and with pneumatic hardware that made the decision hard to reverse [13]. A plant-wide industrial process control problem was published in 1993 as a common test [8], the plant of Chapter 8, and in the same year the correction of a storage ring's orbit was written as a singular value decomposition of its response matrix [14]. A paper of 1994 studied snowball effects in reactor and separator processes with recycle [9].

> ⧖ HISTORY **1964: A Chapter on Overall Process Control**
>
> The first comprehensive discussion of plantwide control was a chapter on overall process control in a book of 1964 by Page Buckley. It introduced material balance control in the direction of flow and opposite to it, production rate control, and buffer tanks as low-pass filters. It discussed the recycle and the need to purge impurities, and it pointed out that inventory and flow cannot be controlled independently at one point of a plant, because the material balance relates them. A review of 2000 found the industrial approach still very much along those lines, with relatively little development in the area since [2]. Section 1.4.3 states the same rules on the example.

Supply chains have their own line. A production-distribution simulation with retail, distributor and factory stages was part of a university summer session from 1958 [5], an ancestor of the game behind Chapter 11, and a paper of 1997 described information distortion in a supply chain as the bullwhip effect [10]. Consistent inventory control, the rule of Section 1.4.3, received its own treatment in 2009 [11].

> ⧖ HISTORY **1973: The Structure Before the Theory**
>
> In the 1970s several articles criticized the gap between theory and practice in process control. The most famous, of 1973, observed that in many areas application was ahead of theory. It named the determination of the control system structure as the central issue: which variables should be measured, which inputs should be manipulated and which links should be made between the two sets. The gap was real, it said, but the theoretician must close it. A review of 2000 still found a systematic method for generating promising structures lacking, because the plantwide problem had not been well understood or even acknowledged as important [2]. Section 1.4 works through measures that can screen such structures.

1.3 Notation and Conventions

The chapters share the symbols of the notation page. The setpoint is SP, the process variable is PV and the manipulated variable is MV, with controller settings K_c and τ_I. A plant with inputs u_j and outputs y_i has the steady-state gain matrix K with elements $K_{ij} = \partial y_i / \partial u_j$ and, where dynamics matter, the transfer matrix $G(s)$. A gain carries the units of its output per unit of its input, so a matrix in engineering units mixes several units. A relative gain is written λ_{ij} and the array of them Λ. The singular values of a matrix are $\sigma_1 \geq \sigma_2 \geq \ldots$, and their ratio $\gamma = \sigma_1 / \sigma_{\min}$ is the condition number. The throughput manipulator (TPM) is the flow, speed or setpoint that sets the production rate.

Symbols differ between chapters. Chapter 2 writes the relative gain λ_{11} and Chapter 3 writes Λ_{11}. Chapter 10 uses κ for the cabin's moles per kilopascal, which is why this chapter writes the condition number as γ. Chapter 11 uses λ for the length of a supply line in days and Γ for the ratio of a stage's orders to the orders it receives. Chapter 6 uses λ for the pole of its per-mode loops and μ for the Tikhonov weight. Chapter 7 lists five quantities that each carry the name snowball gain and defines each by its formula. The running example uses hours, kilomoles and kilomoles per hour.

1.4 Structure Theory on One Plant

1.4.1 The Running Example

The running example is a small reaction plant with a recycle, drawn in Figure 1.2. Pure A enters at the fresh feed F_0 and meets the recycle D in a stirred reactor of holdup M, where it reacts to B at the rate kMz, with z the mole fraction of A in the reactor and $k = 1\,\mathrm{h}^{-1}$. The effluent F goes to a separator that is taken as ideal: it returns all of the A as the recycle and sends all of the B out as the product B, and it holds no material of its own. The reactor balances are Equation (1.1), with $n = Mz$ the kilomoles of A in the reactor.

$$\frac{\mathrm{d}M}{\mathrm{d}t} = F_0 + D - F, \qquad \frac{\mathrm{d}n}{\mathrm{d}t} = F_0 + D - Fz - kn, \qquad D = Fz \tag{1.1}$$

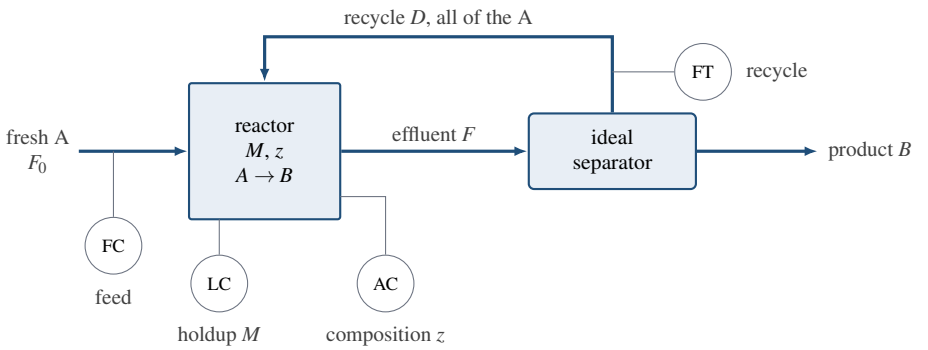

Figure 1.2 The Running Example. Fresh A Meets the Recycle in a Stirred Reactor. An Ideal Separator Returns All of the A and Sends the B Out as Product. The Instruments Measure the Fresh Feed, the Reactor Holdup, the Reactor Composition and the Recycle Flow.

At the design point $F_0 = 200\,\text{kmol/h}$ and $M = 300\,\text{kmol}$. Because every mole of A that leaves the reactor returns, A leaves the plant only by reaction, and the steady state needs $kn = F_0$. The reactor therefore holds $n = F_0/k = 200\,\text{kmol}$ of A whatever its size, which gives $z = 2/3$, a recycle $D = 400\,\text{kmol/h}$, an effluent $F = 600\,\text{kmol/h}$ and a product $B = F_0 = 200\,\text{kmol/h}$. The holdup only dilutes the A with B.

Two simplifications keep the arithmetic exact, and both are stated where they matter. The separator is ideal and holds nothing. The reactor holdup follows its level setpoint as a first-order lag of 1 h, the speed at which the level controller is allowed to move the effluent. Simulations use the velocity-form PI of Section 1.6.2, executed every 0.01 h, with the recycle flow and the reactor composition read through transmitter lags of 0.02 h. The plant is integrated by the fourth-order Runge–Kutta method on a step of 0.001 h. Chapter 7 runs the full version of this plant, with a twenty-tray column, a vapor limit and a recycle line.

1.4.2 Degrees of Freedom

A control structure starts with a count. Each valve, pump speed or heater duty that can be set independently is a dynamic degree of freedom. A real version of the example has seven: the fresh feed, the effluent, the recycle, the product, the separator's reflux and boilup, and the reactor coolant. Some of them must stabilize inventories, because a liquid level on a pumped outlet is an integrator and drifts unless a loop holds it. At steady state, a level whose value changes nothing else, such as the separator's drum and sump, uses up a degree of freedom without giving a choice. The steady-state count is therefore Equation (1.2) [3].

$$N_{ss} = N_{\text{valves}} - N_0 \tag{1.2}$$

where N_0 is the number of inventories with no steady-state effect. Here $N_{ss} = 7 - 2 = 5$, and the five choices are the production rate, the reactor holdup, the reactor temperature and the two purities of the separator. The reactor holdup is a choice because it changes the conversion.

> **+ DID YOU KNOW? Thousands of Loops, a Few Decisions**
>
> A chemical plant may have thousands of measurements and control loops. Plantwide control is not the tuning of each one but the structural decisions about the whole plant [2], and those decisions are in most cases made ad hoc, from experience and engineering insight [3].

The five are then spent. Some go to active constraints, such as a purity held at its sales limit or a temperature held at its upper limit, and the rest go to variables whose constant setpoints keep the

plant near its optimum as disturbances change, which is the idea of self-optimizing control [15]. In the example the reactor temperature and both purities are held, which leaves two handles: the fresh feed F_0 and the holdup M. The rest of this section is about those two.

The count is the first step of most chapters. Chapter 2 has two actuators and six variables worth holding, so two receive setpoints and four follow. Chapter 8 has twelve valves and several inventories that are integrating or unstable. Chapter 7 spends its handles on five level and purity loops and a flow. Chapter 9 shows that choosing the controlled variable is part of the count. Its chromatography cut holds the purity of the pool, while the filter two units later responds to one particular impurity that the purity does not single out.

1.4.3 Inventory Control and the Throughput Manipulator

Every inventory must be held, and the production rate must be set at one place only. These two requirements fix the direction of the level loops. A vessel downstream of the TPM is held by its outflow, in the direction of flow, and a vessel upstream of it by its inflow, against the flow. A structure that obeys the rule for every vessel and every group of vessels is consistent [11, 16]. Placing the TPM is the larger decision [17]. A TPM at the feed sends every change of production through each unit in turn, while a TPM inside the plant changes the units around it first and leaves the vessels upstream to adjust their inflows.

The example has two common structures, with the names Chapter 7 gives them. In CS1 the fresh feed is on flow control and is the TPM. The reactor level then sets the effluent, and the separator's levels set the recycle and the product, all in the direction of flow. In CS2 the effluent is on flow control. The reactor level sets the fresh feed, against the flow, and the production rate is moved by the reactor level setpoint. Both are consistent.

An inconsistent structure is easy to build. With the fresh feed and the product both on flow control, the plant's total holdup integrates $F_0 - B$, and any difference between the two settings fills or empties some vessel. Consistency must also hold for each component. In the example the reaction removes A at the rate kn, so the A inventory regulates itself at $n = F_0/k$. An inert that neither reacts nor leaves with a product has no such exit and accumulates until a purge removes it. That is the pressure of Chapter 8, which trips its interlock within hours with every valve fixed [8], and closes its inventories before anything else.

The chapters place the TPM in several ways. The production index of Chapter 8 and the line speed of Chapter 5 set the flows or the pump in ratio, so one setting moves the whole plant at once. Chapter 3 moves the machine speed to hold a tonnage, and the permeate tank of Chapter 2 cascades slowly onto the production setpoint. Chapter 4 can hold its product solids with the steam or with the feed, and the second choice makes the throughput the MV. Chapter 10 is a closed inventory whose machines can only add gas, so nothing removes an excess except the crew and the leak.

1.4.4 Recycle and the Snowball Effect

With the holdup fixed, the steady state of the example follows from the balances alone. The reactor holds $n = F_0/k$, so $z = F_0/(kM)$, and the recycle that the separator returns is Equation (1.3).

$$D = \frac{F_0 z}{1 - z} = \frac{F_0^2}{kM - F_0}, \qquad \frac{\mathrm{d}\ln D}{\mathrm{d}\ln F_0} = 2 + \frac{F_0}{kM - F_0} \tag{1.3}$$

The denominator is the margin between the reactor's capacity to convert A, $kM = 300\,\mathrm{kmol/h}$, and the feed. At the design point the elasticity is $2 + 200/100 = 4$, so a percent of feed is four percent of recycle, and the slope is 8 kmol/h of recycle per kmol/h of feed. A 10 % step of fresh feed, from 200 to 220 kmol/h, takes the recycle from 400 to 605 kmol/h, 51.25 % more, and the separator's feed from 600 to 825 kmol/h. This is the snowball effect [9]. The recycle grows faster than the production because a reactor of fixed size can convert more A only by running richer in A, and the separator returns that richer effluent's A. At $F_0 = kM$ no steady state exists.

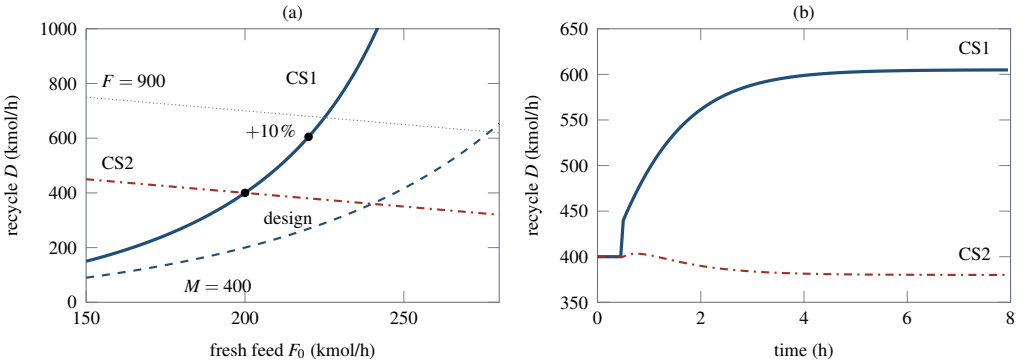

Figure 1.3 The Snowball on the Example. (a) Steady Recycle Against Fresh Feed: CS1 at the Design Holdup (Thick Solid) and at 400 kmol (Dashed), CS2 With the Effluent Fixed (Dash-Dotted) and the Separator Limit $F_0 + D = 900\,\text{kmol/h}$ (Dotted). (b) The Recycle After a 10 % Production Increase at 0.5 h, by the Fresh Feed in CS1 (Thick Solid) and by the Holdup Setpoint in CS2 (Dash-Dotted).

The snowball sets the capacity of the plant. If the separator can take 900 kmol/h, half again its design feed, CS1 reaches that limit at $F_0 = kMF/(kM+F) = 225\,\text{kmol/h}$, only 12.5 % above the design. The holdup is the other handle. In CS2 the effluent stays at 600 kmol/h, so $D = 600 - F_0$ and the recycle falls as production rises. The holdup must grow instead: 10 % more production needs $M = 347.37\,\text{kmol}$, 15.79 % more, and a vessel limited to 400 kmol allows 240 kmol/h. With both the vessel and the separator at their limits the plant makes 276.92 kmol/h, which is why the analysis of this plant in the literature places the reactor at its largest holdup [18]. Placing one flow inside the recycle loop on flow control, as CS2 does, is a general rule [19]. It does not remove the snowball. It moves it from the recycle into the reactor holdup. Figure 1.3(a) draws both structures.

A recycle also changes the dynamics, because it is a feedback path inside the plant [20]. Without the recycle the A in the reactor would settle with the time constant $M/(F+kM) = 0.333\,\text{h}$. With every mole of A returned, A can leave only by reaction, and the time constant becomes $1/k = 1\,\text{h}$, three times longer. A separator that holds some of the recycled material lengthens it further. With a lag of 0.1 h in the separator the slow time constant of the linearized plant is 1.73 h. Figure 1.3(b) shows the recycle after the production step. In CS1 it jumps from 400 to 440 kmol/h as the extra feed passes through, then climbs toward 605 kmol/h and covers 63 % of the change in 0.95 h. In CS2 it falls slowly toward 380 kmol/h, at the pace the holdup is allowed to grow.

Chapter 7 measures the snowball on its column, finds the vapor limit at 10.7 % more feed and switches among three structures during a run. The recycle of Chapter 8 carries gas and an inert, and its pressure is held by the purge.

1.4.5 The Relative Gain Array and Pairing

Suppose the operator of the example wants two things held. The recycle D is the load on the separator, and the reactor composition z sets the selectivity of a real reaction. The handles are the fresh feed F_0 and the holdup M. Differentiating the steady state at the design point gives the gain matrix of Equation (1.4), with z in mol %.

$$\begin{bmatrix}\Delta D\\ \Delta z\end{bmatrix} = \begin{bmatrix}8 & -4\\ 1/3 & -2/9\end{bmatrix}\begin{bmatrix}\Delta F_0\\ \Delta M\end{bmatrix}, \qquad G(s) = \frac{1}{s+1}\begin{bmatrix}2s+8 & -(2s+4)\\ 1/3 & -2/9\end{bmatrix} \tag{1.4}$$

The transfer matrix carries the dynamics. Extra feed reaches the recycle at once, and the rest follows the A inventory with its 1 h time constant. A larger holdup takes effluent while it fills, which lowers the recycle at once.

The relative gain of a pairing compares the gain of a loop with the other loop open against its gain with the other loop closed perfectly, as in Equation (1.5) [7].

$$\lambda_{11} = \frac{(\partial y_1/\partial u_1)_{u_2}}{(\partial y_1/\partial u_1)_{y_2}} = \frac{1}{1 - \dfrac{K_{12}K_{21}}{K_{11}K_{22}}}, \qquad \Lambda = K \circ \left(K^{-1}\right)^{\mathsf{T}} \tag{1.5}$$

where $\circ$ multiplies element by element. Each row and column of Λ sums to one, and Λ does not depend on the units. For the example $\lambda_{11} = 1/[1-(-4)(1/3)/(8\cdot(-2/9))] = 1/(1-3/4) = 4$, so

$$\Lambda = \begin{bmatrix} 4 & -3 \\ -3 & 4 \end{bmatrix} \tag{1.6}$$

Pairing 1 holds the recycle with the feed and the composition with the holdup, on the diagonal, with $\lambda = 4$. Each loop then has four times the gain alone that it has with the other loop closed, so closing the second loop makes the first one slower. Pairing 2 swaps them and has $\lambda = -3$. A negative relative gain means that the steady gain of a loop changes sign when the other loop is closed. With the recycle held by the holdup, more feed makes the holdup grow so much that the reactor becomes leaner, while with the holdup fixed more feed makes it richer.

The consequence is stated by the Niederlinski index [12], Equation (1.7), for a pairing whose loops have integral action and are each stable alone.

$$\mathrm{NI} = \frac{\det K}{\prod_i K_{ii}} \tag{1.7}$$

with the inputs ordered so that each output's paired input is on the diagonal. A negative index means the closed loop is unstable whatever the tuning. For two loops $\mathrm{NI} = 1/\lambda_{11}$, which is 0.25 for pairing 1 and $-1/3$ for pairing 2. The usual rules follow: pair on positive relative gains, prefer values near one, and treat values much larger than one as a sign of sensitivity to model error [21].

Table 1.1 and Figure 1.4 test the prediction. Every loop is tuned alone, with the other in Manual, by the IMC rule of Volume 2 on its own steady gain, a time constant of 1 h and $\tau_c = 0.5$ h, with $\tau_I = 1$ h. The recycle setpoint steps from 400 to 440 kmol/h at 0.5 h with the composition setpoint unchanged. The recycle loop of pairing 1 alone reaches 90 % of the step in 1.30 h, and the uncontrolled composition drifts up by 1.58 points. With both loops closed the recycle takes 5.69 h, 4.4 times as long, and its IAE rises from 18.0 to 72.1 kmol, the factor of about four that $\lambda = 4$ predicts. Pairing 2 runs away. The composition rises, the fresh feed falls to correct it, the recycle falls, and the holdup falls to restore the recycle, which raises the composition further. By 8 h the feed is down from 200 to 38 kmol/h, the holdup setpoint has reached its 60 kmol limit at 7.67 h, and the recycle is 216 kmol/h below its setpoint. Each loop of pairing 2 is stable alone. The recycle loop on the holdup alone settles with an IAE of 17.2 kmol.

The relative gain can also say nothing is wrong when something is. A triangular gain matrix, with one zero off the diagonal, has $\lambda_{11} = 1$ however large the other off-diagonal gain is. The coupling then runs one way. Chapter 2 reports $\lambda_{11} = 1.00005$ on a train whose pump moves the recovery strongly, and Chapter 10 has an oxygen loop and a total pressure loop coupled through one sum, where added oxygen raises the total and the nitrogen does not move the oxygen. Chapter 3 has two-way coupling with $\Lambda_{11} = 1.517$ and a swapped pairing at -0.517 that its run cannot hold. Chapter 8 runs a decentralized structure of seventeen PI loops, and Chapter 4 compares the steam and the feed as the handle on its product solids.

1.4.6 Decoupling

A decoupler adds to each actuator a correction computed from the other controller's output, so that each controller sees only its own loop. The simplified static form of Equation (1.8) needs only two

Table 1.1 Pairings and Decouplers on the Example. Left: the Recycle Setpoint Steps From 400 to 440 kmol/h. Right: the Catalyst Activity Falls 10 %. Both at 0.5 h, Runs of 8 h. t_{90} Is the Time to 90 % of the Step, IAE Is That of the Recycle and $|\Delta z|$ the Largest Composition Error.

	Recycle setpoint step			Activity loss	
Structure	t_{90} (h)	IAE (kmol)	$\|\Delta z\|$ (pts)	IAE (kmol)	$\|\Delta z\|$ (pts)
Pairing 1, recycle loop alone	1.30	18.0	1.58	61.1	2.73
Pairing 1, both loops	5.69	72.1	0.77	1.5	1.74
Both loops, decoupler	3.81	78.6	0.00	247.1	3.51
Both loops, decoupler, retuned	1.40	39.1	0.00	155.6	1.76
Both loops, half decoupler	1.62	25.4	1.11	69.7	2.39
Pairing 2, both loops	never	414.9	7.25	221.3	3.80

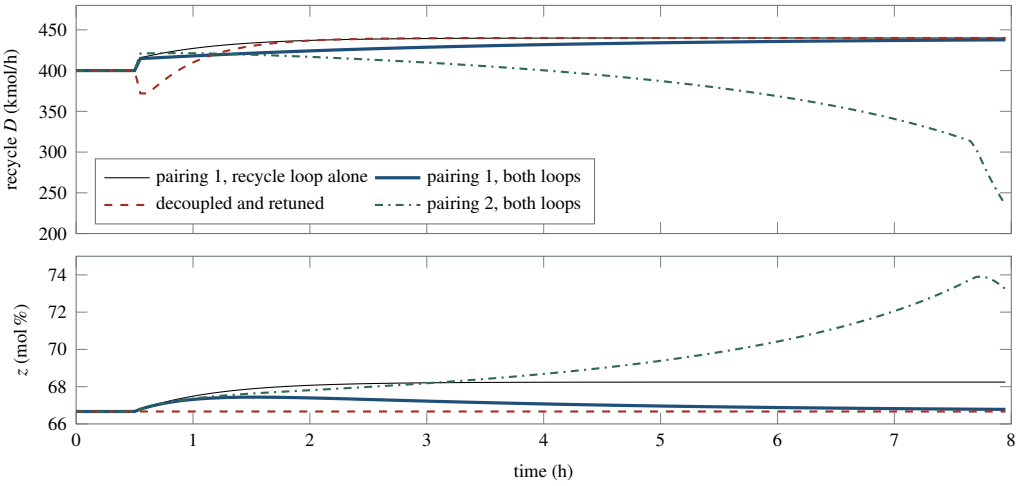

Figure 1.4 The Recycle Setpoint Step of Table 1.1. Thin Solid: Pairing 1, Composition Loop in Manual. Thick Solid: Pairing 1, Both Loops. Dashed: Pairing 1 With the Decoupler, Both Loops Retuned, Whose Recycle First Moves the Wrong Way. Dash-Dotted: Pairing 2, Which Drifts Away From Both Setpoints.

gain ratios [22].

$$u_1 = c_1 + d_{12}\,(c_2 - \bar{u}_2), \qquad u_2 = c_2 + d_{21}\,(c_1 - \bar{u}_1), \qquad d_{12} = -\frac{K_{12}}{K_{11}}, \quad d_{21} = -\frac{K_{21}}{K_{22}} \tag{1.8}$$

where c_i are the controller outputs and $\bar{u}_i$ the design values of the actuators. For the example $d_{12} = 0.5$ kmol/h of feed per kmol of holdup and $d_{21} = 1.5$ kmol of holdup per kmol/h of feed. At steady state each controller then sees the gain K_{ii}/λ_{ii}, which is 2 kmol/h per kmol/h for the recycle loop and $-1/18$ mol % per kmol for the composition loop. These are the gains each loop had with the other loop closed. A decoupler removes the interaction, not the loss of gain, so each loop must be tuned on the smaller gain.

The decoupler is static, and the two paths it balances do not have the same dynamics. In the example both paths to the composition share the 1 h time constant, so the composition is decoupled exactly at every instant. The paths to the recycle differ. The feed acts at once and the holdup's correction acts at once in the opposite direction, so the decoupled recycle loop sees $(2-s)/(s+1)$, a gain of 2 with an inverse response. Table 1.1 shows the result. With the old tunings the composition no longer moves at all, but the recycle reaches 90 % after 3.81 h. Retuned on the decoupled gains, with the inverse response counted as 0.5 h of dead time in the SIMC rule, the recycle loop takes

1.40 h and the composition still does not move. Half the decoupler, the term that moves the feed with the holdup, gives the recycle a t_{90} of 1.62 h and leaves the composition an error of 1.11 points.

The disturbance columns tell the other half. A 10 % loss of catalyst activity is cancelled exactly by raising the holdup to 333.3 kmol, because k and M enter the plant only as the product kM. The undecoupled pairing does exactly that, since its composition loop moves the holdup alone, and leaves the recycle an IAE of 1.5 kmol. Each decoupler tested here adds a feed move to that holdup move and so disturbs the recycle, from 69.7 kmol with half the decoupler to 247.1 kmol with the whole of it untuned. A decoupler designed for setpoint changes is not neutral to disturbances.

> ✦ DID YOU KNOW? **Retuned, Still Not Neutral**
>
> From the simulator: after the 10 % loss of catalyst activity, the decoupler retuned on its decoupled gains still leaves the recycle an IAE of 155.6 kmol, against 1.5 kmol for the plain pairing, while the largest composition errors are 1.76 and 1.74 points (Table 1.1).

Chapter 3 finds the same split. Its full decoupler cuts a moisture excursion from 1.80 to 0.23 points, one of its two terms carries nearly all of that, and the weight loop becomes slower. Chapter 2 needs only one term, because its plant is triangular, and its sweep of that term finds the least error near the design gain. Chapter 3 also notes that a static decoupler puts the right correction on at the wrong time.

1.4.7 Singular Values and Directionality

A gain matrix acts more strongly in some directions than in others. The singular value decomposition $K = U\Sigma V^{\mathsf{T}}$ finds them. Column j of V is a direction of the inputs, column j of U the direction of the outputs it produces, and σ_j the gain along it. The condition number $\gamma = \sigma_1/\sigma_{\min}$ is the ratio of the strongest to the weakest gain [1].

Singular values depend on units, so the matrix must first be scaled. Scaled by the design values, so that every entry is a percent per percent, the example's gain matrix and its inverse are Equation (1.9).

$$K_r = \begin{bmatrix} 4 & -3 \\ 1 & -1 \end{bmatrix}, \qquad K_r^{-1} = \begin{bmatrix} 1 & -3 \\ 1 & -4 \end{bmatrix}, \qquad \sigma_1 = 5.19, \quad \sigma_2 = 0.193 \tag{1.9}$$

The condition number is 27.0. The strong input direction is $(0.79, -0.61)$, more feed with less holdup, which closes the margin $kM - F_0$ and moves mainly the recycle. The weak direction is $(0.61, 0.79)$, feed and holdup together, which keeps that margin and moves mainly the composition, as Figure 1.5 shows. The inverse says the same in whole numbers. A percent more recycle at constant composition needs one percent more feed and one percent more holdup, a move of length 1.41. A percent more composition at constant recycle needs three percent less feed and four percent less holdup, a move of length 5.

Three consequences follow. A controller that inverts the plant asks for five times the move of the output change in the weak direction, so noise and model error in that direction are amplified five times and actuators saturate sooner. A disturbance that lies along the strong direction is cheap to reject and one along the weak direction is expensive. The catalyst loss lies exactly along a holdup move, which is why the plain pairing rejects it so well. Finally, the condition number depends on the scaling, while the relative gain does not. In engineering units the same matrix has $\gamma = 180$. No choice of units brings it below 13.9, and for two by two plants that floor is fixed by the relative gains, here $7 + \sqrt{48}$, where 7 is the sum of the magnitudes in one row of Λ [21]. A large relative gain means a large condition number in every scaling. The reverse does not hold. The triangular sum of Chapter 10, a matrix with rows $(1, 0)$ and $(1, 1)$, has $\gamma = 2.62$ and a floor of one, and the gains of Chapter 2 have $\gamma = 22.3$ in their units and a floor of 1.01.

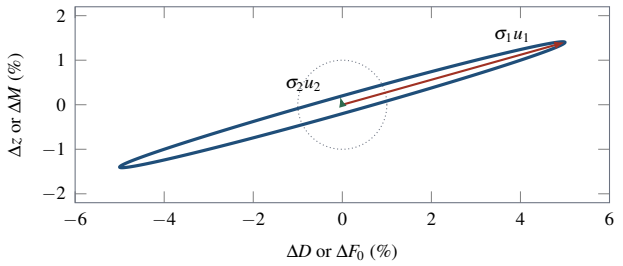

Figure 1.5 Directionality of the Example. Every Input Move of Length One Percent (Dotted Circle) Produces an Output Change on the Thick Ellipse, Whose Half-Axes Are $\sigma_1 = 5.19$ Along u_1 and $\sigma_2 = 0.193$ Along u_2.

Chapter 6 is this subsection at scale. Its 96 by 96 response has $\gamma = 405$, and the chapter chooses how many modes to correct by truncating the inverse or by Tikhonov regularization, trading residual orbit against corrector noise.

1.4.8 Disturbance Propagation

A control system does not remove a disturbance. It moves the variation from the controlled variables to the manipulated ones, and the structure decides which ones. After the 10 % catalyst loss with the holdup and feed fixed, CS1 sends the upset into the recycle, which settles at 571.43 kmol/h, 42.86 % above design. CS2 with the effluent and holdup fixed sends it into production, which falls to 186.21 kmol/h, 6.90 % below. Neither structure is better in general. The choice is which variable the plant can afford to move.

Inventories pass variation from unit to unit. A vessel whose level is held by its outflow passes its inflow to the next unit, filtered by its level loop. With proportional control the outflow follows the inflow as a first-order lag and never exceeds it. With PI control the level returns to its setpoint, and the outflow follows Equation (1.10). Tuning such loops to use the vessel as a buffer is averaging level control [23].

$$\frac{Q_{\text{out}}(s)}{Q_{\text{in}}(s)} = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \qquad \left|\frac{Q_{\text{out}}}{Q_{\text{in}}}(j\omega_n)\right| = \sqrt{1 + \frac{1}{4\zeta^2}} \tag{1.10}$$

where $\omega_n = \sqrt{K_c'/\tau_I}$ and $\zeta = \frac{1}{2}\sqrt{K_c'\tau_I}$, with K_c' the controller gain in outflow per unit of held volume. Its amplitude ratio exceeds one at every frequency below $\sqrt{2}\,\omega_n$, whatever the damping. The peak is 1.47 at $\zeta = 0.5$ and 1.15 at $\zeta = 1$, and three such vessels in series pass 3.16 times a flow variation at $\zeta = 0.5$, as Figure 1.6 shows. The gain at zero frequency is exactly one, because a vessel must in the end pass on what it receives.

> **+ DID YOU KNOW? Less Buffer, More Structure**
>
> Setting production at the feeds and designing the loops unit by unit, in sequence, works for processes with much inventory between the units, including buffer tanks, and little recycle. As plants hold less inventory, the plantwide structure becomes a more important part of the design [13].

A chain whose every stage amplifies some frequency is not string stable, and a disturbance grows as it travels [24]. A supply chain is the same structure with orders in place of outflows, and the growth of order variation upstream is the bullwhip effect [10]. The recycle is a third path. In the example a slow change of feed reaches the recycle four times enlarged in relative terms, while at 1 rad/h the factor is 2.9 and at 10 rad/h it is 1.07, because the recycle's own dynamics take an hour.

Chapter 11 computes the amplitude ratio of each ordering rule, finds a peak of 1.89 for its human rule and no rule that is string stable, and measures the variance growing stage by stage. Chapter 4

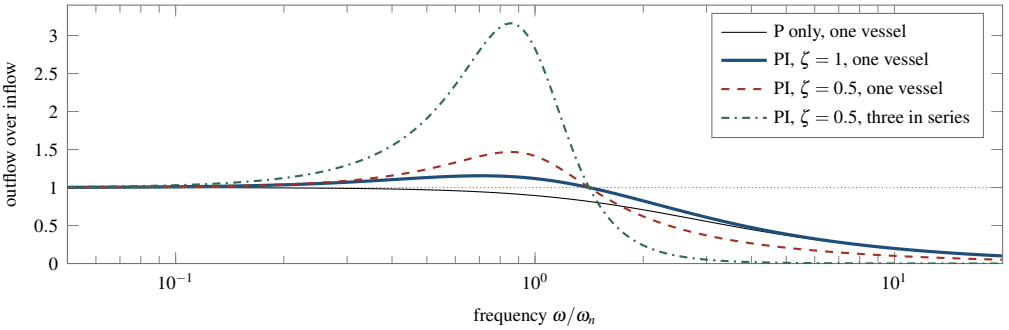

Figure 1.6 How Level Loops Pass Flow Variation Downstream. Thin Solid: Proportional Control, With the Same Gain as the PI at $\zeta = 1$. Thick Solid and Dashed: PI at $\zeta = 1$ and 0.5. Dash-Dotted: Three PI Vessels at $\zeta = 0.5$ in Series.

Table 1.2 The Structural Question of Each Chapter and the Measure It Reports

Ch.	Structural question	Measure reported
2	Which actuator holds production and which recovery	Relative gain of a triangular plant, one-way decoupler gain
3	Pairing of weight and moisture, and which half of a decoupler	Relative gain of two-way coupling, off-specification tonnage
4	Steam or feed on the product solids, tight or averaging levels	Steady cost of each handle, flow variation passed to the dryer
5	Line speed as the throughput handle, with feedforward to tension and coating	Scrap in meters after a speed change
6	How many directions of the inverse to correct	Singular values, condition number, noise gain per mode
7	Which flow in the recycle is fixed and where production is set	Snowball gain, vapor margin
8	Inventory loops before quality, one handle for production	Time to shutdown, operating cost
9	Which variable the upstream cut should hold	Late impurity in the pool against filter time
10	Two loops on one sum in a closed inventory	Valve opening as the signal of a leak
11	The ordering rule at each stage of a chain	Peak amplitude ratio, variance ratio by stage

compares tight and averaging level loops by the flow they pass to the dryer, and Chapter 8 finds that a stripper level held tight passes the disturbance on to production. Chapter 5 follows one speed change along two paths, and Chapter 9 follows one cut to the filter.

1.4.9 What Predicts a Structure

The measures of this section answer different questions. The degree-of-freedom count says how many choices exist. The consistency rule says whether every inventory is held. The snowball elasticity says how much a recycle magnifies a production change. The relative gain and the Niederlinski index say whether a pairing works and how much its loops slow each other. The decoupled gains K_{ii}/λ_{ii} say what a decoupler leaves. The singular values say which directions are expensive, and the amplitude ratio of an inventory says whether a chain grows a disturbance. On the example the measures computed from the gains predicted what the simulations then showed. Table 1.2 lists the structural question of each chapter and the measure it reports, as groundwork for the comparison in Chapter 12.

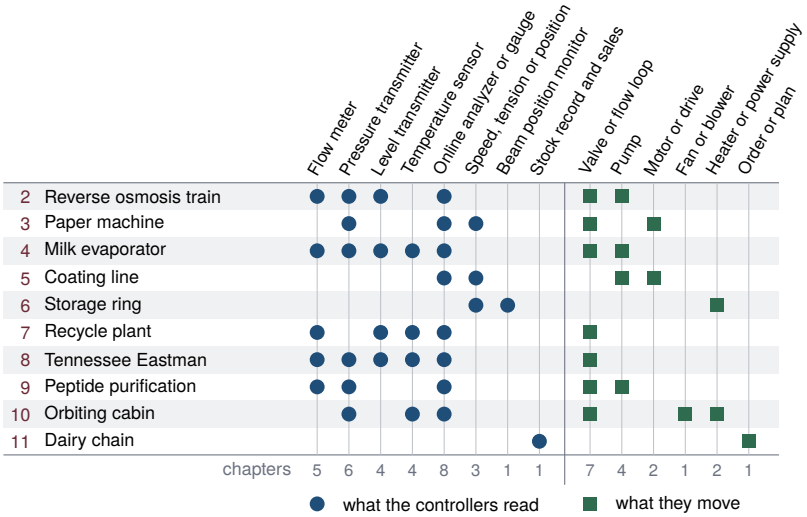

Figure 1.7 What the controllers of each chapter read and move: a round mark where a controller reads the instrument and a square mark where it moves the element, as the chapter's variables table and its technology section name them. A flow set through a flow station counts as a valve or flow loop, and a computed value, such as a recovery or a tonnage, is not marked. Notice that the online analyzers and gauges appear in more chapters than any other instrument.

1.5 Instruments and Final Elements in This Volume

A pairing joins an instrument to a final element, so the devices of a plant bound the structures it can have. Figure 1.7 lists what the controllers of each chapter read and what they move. An online analyzer or gauge appears in eight chapters and a pressure transmitter in six, and a valve or flow loop is the final element in seven. Two chapters stand apart. The controllers of Chapter 6 read beam position monitors and move the supplies of corrector magnets, and those of Chapter 11 read a stock record kept from deliveries and till sales and move an order.

The production handle is usually a flow or a speed. A flow handle sits on a flow loop. A magnetic flowmeter, top of Figure 1.8, reads the voltage that a moving liquid induces across a magnetic field, in proportion to its velocity, and asks only that the liquid's electrical conductivity meet a certain minimum [25]. Chapter 2 reads its recovery as the ratio of two flow measurements. A speed handle moves a machine. Varying the speed of the machine that pushes a fluid is more energy-efficient than throttling its discharge or recycling fluid to its suction, and it removes a control valve from the line [25]. The bottom of the figure shows such a drive on a pump, the arrangement of the high-pressure pump of Chapter 2. The machine speed of Chapter 3 and the line speed of Chapter 5 are handles of the same kind.

The quality measurements are the slow ones. A process chromatograph separates species in time, so cycle times of minutes are not uncommon even online, and some measurement dead time is inevitable [25]. The analyzers of Chapter 8 report every 0.1 h and 0.25 h, the scanning gauge of Chapter 3 posts one average per pass, and the cabin of Chapter 10 hears from its mass spectrometer once a minute. In the first two, the flow, pressure and level loops act many times between two results.

The structures sit in layers. The control system of a chemical plant is usually divided by time scale into scheduling in weeks, site-wide optimization in days, local optimization in hours, supervisory control in minutes and regulatory control in seconds [3], and a survey of predictive control draws basic control executing about every second under constraint control about every minute [26]. A distributed control system mediates a group of localized controllers through a supervisory level, programmable logic controllers can serve as those local controllers, and the loops run with cycle

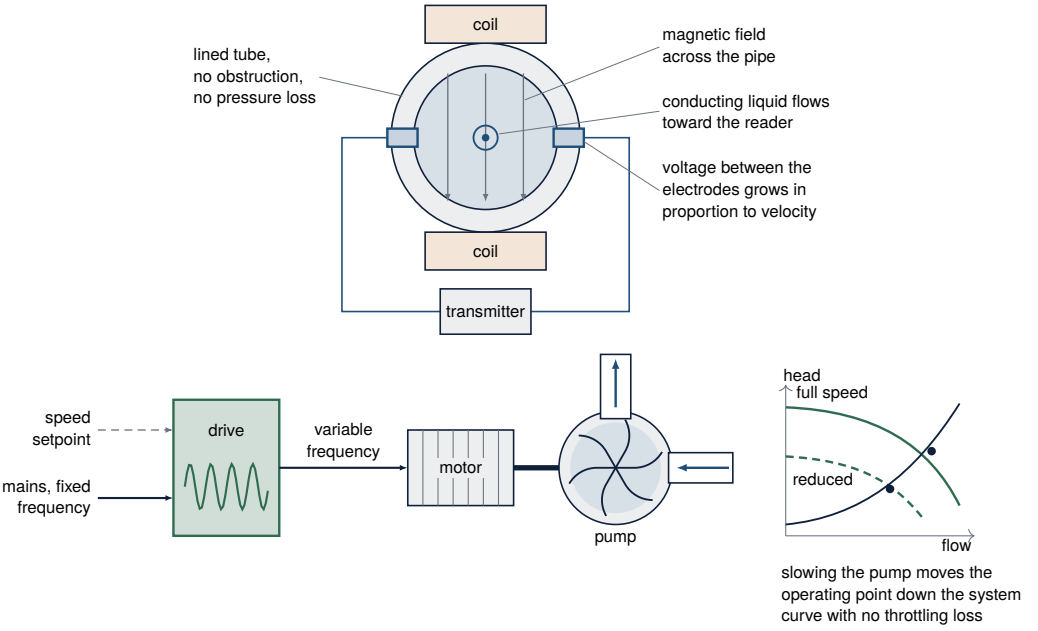

Figure 1.8 Top: principle of an electromagnetic flow meter. Bottom: principle of a variable-speed drive on a centrifugal pump. Notice that the meter reads the stream through electrodes at the wall of the tube, and that the drive moves the pump onto a new curve on the same system curve, so the flow changes with no valve in the line.

Layer	What it does in this volume	Where it runs	How often
Plan or ordering calendar	a power plan for the cabin's beds and oxygen, the orders of a dairy chain	**Local optimization** over the loops; a store's **stock management system**	at each swap or sunrise in the cabin; once a day in the dairy chain
Supervisory structure	a production index with ratio stations, selectors, a decoupler, a predictive controller	**Supervisory level** of a DCS over its localized controllers	0.05 s on the coating line to 0.05 h on the Tennessee Eastman plant
Base loops	PI loops on flows, levels and qualities, a controller per singular mode	**Localized controllers** and **PLCs**; **cell controllers** or arrays of signal processors at a ring	1 ms to 10 ms on the ring and the coating line; 0.5 s to 6 s in the plants; 1 min in the cabin
Sensors and final elements	flow meters, analyzers, gauges, monitors; valves, pumps, drives, magnet supplies	wires, sample lines, a scanning frame, a stock record and a truck	a gauge pass every 20 s; analyzer results every 0.1 h to 0.25 h

Figure 1.9 The control layers of this volume, the platform that hosts each and how often each acts. Platforms are from [27] and [3] and from the chapters' technology sections, and the periods from the chapters. Notice that the supervisory structure of a fast line can act more often than the base loops of a slow plant.

times from milliseconds to minutes [27]. Figure 1.9 places the loops of this volume on their platforms. They run from the millisecond of the fast orbit loop of Chapter 6 to the daily ordering decision of Chapter 11.

Table 1.3 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.

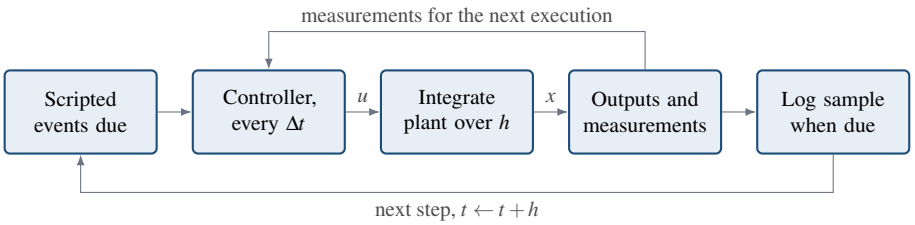

Figure 1.10 Work of the Engine in One Integration Step

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 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

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

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.10. 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 (1.11) to (1.14) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.11}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.12}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.13}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.14}$$

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 (1.14) drives I toward the value that reproduces the output applied, which is back-calculation [28]. With no integral action, Equation (1.14) 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 (1.15).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.15}$$

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 (1.16) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^\top Q\,(y_k - r_k) + \sum_j \Delta u_j^\top R\,\Delta u_j + \sum_i \rho_i s_i^2 \tag{1.16}$$

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 [29]. 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 [30]. 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 [31].

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 12.

Pairing: which handle, which target	Directions: how many to correct	Inventories: where production is set	Propagation: where an upset goes
② a relative gain of one on a plant that interacts	⑥ an ill-conditioned inverse, modes to keep	④ steam or throughput as the handle	⑤ one speed change on two paths
③ a static decoupler through a scanning gauge		⑦ the snowball and the production handle	⑨ a cut judged two units downstream
⑩ two loops on one sum in a closed cabin		⑧ inventories first, then one handle	⑪ orders that grow along a chain

Figure 1.11 Method map of the volume. Each chapter is placed by the part of Section 1.4 that its structural question in Table 1.2 rests on, and labeled with its lesson, shortened from Table 1.4. The placement is the editor's reading of the two. Notice that three of the four questions draw three chapters each, and that only the storage ring turns on directions.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The structural lesson	Level	Time
2	Seawater reverse osmosis train	A pairing chosen against a gain that changes sign. A relative gain of one on a plant that interacts.	introductory	3 h
3	Paper machine with a scanning gauge	Pairing and a static decoupler through a sampled, averaged measurement.	intermediate	3.5 h
4	Three-effect milk evaporator	Steam or throughput as the handle, and averaging levels along a chain.	intermediate	2.5 h
5	Roll-to-roll coating line	One speed change on two paths, held by feedforward and compensation.	intermediate	2.5 h
6	Storage ring orbit	An ill-conditioned inverse and how many modes to correct.	advanced	2 h
7	Reactor, column and recycle	The snowball and the choice of the production handle.	introductory	2 h
8	Tennessee Eastman plant	Inventories first, then one handle for production, on a benchmark.	advanced	2.5 h
9	Peptide purification route	An upstream decision judged by a unit two steps downstream.	intermediate	2 h
10	Sealed orbiting cabin	Two loops on one sum, and a closed inventory that can only be added to.	intermediate	2 h
11	Perishable dairy chain	Inventories in series and the growth of variation upstream.	intermediate	2.5 h

1.8 Chapter Roadmap

Table 1.4 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. Figure 1.11 places the same chapters by the part of Section 1.4 that each one's structural question rests on.

Chapters 2 and 3 form a first course in pairing and are best read in order, with Sections 1.4.5 and 1.4.6. Chapter 7 then carries Sections 1.4.3 and 1.4.4 to a full column, and Chapter 8 is best read after it. Chapter 6 needs Section 1.4.7 and stands alone otherwise. Chapters 4, 5, 9, 10 and 11 can be read in any order after Section 1.4.8. A reader who wants only the plantwide thread can take Chapters 7, 8, 4 and 11.

KEY POINTS

- The control structure of a plant is chosen before any loop is tuned, and a wrong pairing or a recycle that snowballs cannot be repaired by tuning.
- Every inventory must be held and the production rate set at one place only, which fixes the direction of every level loop.
- A negative relative gain means the closed loop is unstable whatever the tuning, and a decoupler removes the interaction but not the loss of gain.
- A control system does not remove a disturbance. It moves the variation to the manipulated variables, and level loops in a chain can amplify it.

Where else this applies. Recycle amplification and inventory pairing in an ammonia synthesis loop in Volume 13, Chapter 2, degrees of freedom and triangular interaction on a crude tower in Volume 11, Chapter 4, and offset-free predictive control of a boiler and turbine in Volume 9, Chapter 4.

References

[1] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. Wiley, 2005.

[2] T. Larsson and S. Skogestad, "Plantwide control – A review and a new design procedure," *Modeling, Identification and Control: A Norwegian Research Bulletin*, vol. 21, no. 4, pp. 209–240, 2000.

[3] S. Skogestad, "Control structure design for complete chemical plants," *Computers & Chemical Engineering*, vol. 28, no. 1–2, pp. 219–234, 2004.

[4] N. Rillieux, "Improvement in sugar-works," United States Letters Patent No. 3,237, Aug. 1843, patented August 26, 1843; image of the specification and drawings served by the United States Patent and Trademark Office.

[5] I. J. Martinez-Moyano, "History of the beer game," *System Dynamics Review*, vol. 40, no. 2, p. e1767, 2024.

[6] K. J. Åström, "Computer control of a paper machine—an application of linear stochastic control theory," *IBM Journal of Research and Development*, vol. 11, no. 4, pp. 389–405, 1967.

[7] 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.

[8] J. J. Downs and E. F. Vogel, "A plant-wide industrial process control problem," *Computers & Chemical Engineering*, vol. 17, no. 3, pp. 245–255, 1993.

[9] W. L. Luyben, "Snowball effects in reactor/separator processes with recycle," *Industrial & Engineering Chemistry Research*, vol. 33, no. 2, pp. 299–305, 1994.

[10] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[11] E. M. B. Aske and S. Skogestad, "Consistent inventory control," *Industrial & Engineering Chemistry Research*, vol. 48, no. 24, pp. 10 892–10 902, 2009.

[12] A. Niederlinski, "A heuristic approach to the design of linear multivariable interacting control systems," *Automatica*, vol. 7, no. 6, pp. 691–701, 1971.

[13] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[14] Y. Chung, G. Decker, and K. Evans, Jr., "Closed orbit correction using singular value decomposition of the response matrix," in *Proceedings of International Conference on Particle Accelerators*. IEEE, 1993, pp. 2263–2265.

[15] S. Skogestad, "Plantwide control: the search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[16] R. M. Price and C. Georgakis, "Plantwide regulatory control design procedure using a tiered framework," *Industrial & Engineering Chemistry Research*, vol. 32, no. 11, pp. 2693–2705, 1993.

[17] R. M. Price, P. R. Lyman, and C. Georgakis, "Throughput manipulation in plantwide control structures," *Industrial & Engineering Chemistry Research*, vol. 33, no. 5, pp. 1197–1207, 1994.

[18] T. Larsson, M. S. Govatsmark, S. Skogestad, and C. C. Yu, "Control structure selection for reactor, separator, and recycle processes," *Industrial & Engineering Chemistry Research*, vol. 42, no. 6, pp. 1225–1234, 2003.

[19] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, "Plantwide control design procedure," *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

[20] J. Morud and S. Skogestad, "Dynamic behaviour of integrated plants," *Journal of Process Control*, vol. 6, no. 2–3, pp. 145–156, 1996.

[21] P. Grosdidier, M. Morari, and B. R. Holt, "Closed-loop properties from steady-state gain information," *Industrial & Engineering Chemistry Fundamentals*, vol. 24, no. 2, pp. 221–235, 1985.

[22] E. Gagnon, A. Pomerleau, and A. Desbiens, "Simplified, ideal or inverted decoupling?" *ISA Transactions*, vol. 37, no. 4, pp. 265–276, 1998.

[23] K. A. McDonald, T. J. McAvoy, and A. Tits, "Optimal averaging level control," *AIChE Journal*, vol. 32, no. 1, pp. 75–86, 1986.

[24] D. Swaroop and J. K. Hedrick, "String stability of interconnected systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 3, pp. 349–357, 1996.

[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] 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.

[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.

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

[29] 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.

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

[31] 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 · 6 min

A blending tank makes $10\,\mathrm{m^3/h}$ of product at $20\,\mathrm{g/L}$ from a concentrate at $100\,\mathrm{g/L}$, flow F_c, and water, flow F_w. The outputs are the product flow $F = F_c + F_w$ and its concentration c.

(a) Find F_c, F_w and the four steady-state gains of F and c to F_c and F_w. *Checkpoint:* $F_c = 2\,\mathrm{m^3/h}$, $F_w = 8\,\mathrm{m^3/h}$, gains 1 and 1 on F, $+8$ and -2 g/L per $\mathrm{m^3/h}$ on c.

(b) Find the relative gain of pairing F with F_c, and choose a pairing. *Checkpoint:* $\lambda = 0.2$, so pair F with F_w ($\lambda = 0.8$) and c with F_c.

1.2 Concept check · 2 min

A two-by-two pairing has a relative gain of -3. What does that predict when both loops have integral action and each is stable alone?

(a) The loops interact weakly and can be tuned one at a time

(b) The closed loop is unstable for every tuning of the two controllers

(c) The loops are stable but three times slower than alone

(d) The pairing works once a static decoupler is added

1.3 Theory and calculation · 4 min

A plant has the steady-state gain matrix

$$K = \begin{bmatrix} 2 & -1 \\ 1 & -1 \end{bmatrix}.$$

Find the relative gain of the diagonal pairing and the Niederlinski index of both pairings.

1.4 Theory and calculation · 5 min

For the plant of the previous exercise, paired on the diagonal, find the two gains of the simplified static decoupler $u_1 = c_1 + d_{12}(c_2 - \bar{u}_2)$, $u_2 = c_2 + d_{21}(c_1 - \bar{u}_1)$ and the steady gain each controller then sees.

1.5 Concept check · 2 min

What does a simplified static decoupler remove from a pair of interacting loops?

(a) The loss of gain each loop suffers when the other is closed

(b) The effect of every disturbance on both outputs

(c) The steady effect of each controller's moves on the other output

(d) The difference in speed between the two paths it balances

1.6 Theory and calculation · 6 min

In the running example the reactor's capacity is changed to $kM = 250\,\mathrm{kmol/h}$ with the fresh feed at 200 kmol/h.

(a) Find the reactor composition, the recycle and the elasticity of the recycle to the feed. *Checkpoint:* $z = 0.80$, $D = 800\,\mathrm{kmol/h}$, elasticity 6.

(b) Find the recycle after 5 % more feed. *Checkpoint:* 1102.5 kmol/h, 37.8 % more.

1.7 Theory and calculation · 3 min

An evaporator train has a feed valve, three effects whose sumps are each drained by a pump, and a steam valve. The sump levels affect nothing else at steady state. How many steady-state degrees of freedom does the train have, and what are they?

1.8 Concept check · 2 min

A plant has a fresh feed, three vessels in series and a product stream. Which structure cannot hold every inventory?

(a) Feed on flow control, each vessel's level on its outflow

(b) Product on flow control, each vessel's level on its inflow

(c) Second vessel's outflow on flow control, upstream levels on inflows, downstream levels on outflows

(d) Feed and product both on flow control, each level on its outflow

1.9 Theory and calculation · 4 min

Two gas loops share one total pressure, so that one unit of oxygen raises both the oxygen partial pressure and the total by one unit, and one unit of nitrogen raises only the total. Find the relative gain and the singular values of the gain matrix

$$K = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

1.10 Theory and calculation · 5 min

A vessel's level is held by its outflow with a PI controller of gain $K_c' = 2\,\mathrm{h}^{-1}$ (outflow per unit of held volume) and $\tau_I = 0.5\,\mathrm{h}$.

(a) Find ω_n and ζ of the outflow's response to the inflow. *Checkpoint:* $\omega_n = 2\,\mathrm{rad/h}$, $\zeta = 0.5$.

(b) Find the amplitude ratio of outflow to inflow at ω_n, for one vessel and for three identical vessels in series. *Checkpoint:* 1.41 and 2.83.

CHAPTER 2

Seawater Reverse Osmosis Virtual Laboratory to Teach Loop Pairing and Decoupling

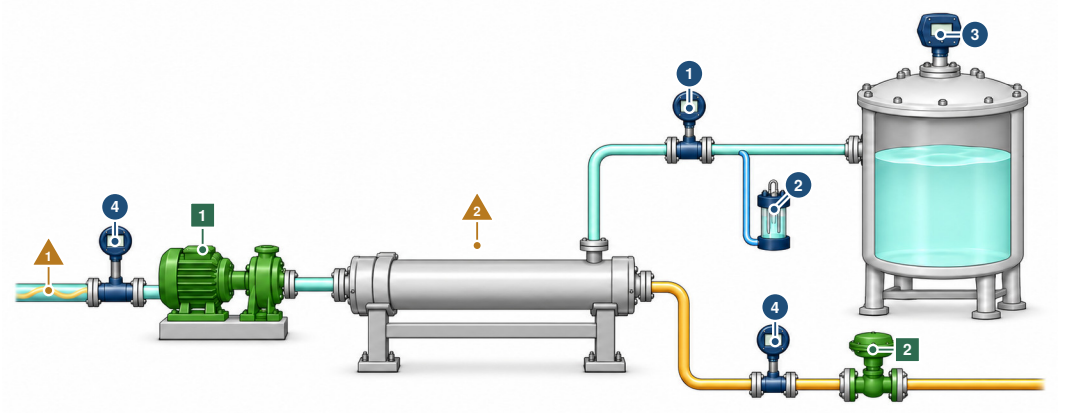

The controller moves 1 high-pressure pump, whose speed sets the push on the membrane
2 concentrate valve, which lets the salty water that is left flow away

The controller reads 1 permeate flow meter: the production 2 conductivity cell: how salty the permeate is
3 level of the permeate tank 4 feed and concentrate flow meters: the share of seawater recovered

What disturbs it 1 seawater whose saltiness and temperature change 2 fouling that builds up on the membrane

Notice that one pump and one valve set every flow, salinity and pressure of the train, so pairing each with a target is the first design decision. Illustration, not to scale.

IN PLAIN TERMS **Two Handles on a Plant That Makes Fresh Water**

The main idea. A seawater plant pushes seawater hard against a thin film that lets water through and holds most of the salt back. A pump sets the push, and a valve lets the salty water that is left flow away.

The challenge. The usual test for whether two control loops will fight says these two will not. Yet each pump move still changes how much seawater turns into fresh water, and giving production to the wrong handle sends the plant far off target.

Why it matters. Which handle holds which target is decided before any tuning, and no tuning can repair a wrong choice.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on energy (Section 2.1), the history in Section 2.2, the two try-its and the fact on fouling (Section 2.6) and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 2.6), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the two open challenges

The control question A seawater reverse osmosis train has two actuators, a high-pressure pump and a concentrate valve, and its relative gain reads 1.00005, the value that means no interaction. Which actuator should hold production, why do the loops still interact, and what keeps the membrane under its pressure rating as it fouls?

Learning objectives

- Derive the permeate flow and salinity at the design point from the solution-diffusion and polarization equations.
- Explain why the valve gain on permeate flow passes through zero and changes sign at the design point.
- Compute the relative gain and the Niederlinski index of both pairings from the gain matrix, and tell a triangular plant from a diagonal one.
- Design the one-way static decoupler from two gains and judge its calibration from a sweep.
- Compare the decoupled loops with the MPC on a pressure limit and name the mechanism that holds it.

Prerequisites Sections 1.4.5 (relative gain and pairing), 1.4.6 (simplified decoupling), 1.4.7 (directionality), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ro-desalination/

See also Volume 7, Chapter 4 for a setpoint at the maximum of the plant characteristic, where the process gain reverses sign. Volume 9, Chapter 4 for offset-free predictive control on a two-by-two plant with a limit, and Volume 9, Chapter 8 for a cost optimum that sits on a moving constraint.

2.1 Introduction

Seawater reverse osmosis recovers fresh water from a saline feed by forcing it through a semipermeable membrane at a pressure above the osmotic pressure of the concentrate, and it now supplies a large share of the municipal water produced from the sea [1]. A train has two manipulated variables. The high-pressure pump sets the pressure applied at the membrane, and the concentrate valve sets the fraction of the feed that leaves without passing through it. Permeate flow, recovery, permeate salinity, feed pressure, specific energy and the margin to scaling are all functions of that pair, so two of the six receive setpoints and four follow. Which two receive them is a design decision taken before any controller is tuned, and it is the one decision that tuning cannot repair.

Sweeping the concentrate valve at a fixed pump speed produces an interior maximum in permeate flow rather than a monotone curve. Permeate flow rises from 438.71 m^3/h at 40 % valve travel to 450.00 m^3/h at 50 % and falls to 436.51 m^3/h at 60 %, and the maximum is the design point. A production controller paired with that valve has no steady-state gain to work against at the operating point and a gain of the opposite sign a few percent to either side. The pump carries no such defect, so production belongs on the pump and recovery on the concentrate valve.

The scenarios of the app cover loop pairing and the relative gain array, one-way static decoupling, gain sign reversal inside the normal operating envelope, override control through a low selector with external reset, cascade control from a surge tank onto a production setpoint, and the drift of every process gain as an asset ages [2, 3]. The physical fundamentals behind the model are a total and a component mass balance over one membrane array, a driving force written as the applied pressure less the osmotic pressure at the membrane wall, and a separation of time scales that runs from a hydraulic node with a three second time constant to a fouling state that moves over weeks [4, 5]. Specific energy is the largest operating cost a desalination plant carries, and the recovery that minimizes it lies close to the concentration at which the sparingly soluble salts crystallize [6], so the distance between an optimum and a constraint is what the control system is defending.

DID YOU KNOW? The Energy Price of Water From the Sea

A research roadmap of the United States Department of Energy gives the energy of seawater desalination in that country as 3.0 to 3.8 kWh/m^3. Supplying potable water in large cities of the country takes 0.34 kWh/m^3 on average, in a study the same roadmap cites [7].

Assessments of process control education call for problems in which the control structure is decided rather than supplied [8, 9]. Interaction is usually demonstrated on plants whose relative gain array announces it, which omits the case in which the standard tool reports no interaction and the loops interact anyway.

This chapter works with a seawater reverse osmosis train on which the pairing decision, its failure mode and its remedy are all measurable. Section 2.2 reviews desalination practice and reverse osmosis control. Section 2.3 gives the model and every constant it uses. Section 2.4 derives the gain matrix, its relative gain and the one-way coupling that the relative gain cannot show. Section 2.5 describes the two regulatory loops, the static decoupler, the supervisory layer and an optional model predictive controller (MPC) of both actuators. Section 2.6 reports the scenario results, from the production loop alone and the reversed pairing through decoupling, warm feed, the energy optimum, fouling and the graded shift, and Section 2.6.8 measures which mechanism of the MPC keeps the membrane under its pressure rating. Section 2.7 describes the real instruments, drives and controllers. Sections 2.8 to 2.10 state the limitations, the three results that carry past this plant and the conclusions.

2.2 Background

Reverse osmosis displaced thermal distillation as the dominant seawater desalination technology over three decades, and energy rather than capacity drove the change. A thermal plant pays for a phase change, while a membrane plant pays only for the pressure applied to the water it processes. Energy recovery closed most of the remaining gap. A pressure exchanger returns the concentrate pressure to the feed at high efficiency, so the high-pressure pump raises only the water that becomes permeate [10], and the specific energy of seawater trains has fallen toward the thermodynamic limit as a result [6]. Desalination now supplies water to a large fraction of the world’s coastal population and consumes electricity in proportion [1].

HISTORY 1952: A Federal Program for Salt Water

Federal funding of desalination research in the United States began in 1952, when the Saline Water Act was passed. An Office of Saline Water was created in 1954, and an Office of Water Research and Technology followed in 1974. A federal handbook for planners later judged these programs primarily responsible for developing reverse osmosis, at a time when few water supply problems on land had even been identified in the country, and noted that the same research gave other fields kidney dialysis with membranes. The Bureau of Reclamation served as a testing laboratory for the office, with membrane pilot plants in the field, and federal funding for desalination largely stopped in the early 1980s [11].

Control of reverse osmosis trains has been studied mainly as a multivariable regulation problem. System identification and decentralized control of a small train date to the late nineteen eighties [12], and reviews of the subject group the published work into empirical transfer function models closed with PID loops and nonlinear model-based designs [13]. Nonlinear model-based control has been demonstrated on an experimental unit operating against pressure and flow constraints [14]. The steady-state trade-off that fixes the operating point, between the feed that has to be pressurized and the osmotic pressure of the concentrate, produces an interior minimum in specific energy that is well established for spiral-wound seawater trains [15].

Membrane performance is not stationary. Permeability declines as foulant accumulates, and on a seawater train the decline is fast enough to move every process gain within one cleaning interval [16]. Sparingly soluble salts crystallize on the membrane once the concentrate passes saturation, which

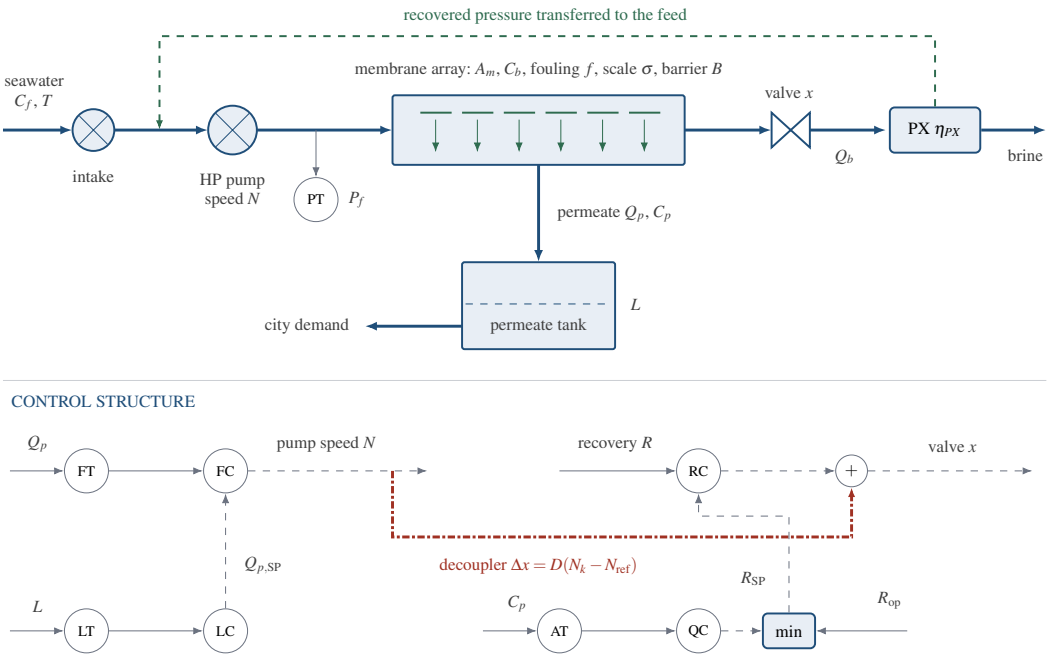

Figure 2.1 The simulated train. Production is controlled on the high-pressure pump (FC) and recovery on the concentrate valve (RC), which is the pairing established in Section 2.4. The permeate tank level cascades onto the production setpoint (LC) and the permeate conductivity low-selects onto the recovery setpoint (QC), so the recovery target R_{SP} is the lower of the operator value and the quality demand and the production target $Q_{p,\mathrm{SP}}$ comes from the tank. The thick dash-dotted path is the one-way static decoupler that carries the pump command directly to the valve, which is the compensation the triangular gain matrix of (2.13) makes possible.

no rinse reverses and which practice controls through antiscalant dosing and through the recovery setpoint [17]. Aging therefore makes the plant a slowly varying system, and a controller commissioned on a clean membrane is running a different plant a month later.

2.3 Process Model

The plant is a single-stage seawater train with a pressure exchanger, sized at 33 000 m^2 of membrane area in one lumped array. Figure 2.1 shows the flowsheet, the instruments and the four controllers.

2.3.1 Working Units and the Design Point

The clock is in minutes, flows in m^3/h, pressures in bar, salinities in g/L and flux in LMH, which is liters per square meter per hour and the unit used on membrane data sheets. At the design point the feed is 35.0 g/L at 25.0 °C, the pump runs at 88 % and the concentrate valve sits at 50 %. That gives 1000 m^3/h of feed, 450.0 m^3/h of permeate at 45.00 % recovery, 62.00 bar of feed pressure, 13.636 LMH of flux, 250.0 mg/L of product and 2.4627 kWh/m^3.

Table 2.1 lists every constant the model uses, so the simulator is reproducible from this chapter alone.

2.3.2 Solution-Diffusion Transport

Water and salt cross the membrane by independent mechanisms [4]. Water crosses in proportion to the applied pressure remaining after the osmotic pressure is subtracted. Salt crosses in proportion to the concentration difference across the membrane and does not depend on pressure,

Table 2.1 Model constants. Every value is the application default.

Symbol	Meaning	Value
A_m	active membrane area	33000 m^2
A_{25}	water permeability at 25 °C	0.6922 LMH/bar
B_{25}	salt permeability at 25 °C	0.06513 LMH
k	polarization mass transfer coefficient	143.11 LMH
$E_{A,w}$	water permeability Arrhenius coefficient	2700 K
$E_{A,s}$	salt permeability Arrhenius coefficient	5100 K
k_π	osmotic coefficient of seawater	0.771 bar/(g/L)
P_p	permeate back pressure	0.5 bar
V_{mem}	feed-channel volume of the array	40 m^3
k_p	pump shut-off head at full speed	90.393 bar
k_Q	pump curve droop	8.0×10^{-6} bar/(m^3/h)^2
C_v	concentrate valve size	393.86 (m^3/h)/bar$^{1/2}$
R_v	concentrate valve rangeability	30
P_{back}	concentrate discharge pressure	1.0 bar
ΔP_d	array pressure drop at design flow	2.50 bar
$n_{\Delta P}$	pressure-drop flow exponent	1.7
$k_{\Delta P}$	pressure-drop increase per unit fouling	3.0
C_{hp}	high-pressure loop capacitance	0.075 m^3/bar
τ_N	pump drive lag	0.10 min
τ_x	concentrate valve travel time	0.15 min
η_{HP}	high-pressure pump efficiency	0.82
η_B	booster pump efficiency	0.78
η_{PX}	pressure-exchanger efficiency	0.96
η_I	intake pump efficiency	0.80
P_{in}	intake and pre-treatment pressure	2.5 bar
k_f	fouling rate at design flux	7.0×10^{-4} min^{-1}
J_0	reference flux for the fouling law	13.636 LMH
Q_{cross}	crossflow below which fouling accelerates	300 m^3/h
k_ϕ	low-crossflow fouling multiplier	2.5
$C_{\text{sat},0}$	saturation salinity, no antiscalant	57.4 g/L
a_s	antiscalant effectiveness	0.45
$D_{1/2}$	antiscalant half-saturation dose	1.5 mg/L
k_σ	scaling rate above saturation	0.15 min^{-1}
k_{ox}	oxidant damage rate	0.020 min^{-1}/(mg/L)
A_t, H_t	permeate tank area and height	60, 6.0 m^2, m
C_{div}, w_{div}	diversion salinity, switch width	600, 25 mg/L
$P_{\max}$	membrane pressure rating	70 bar
$J_{\max}$	design flux limit (alarm only)	17.0 LMH
θ_C	permeate conductivity transport delay	0.45 min
τ_C	permeate conductivity cell response	0.35 min
τ_m	flow, recovery and level transmitter filters	0.05 min
τ_T	intake temperature first-order lag	8.0 min
Δt, Δt_c	integration step, controller scan	0.005, 0.02 min

$$J_w = A\,(\bar{P} - P_p - \pi_w), \qquad J_s = B\,(C_w - C_p). \tag{2.1}$$

That asymmetry governs reverse osmosis quality control. Raising the applied pressure raises the water flux without raising the salt flux, so the permeate becomes fresher, and the fastest correction for an off-specification permeate is more flux rather than less. The permeate composition follows from equating the two fluxes,

$$C_p = \frac{B C_w}{J_w + B}, \qquad Q_p = \frac{J_w A_m}{1000}, \qquad R = \frac{Q_p}{Q_p + Q_b}. \tag{2.2}$$

At design $A = 0.6922\,\text{LMH/bar}$ and $B = 0.065\,13\,\text{LMH}$, which returns 250.0 mg/L of product and 99.286 % rejection.

2.3.3 Concentration Polarization and Temperature

The osmotic pressure that opposes the flux is the one at the membrane wall rather than the one in the bulk. Permeate removal concentrates the layer against the surface, and the two together give

$$C_w = \bar{C}\exp\!\left(\frac{J_w}{k}\right), \qquad \bar{C} = \frac{C_b - C_f}{\ln(C_b/C_f)}, \tag{2.3}$$

with the log-mean standing in for the bulk profile along an array whose ends are C_f and C_b. At design $\bar{C} = 47.82\,\mathrm{g/L}$, the polarization factor is 1.100 and the wall sits at 52.60 g/L, so the wall osmotic pressure is 40.55 bar against a net driving pressure of 19.70 bar. The polarization layer is thickest when the flux is highest, which is a negative feedback that makes the train self-limiting. The layer is also where scale nucleates [17].

Osmotic pressure and both permeabilities carry temperature,

$$\pi = k_\pi C\,\frac{T + 273.15}{298.15}, \quad A(T) = A_{25}\,\mathrm{e}^{E_A\left(\frac{1}{298.15} - \frac{1}{T+273.15}\right)} \tag{2.4}$$

with $E_A = 2700\,\mathrm{K}$ for water and 5100 K for salt. The two activation energies differ, and the salt one is larger, so a warm feed is pressurized more cheaply and produces a saltier permeate while a cold feed costs more pressure and produces a fresher one. Holding both actuators fixed and moving the intake from 16 °C to 32 °C takes the permeate from $400.5\,\mathrm{m^3/h}$ at 159 mg/L to $482.4\,\mathrm{m^3/h}$ at 353 mg/L.

2.3.4 The Hydraulic Node

Feed pressure is the value at which the pump curve, the membrane and the concentrate valve balance,

$$P_{\mathrm{pump}} = k_p\left(\frac{N}{100}\right)^2 - k_Q Q_f^2, \tag{2.5}$$

$$Q_b = C_v R_v^{x/100-1}\sqrt{P_b - P_{\mathrm{back}}} \tag{2.6}$$

with an equal-percentage characteristic of rangeability 30 on the concentrate valve. The node is integrated rather than solved algebraically so that its transient is represented,

$$\frac{\mathrm{d}P_f}{\mathrm{d}t} = \frac{Q_{\mathrm{pump}}(N, P_f) - Q_f}{60\,C_{hp}}, \tag{2.7}$$

$$\frac{\mathrm{d}C_b}{\mathrm{d}t} = \frac{Q_f C_f - Q_b C_b - Q_p C_p}{60\,V_{\mathrm{mem}}}. \tag{2.8}$$

With $C_{hp} = 0.075\,\mathrm{m^3/bar}$ the pressure state has a time constant near three seconds, which is the fastest state in the application and four orders of magnitude faster than the fouling state of (2.10). The array pressure drop is 2.50 bar at design, scales with feed flow to the power 1.7 and is multiplied by $(1 + k_{\Delta P} f)$ as the membrane fouls.

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

2.3.5 Energy Recovery and Specific Energy

A pressure exchanger returns the concentrate pressure to most of the feed, so the high-pressure pump raises only the water that becomes permeate [10],

$$W = \frac{Q_p(P_f - P_{in})}{36\,\eta_{HP}} + \frac{Q_b\left(P_f - \eta_{PX}(P_b - P_{\mathrm{back}})\right)}{36\,\eta_B} + \frac{Q_f P_{in}}{36\,\eta_I}. \tag{2.9}$$

At design the three terms are 907.0 kW, 114.4 kW and 86.8 kW for 2.4627 kWh/m^3. Setting $\eta_{PX} = 0$ in the same expression gives 4.907 kWh/m^3, which is the difference the device makes and is consistent with the reported industrial trend [6].

2.3.6 Fouling, Scaling and Oxidation

Fouling, scaling and oxidation are separate states with separate rate laws,

$$\dot{f} = r_f \, k_f \, \phi_x \left(\frac{J_w}{J_0} \right)^{1.5} (1 - f), \tag{2.10}$$

$$\dot{\sigma} = k_\sigma \left(\frac{C_b}{C_{\text{sat}}} - 1 \right)_+^{2} (1 - \sigma), \tag{2.11}$$

$$\dot{B} = k_{ox} \, c_{\text{ox}}. \tag{2.12}$$

Fouling grows with flux to the three-halves power and is multiplied by $\phi_x = 1 + k_\phi (1 - Q_b / Q_{\text{cross}})$ when the concentrate flow falls below 300 m^3/h, so a starved crossflow ages the membrane faster than a high flux does [16]. The fouling clock r_f is zero unless a scenario starts it. Scaling is exactly zero below saturation and quadratic above it, and no rinse reverses it. Oxidation raises the salt permeability alone, leaves the water permeability untouched and does not reverse. The antiscalant dose D raises the saturation salinity as $C_{\text{sat}} = C_{\text{sat},0} \left[1 + a_s D / (D_{1/2} + D) \right] / h$, where h is the scaling tendency of the feed, so the benefit saturates in dose, from 57.40 g/L untreated to 74.62 g/L at 3 mg/L and only 79.15 g/L at 8 mg/L.

2.3.7 Permeate Tank and Distribution Demand

Permeate above a diversion salinity of 600 mg/L is sent to waste rather than to the tank through a logistic switch of width 25 mg/L, narrow enough to act as a switch and smooth enough for a fixed-step integrator. The tank is a volume and a salt inventory driven by the on-specification fraction and the distribution demand, and it converts a demand set outside the plant into a production setpoint the train can follow.

2.3.8 Initialization

Initialization solves the equilibrium rather than approximating it. Given a pump speed and a valve position, the feed pressure is bisected until the pump curve and the plant agree, and the concentrate salinity is closed by a damped fixed point on the array salt balance. Both converge to well below display resolution, so the initial condition is a true equilibrium and every scenario starts at steady state. The same solver backs the two manual-mode calculators, which return the pump speed for a chosen permeate flow and the valve position for a chosen recovery.

2.3.9 Assumptions

The array is one well-mixed feed channel with a log-mean bulk salinity rather than a discretized series of pressure vessels. Concentration polarization uses a fixed mass-transfer coefficient. Fouling and scaling are scalar states with empirical laws. Salinity is one equivalent salt. Integration is fixed-step RK4 at 0.005 min with a 1.2 s controller scan, so the trends and the exported CSV reproduce exactly on any machine (Section 1.6.1).

2.4 Gain Structure and Loop Pairing

The steady-state gains of the model are obtained by re-solving the equilibrium of Section 2.3 on either side of the design point with a central difference of 0.5 % of actuator travel, a step small enough that halving it changes no tabulated digit. Table 2.2 lists the gain of the concentrate valve on permeate flow across the maximum.

Table 2.2 Gain of the concentrate valve on permeate flow, pump held at 88 %, 35.0 g/L feed at 25.0 °C. The sign reverses at the design point.

Valve x (%)	Recovery (%)	$\partial Q_p/\partial x$ (m^3/h/%)
42	50.9	+1.772
46	48.0	+0.938
50	45.0	−0.001
54	41.9	−1.045
58	38.6	−2.192

Two effects cancel at the design point. Closing the valve raises the feed pressure, which raises the flux through (2.1), and it raises the concentrate salinity, which raises the osmotic pressure at the membrane wall through (2.3) and lowers the flux. At 45.0 % recovery on this train the two are equal to within the resolution of the numerical derivative. A production controller wired to that valve has no gain to work against at the operating point and a reversed plant one step away from it, which is a defect that no tuning repairs [2]. Volume 7, Chapter 4 meets the same defect at the friction peak of a braking tire, where the slip setpoint sits on the maximum of the plant characteristic. The pump gain on permeate flow is 17.147 m^3/h/% at the design point and does not change sign anywhere in the operating envelope.

The relative gain array of Section 1.4.5 is the standard tool for the pairing decision [2, 18]. For the correct pairing on this train the four steady-state gains at the design point are

$$K = \begin{bmatrix} \partial Q_p/\partial N & \partial Q_p/\partial x \\ \partial R/\partial N & \partial R/\partial x \end{bmatrix} = \begin{bmatrix} 17.147 & -0.001 \\ 0.7029 & -0.7687 \end{bmatrix} \tag{2.13}$$

in units of m^3/h/% and %/%, and the array follows from

$$\lambda_{11} = \left(1 - \frac{K_{12}K_{21}}{K_{11}K_{22}}\right)^{-1} = 1.00005. \tag{2.14}$$

A relative gain of one is the textbook signature of a plant with no interaction, and it is returned here by a plant that interacts strongly. Every term in (2.14) is a product of the two off-diagonal gains, so setting either one to zero returns $\lambda_{11} = 1$ no matter how large the other is. The matrix in (2.13) is triangular rather than diagonal, and a triangular plant has coupling that runs in one direction only.

The direction is the useful part. The valve does not move the permeate flow, so closing the recovery loop leaves the production loop undisturbed. The pump moves the recovery by 0.703 %/% of speed, so closing the production loop disturbs the recovery loop every time production changes. Section 2.6 measures both halves of that statement.

Triangularity also carries its own remedy. A disturbance that arrives entirely from one known signal is cancelled from that signal rather than corrected after it is measured, and one static feedforward is sufficient. In the terms of Section 1.4.6 it is the second term of the simplified decoupler alone, and the first term vanishes with K_{12}. A fully interacting plant requires two decouplers that interact with each other, which is a large part of why plants with genuine two-way coupling are given a model-based controller instead [13, 14]. Chapter 3 has such a plant, with a relative gain above one and a decoupler with both terms.

2.5 Controller Design

2.5.1 Velocity Form PID

All four loops share the velocity-form PID of Section 1.6.2, Equations (1.11) to (1.14), with the derivative taken on the measurement [19]. The controller scan interval is $\Delta t_c = 1.2$ s. Anti-windup

is back-calculation applied to the integral alone, with a tracking time T_t equal to τ_I on every loop. Derivative action is provided only on the tank level loop, where the measurement is filtered with a time constant of 0.05 min before it is differenced, and it is zero in every shipped tuning. Tunings follow the SIMC rules [20] applied to gains measured on the app itself.

2.5.2 The Two Regulatory Loops

The production loop takes permeate flow to the high-pressure pump with $K_c = 0.02\,\%/\mathrm{m}^3/\mathrm{h}$ and $\tau_I = 0.4\,\mathrm{min}$. It is reverse acting, because the process gain is positive and the pump command falls as the measured permeate flow rises. The recovery loop takes recovery to the concentrate valve with $K_c = 0.6$ and $\tau_I = 0.6\,\mathrm{min}$. It is direct acting, because opening the valve lowers the recovery and the valve command therefore rises as the measured recovery rises. Neither loop carries appreciable dead time, and the production loop reaches 95 % of a $60\,\mathrm{m}^3/\mathrm{h}$ step in 4.01 min. Each controller is seeded on the actuator it drives in the chosen pairing, and a pairing change in automatic re-seeds it by output tracking, so changing the structure moves no actuator.

2.5.3 Static Decoupler

The pump moves recovery by +0.7029 %/% and the valve by −0.7687 %/%, so a valve move of $0.7029/0.7687 = 0.914$ times the pump move cancels the coupling before the recovery controller sees an error. The same ratio can be measured in manual mode. The calculator that returns the valve position for a chosen recovery, run for 45 % recovery at 87 and 89 % pump speed, gives $-K_{21}/K_{22} = 0.915$ per percent of speed from the two valve positions it reports. The app ships 0.95. The implementation is given by Equation (2.15),

$$\Delta x = D\,(N_k - N_{\mathrm{ref}}), \qquad x_k = \mathrm{clamp}\big(x_k^{\mathrm{PID}} + \Delta x\big) \tag{2.15}$$

where x_k^{PID} is computed with the controller's own output limits shifted by Δx rather than back-calculated against the applied value. Back-calculating against the applied value feeds the feedforward into the integral, which cancels the very correction it makes and leaves a steady offset. The reference speed $N_{\mathrm{ref}} = 88\,\%$ is calibrated on a clean membrane, which is the assumption the decoupler exposes.

2.5.4 Supervisory Layer

The permeate tank level cascades onto the production setpoint with $K_c = 30\,\mathrm{m}^3/\mathrm{h}/\%$ and $\tau_I = 25\,\mathrm{min}$, deliberately slow so that the tank absorbs demand rather than passing it to the membrane. The permeate conductivity drives a second controller whose output is a recovery setpoint, low-selected against the operator target,

$$R_{\mathrm{SP}} = \min\big(R_{\mathrm{op}},\, R_{\mathrm{quality}}\big)\,, \tag{2.16}$$

so the override can only ask for less recovery. The quality controller ($K_c = 0.2\,\%$ of recovery per mg/L, $\tau_I = 12\,\mathrm{min}$) receives an external reset while the operator target wins the selector. Its integral term is set to the applied recovery setpoint in place of its own integration, so its output is that setpoint plus its proportional term. It cannot wind up while it is out of control, and with a constant operator target it takes over without a bump on the first scan at which the measured salinity exceeds the limit [19]. The selected value is never written back into the operator target, so the recovery returns to R_{op} once the quality demand clears.

2.5.5 Model Predictive Control

A separate structure selector can replace the flow and recovery loops, the decoupler and the quality override with one model predictive controller (MPC) for both actuators, the receding-horizon method behind most industrial multivariable control [21], set up as Section 1.6.3 describes in general. Scenario 11 alone loads with the MPC, and the application and its other scenarios start with the loops. Every $\Delta t_m = 0.1\,\mathrm{min}$, which is five scans of Δt_c, the controller solves a quadratic program for the

pump speed and the valve position over $H = 90$ steps (9 min) in twelve move blocks, the shortest a single step near the present and the longest 45 steps at the far end of the horizon, 24 decisions in all. The train receives only the first move, held until the next plan. With the recovery loop open the valve stays where the operator set it and the plan moves the pump alone. The pairing then only chooses which loops would replace a rejected plan, and the level cascade still sets the flow target. Predictive control of reverse osmosis units has been built on identified transfer function models [22], and optimization-based control has been used to minimize the specific energy of a train [23].

The controller's model is the membrane of Section 2.3 rather than an identified two-by-two, because the step responses have two modes. A 1 % pump step raises the permeate flow by 20.55 m^3/h within half a minute and leaves it 17.16 m^3/h higher after the concentrate has salted up, while a 1 % valve step moves it by −3.79 m^3/h at half a minute and by −0.01 m^3/h after twenty. The near-zero valve gain of Table 2.2 is therefore a fast effect and a slow one cancelling, and first-order fits of the eight responses of permeate flow, recovery, feed pressure and permeate salinity to the two actuators leave an RMS residual of 1.0 to 26.3 % of each response. The plan uses Equations (2.1) to (2.6) and the array pressure drop with the constants of Table 2.1 and a design feed of 35 g/L at 25 °C through a new pump, with the dynamics of Equation (2.17),

$$\begin{aligned} \tau_N \frac{\mathrm{d}\hat{N}}{\mathrm{d}t} &= u_N - \hat{N}, \qquad \tau_x \frac{\mathrm{d}\hat{x}}{\mathrm{d}t} = u_x - \hat{x} \\ \frac{\mathrm{d}\hat{C}_b}{\mathrm{d}t} &= \frac{Q_f C_f^0 - Q_b \hat{C}_b - Q_p C_p}{60\, V_{\mathrm{mem}}} \\ \left(Q_p, Q_b, P_f, C_p\right) &= \mathscr{M}\left(\hat{N}, \hat{x}, \hat{C}_b;\ \hat{a}A_{25},\ \hat{b}B_{25}\right) \end{aligned} \tag{2.17}$$

where u_N and u_x are the pump and valve commands, $\hat{N}$, $\hat{x}$ and $\hat{C}_b$ are the model's pump speed, valve position and concentrate salinity, $C_f^0 = 35\,\mathrm{g/L}$ is the design feed salinity, and $\mathscr{M}$ returns the flows, the feed pressure and the permeate salinity at which the pump curve, the membrane and the valve agree, with the declared water and salt permeabilities scaled by the estimates $\hat{a}$ and $\hat{b}$. The feed pressure is treated as settled, because the node of Equation (2.7) has a time constant of 3.2 s. At every plan, central differences about the current estimate give a linear model, discretized over Δt_m with its free motion kept as an affine term. At the design point its steady-state gains are those of Equation (2.13), 17.1473, −0.0010, 0.7029 and −0.7687, its valve gain on permeate flow is +1.772 at 50.9 % and −2.192 at 38.6 % recovery as in Table 2.2, and its salt balance has a time constant of 3.235 min against 3.207 min in the tail of the plant's own pump step. Against ±1 % step tests on the plant at the design point its RMS residual is 0.31 to 2.06 % of each response's range, most of it in the first 0.2 min, while the plant's pressure node is still moving.

Two multipliers and an output bias carry what the membrane has become and remove steady offset [24]. Volume 9, Chapter 4 removes offset with a Kalman estimator that carries a fuel bias instead. At every scan $\hat{a}$ is the measured flux divided by A_{25} and by the net driving pressure that the measured feed pressure leaves over the estimated wall osmotic pressure, and $\hat{b}$ is the conductivity reading divided by the model's permeate salinity passed through the cell's declared 0.45 min delay and 0.35 min response, each filtered over 1 min. The concentrate salinity is integrated from the measured permeate flow, recovery and conductivity. A bias on the permeate flow, the recovery and the feed pressure, filtered over 0.2 min, compares each measurement with the model's output passed through the declared pressure node and transmitter filters. Fouling, scale, temperature, feed salinity and pump wear therefore reach the plan only through these estimates. At the end of scenario 11 $\hat{a} = 0.771$, where 20.89 % fouling leaves 0.791 of the clean permeability, because $\hat{a}$ also absorbs the larger pressure drop of the fouled array. Over the same run the relinearized valve gain on permeate flow moves from +1.476 to −6.802 m^3/h/% as the recovery falls to 33.27 %, crossing the sign change of Table 2.2 without a rejected plan.

Each plan minimizes the objective Φ of Equation (2.18),

$$\begin{aligned}\Phi = &\sum_{k=1}^{H}\left[\left(\hat{Q}_{p,k}-Q_p^*\right)^2+q_R\left(\hat{R}_k-R^*\right)^2\right]\\ &+\mu\sum_{j=1}^{12}\left(s_N\,\Delta u_{N,j}^2+s_x\,\Delta u_{x,j}^2\right)+\sum_{\ell}\left(\rho_\ell\,\varepsilon_\ell^2+\sigma_\ell\,\varepsilon_\ell\right)\\ \text{s.t. }&\hat{P}_{f,k}\le P_{\max}-\delta_P+\varepsilon_P,\quad \hat{C}_{p,k}\le C_{\lim}-\delta_C+\varepsilon_C\\ &\hat{Q}_{p,k}\le J_{\max}A_m/1000-\delta_Q+\varepsilon_Q,\quad \varepsilon_\ell\ge 0\\ &40\le u_{N,j}\le 100,\quad 5\le u_{x,j}\le 100\end{aligned}\tag{2.18}$$

where the hats are predictions at step k, $\Delta u_{N,j}$ and $\Delta u_{x,j}$ are the command changes between consecutive blocks, counted for the first block from the commands now applied, Q_p^* and R^* are the flow and recovery targets after the adjustment described below, $q_R = 4\,w_R$ per %^{2} against 1 per (m^3/h)^2, $s_N = 10$ and $s_x = 1$ per %^{2}, and the move price μ and the recovery weight w_R are both 1 unless the reader changes them. The salinity limit $C_{\lim}$ is the quality limit when level and quality supervision is on and the 500 mg/L specification otherwise, and the flux ceiling $J_{\max}A_m/1000$ is 561 m^3/h. The bands sit $\delta_P = 0.2$ bar, $\delta_C = 2$ mg/L and $\delta_Q = 1$ m^3/h inside the real limits, with prices $\rho_P = 1\times10^5$ per bar^2 and $\sigma_P = 1\times10^4$ per bar on the pressure and 100 and 100 on the salinity and the flow. The shared solver gives each soft limit one slack ε_ℓ for the whole horizon. Where the free response already breaks a limit, that bound starts at the free response and decays to the band with a time constant of 1 min, or 2 min for the salinity, so the program stays feasible. The command bounds are hard and equal the loops' clamps.

The targets are made consistent with the limits before each plan, and this adjustment holds the rating, as Section 2.6.8 measures. At the linear model's steady state the feed pressure and the permeate salinity are linear in the two targets. If holding both would leave less than 0.5 bar under the rating or 5 mg/L under the salinity limit, R^* is lowered until that margin is restored, and Q_p^*, which is also capped 3 m^3/h under the flux ceiling, is lowered only once R^* has reached its 25 % floor or the valve belongs to the operator. Production is what the distribution system draws, and a lower recovery relieves the pressure because a less concentrated channel presents less osmotic pressure at the wall, so recovery is spent first. The steady state used for this step ignores the pump's top speed. The plan carries no price on energy. Two handles holding two targets have no steady freedom left to trade, so a linear price of 1 per kilowatt per step on Equation (2.9), linearized at the design point, only pulls the targets off, and over scenario 8 it raises the flow error from 12.4 to 678.6 m^3/h min for 0.19 % less specific energy.

The controller accepts a plan only if the solver's status is optimal, which is checked even when the solve returns success, every returned array has its shape, the cost, the slacks and the predictions are finite, no slack is negative, all planned commands respect their bounds and the first of them equals the move applied. Otherwise the loops of the chosen pairing take over until a later plan is accepted. Under the MPC those loops follow the commands the plan applies on every scan, without preload, so a rejection forced at 150 min of scenario 11, while the plan holds the pressure under the rating, moves the valve by 0.0007 % on the hand-over scan. No run reported here rejected a plan.

The plant integrates its eleven states with the run's own parameters, the instruments add their own filters, delay and noise, and the plan holds only the declared constants and the design feed, so a change a reader makes to the membrane, the pump or the feed reaches the controller as model mismatch. In a replay test, a fresh controller outside the simulator receives the measurements of 150 min of scenario 11, recorded on a plant with a different hidden membrane, pump, valve and feed, and reproduces all 7500 applied commands bit for bit. In the run measured for this chapter a plan takes 0.403 ms on average in Chrome, a load of 0.40 % of a processor core when the simulation advances one minute per wall-clock second and 8.1 % at its fastest rate of 20 min per second.

2.6 Results

Twelve scenarios run on one plant, each with the briefing, completion rule and data export of Section 1.6.4. Every number below is read from the app at the stated time with the disturbances at their scenario values. The measurement noise generators are off in every shipped scenario, so each run is deterministic and repeats exactly. Each case is selected, any setting named in the text is changed, and the case is run from $t = 0$.

2.6.1 Production Loop Alone

With the recovery loop in manual and the valve pinned at 50 %, a production step from 450 to 510 m^3/h at 10 min in scenario 3 reaches and stays within 5 % of the step after 4.01 min. Recovery, which is uncontrolled in this configuration, drifts from 45.00 to 47.30 % because the pump raises permeate and feed by different fractions. The saturation index rises from 0.850 to 0.887. On this feed the excursion is harmless. On a harder feed or at a lower antiscalant dose it consumes the whole scaling margin, and it arrives as a side effect of a production change that the production loop reports as a success.

2.6.2 The Reversed Pairing

Scenario 5 assigns production to the concentrate valve and recovery to the pump, which is the arrangement a valve-controls-flow reading of the flowsheet suggests. Its production controller is given a gain 15 times larger than the pump loop's, which is the authority the valve has away from the maximum.

The scenario starts in manual and transfers both loops to automatic at 5 min. The transfer is bumpless, because each controller has tracked its own actuator in manual, and nothing moves until the setpoint steps from 450 to 500 m^3/h at 10 min. The flow controller has no gain to work against at the maximum of Table 2.2, so it integrates, and it closes the valve, which is the right direction only on the opening side of the maximum. Moving down the closing side of the curve raises the recovery, the recovery controller answers by slowing the pump, the permeate flow falls further and the two loops drive each other. The valve reaches its 5 % stop 3.9 min after the step, the recovery peaks at 76.3 %, and the concentrate passes saturation (a saturation index of 1.14) long enough to cost 4.07 % of permeability to scale. At 120 min the pump has been dragged from 88 to 68.39 %, feed pressure has collapsed from 62.00 to 42.01 bar and the train is making 82.5 m^3/h of 1241 mg/L permeate against a 500 m^3/h setpoint. The integrated absolute flow error over those two hours is 42 118 m^3/h min. The state at 90 min is shown in Figure 2.3.

Nothing else was at fault. Both loops are stable alone, the tuning is reasonable for the gain the valve has away from the design point, the transfer is bumpless, and neither loop carries dead time. Only the pairing is wrong.

> **TRY IT IN THE APP** **5 · The Pairing You Would Have Guessed**
>
> Load the scenario and press Run. Both loops switch to automatic at 5 min, and the production setpoint steps from 450 to 500 m^3/h at 10 min. Watch the concentrate valve trend. The valve runs to its 5 % stop 3.9 min after the step, the recovery loop slows the pump, and the permeate flow falls instead of rising.

2.6.3 Effect of the Static Decoupler

Table 2.3 carries the central measurement of this chapter. The correct pairing, whose relative gain array reads 1.000 05, deviates 0.937 points in recovery and requires the concentrate valve to travel from 50.00 to 53.44 % to bring it back, taking 8.4 min to return within 0.05 of a point. That is the interaction (2.14) reports as absent. Figure 2.2 draws the table on logarithmic axes.

Adding (2.15) reduces the peak by a factor of 4.4 and the integral error by a factor of 7.2, and it does so with a single multiplication in the control scan. Measured over the 50 min following the step in scenario 6, rather than across the paired step and return of scenario 4, the peak falls from 0.937

Table 2.3 Response of the recovery loop to the 70 m^3/h production step of scenario 4, with the pairing chosen before the run. The peak deviation is taken over the 135 min following the step and the integral error over the whole 180 min scenario. The swapped row uses the scenario 5 flow-loop tuning ($K_c = 0.30\,\%/m^3/h$, $\tau_I = 2.5$ min). It starts without a bump at the design point and leaves it with no step at all, so 885.1 of its 3585.0 % min accrue before the step.

Structure	Peak $\lvert\Delta R\rvert$ (% pts)	Recovery IAE (% min)	Final Q_p (m^3/h)
Direct pairing	0.937	6.87	450.0
Direct plus decoupler	0.213	0.95	450.0
Swapped pairing	24.35	3585.0	450.0

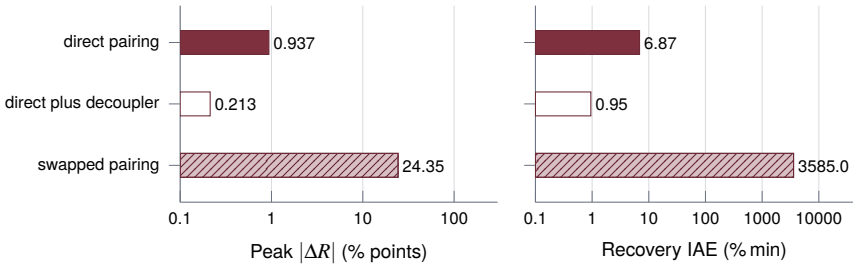

Figure 2.2 The recovery loop on the production step of scenario 4, from Table 2.3, on logarithmic axes: the peak deviation (left) and the integral error (right). The dark bar is the direct pairing, the white bar the direct pairing with the decoupler and the hatched bar the swapped pairing. Notice that the decoupler cuts both measures of an interaction that a relative gain of 1.000 05 reports as absent, while the swapped pairing is worse by orders of magnitude.

Table 2.4 Decoupler gain sweep on the scenario 6 production step, measured over the 50 min following the step. Each gain is set before the run. The design value from the measured gain ratio is 0.914.

Decoupler gain D	0.00	0.50	0.95	1.30	1.60
Peak $\lvert\Delta R\rvert$ (pts)	0.937	0.452	0.160	0.504	0.864
Recovery IAE (% min)	3.44	1.38	0.48	1.92	3.15

to 0.160 points. The swapped row fails without being disturbed. With the recovery loop closed on the pump, opening the valve raises the permeate flow by about $K_{12} - K_{11}K_{22}/K_{21} = 18.75\,m^3/h/\%$, because the pump speeds up to hold the recovery, which is the opposite of the sign the flow controller acts on. Its relative gain is $1 - \lambda_{11} = -0.000\,05$, and the design point is unstable. The recovery leaves it by a full point 4.0 min into the run and settles by 20 min at 22.4 %, with the pump on its 100 % stop and the valve at 83.3 %, while the flow is held on setpoint from the far side of the maximum. With the recovery loop also detuned to $K_{c,R} = 0.05$ and $\tau_{I,R} = 1.5$ min, the drift grows slowly enough (0.1 point at 33 min) for the step to decide the outcome, and the run ends with the valve on its 5 % stop and 82.2 m^3/h of permeate. The row is reported at the same scale because decoupling is a refinement applied to a correct pairing rather than a substitute for one.

The decoupler gain is not a free parameter. Table 2.4 sweeps it across the design value, and over-decoupling at 1.6 is nearly as costly as no decoupling at all. The curve has a minimum rather than a plateau, which is what makes the reference speed N_{ref} worth attention, because it is a static calibration on a plant whose gains drift.

🕮 **Pause.** Try Exercise 2.2 before reading on.

Table 2.5 Specific energy and saturation index against recovery with permeate production held at 450 m^3/h and the antiscalant at 3 mg/L.

R (%)	N (%)	P_f (bar)	SEC (kWh/m^3)	SI
40	88.33	60.41	2.5156	0.780
45	88.00	62.00	2.4627	0.850
49	88.19	63.55	**2.4520**	0.916
52	88.55	64.89	2.4590	0.973
54	88.90	65.88	2.4703	**1.015**
58	89.83	68.13	2.5082	1.111

TRY IT IN THE APP 6 · A One-Way Decoupler

Load the scenario, run it to 60 min and watch the recovery after the production step at 10 min. Then set the Pump-to-valve decoupling gain to 1.6, press Reset and run again. The peak recovery deviation over the 50 min after the step grows from 0.160 points at the shipped 0.95 to 0.864 points (Table 2.4).

2.6.4 Warm Feed and the Quality Override

The intake warms from 25.0 to 32.5 °C over scenario 7. The production loop holds its setpoint throughout, because a warm feed is pressurized more easily. The permeate quality is what changes. With the supervisory layer set to the tank level alone, so that the quality override is out of the loop, the permeate reaches 379 mg/L at the end of five hours against 250 mg/L at design, and the specific energy falls to 2.358 kWh/m^3 because the same production requires only 59.06 bar. A seawater plant therefore produces its cheapest water and approaches its salinity specification most closely in the same season.

With the override active at a 340 mg/L internal limit, well inside the 500 mg/L contractual specification, (2.16) pulls the recovery setpoint from 45.0 to 32.6 % and holds the product at 340 mg/L. It takes the selector at 26.1 min, the first scan at which the conductivity cell reads above the limit, and the permeate peaks at 356.2 mg/L at 37 min before the lower recovery returns it to the limit. The cost is quantified: feed pressure 55.69 bar but specific energy 2.607 kWh/m^3, which is 10.6 % above the unconstrained case. Lower recovery yields less permeate per unit of feed pressurized, which is what the override spends. The trade is favorable only because off-specification water has no value.

2.6.5 Specific Energy Against Recovery

Recovery is the efficiency variable. Recovering more of the feed leaves less seawater to be raised to pressure and discarded, and it also produces a saltier concentrate whose osmotic pressure the pump has to overcome. The two effects produce an interior optimum, which is well established for spiral-wound seawater trains [15] and which the application resolves numerically at fixed production.

The minimum in Table 2.5 sits at 49 % recovery and 2.4520 kWh/m^3, and it is shallow. The whole curve spans 2.6 % between 40 and 58 % recovery, which is why the operating point is a standing argument rather than a calculation. The saturation index crosses one at 53.3 %, so the energy optimum is feasible with 4.3 points of margin. That margin is not a property of the plant. It is a property of an antiscalant dose and of a feed chemistry that is not measured online. At 52 % recovery the index is 1.265 untreated, 1.072 at 1 mg/L and 0.973 at 3 mg/L, for $7.27/h of chemical.

2.6.6 Fouling Drift and Oxidative Damage

Scenario 8 starts the fouling clock at 5 min and runs six hours at constant production. No loop loses control of its variable at any point.

Table 2.6 records a control system holding every setpoint while the plant it holds them on becomes expensive. Specific energy rises from the design value of 2.463 to 2.759 kWh/m^3, 12.0 %, over the

Table 2.6 Fouling with permeate production held at 450 m^3/h, scenario 8 run on past its six hours to 600 min. The element rating is 70 bar and the pump gain $\partial Q_p/\partial N$, by central difference on the fouled steady state, is in m^3/h/%.

t (min)	f (%)	P_f (bar)	N (%)	SEC (kWh/m^3)	$\partial Q_p/\partial N$
10	0.3	62.07	88.05	2.466	17.12
120	7.7	63.94	89.21	2.556	16.61
240	15.2	66.09	90.54	2.656	16.06
360	22.0	68.38	91.93	2.759	15.53
480	28.3	70.83	93.39	2.866	15.01
600	34.1	73.45	94.93	2.978	14.50

six hours of the scenario and by 20.9 % at 600 min, on the largest operating cost a desalination plant carries. Feed pressure crosses the 70 bar element rating at 440.2 min and 26.3 % fouling. The production controller reports no error at any point, because production is not what is degrading. Every gain in the plant moves at the same time. The production loop's own plant gain falls from 17.12 to 14.50 m^3/h/%, a loss of 15.3 %, which detunes the loop in the safe direction here and in the unsafe direction immediately after a clean.

> **+ DID YOU KNOW? Every Setpoint Held, Every Gain on the Move**
>
> From the simulator: in scenario 8 the membrane fouls while every loop holds its setpoint. By 600 min the feed pressure has risen from 62.07 to 73.45 bar, past the 70 bar rating, and the pump's own gain on permeate flow has fallen from 17.12 to 14.50 m^3/h/% (Table 2.6).

Scenario 10 is a disturbance of a different kind. Twenty-five minutes of dechlorination failure at 0.6 mg/L raises the salt permeability by 30.0 % through (2.12), and it remains raised. The product moves from 250 to 324 mg/L and rejection from 99.286 to 99.073 %, permanently, on a train whose permeate flow holds 450.0 m^3/h throughout the event and afterward. Just under a third of the distance to the 500 mg/L specification is consumed in less than half an hour and no loop recovers it. The production measurement alone records nothing at all.

2.6.7 Graded Shift

The graded scenario is five hours on a train that is already fouling, against a feed whose salinity and scaling tendency rise, an intake that warms and a distribution demand that rises from 450 to 510 m^3/h and then falls to 420 m^3/h. The target is 2300 m^3 delivered, production within 8 m^3/h of its setpoint over the last third of the shift, product under 500 mg/L throughout, no scale, no time above 70 bar and under 2.60 kWh/m^3 averaged over the delivered water. It loads with weak tuning, the recovery loop open, both supervisors off and the antiscalant at 2 mg/L. The recommended configuration closes both loops with the tuning of scenario 4 and the decoupler, turns on level and quality supervision with a 380 mg/L limit, and sets 46 % recovery and 5 mg/L of antiscalant.

The as-loaded run fails on volume alone, at 2238 m^3, 62 m^3 short of the target, because with both supervisors off the production setpoint carries no information about the distribution demand. It spends only 69.9 % of the shift inside the flow band, although its worst deviation over the judged last third, 7.5 m^3/h, passes. The tank drains from 60 to 17.3 %, short of the 12 % low-level alarm. Closing the level cascade converts the tank inventory into a production setpoint, and the recommended configuration passes with a margin of 43 m^3. Among the structural elements the cascade alone accounts for the delivered volume and for the flow integral error falling from 1821 to 270 m^3/h min. It needs the tuned flow loop to do so. With the shipped weak tuning, the cascade alone delivers 2426 m^3 but lets the production wander 117 m^3/h from its setpoint over the last third.

Table 2.7 The graded five hour run. Each row after the second changes one element of the recommended configuration before the run. In band is the share of the whole shift with production within 8 m^3/h of its setpoint.

Configuration	Volume (m^3)	In band (%)	SEC (kWh/m^3)	Scale (%)	Max SI
As loaded	2238	69.9	2.480	0.000	1.000
Recommended	2343	98.1	2.537	0.000	0.958
dose 2 mg/L	2343	98.1	2.538	0.677	1.025
no decoupler	2343	98.0	2.537	0.000	0.958
no override	2343	98.1	2.537	0.000	0.958

The remaining rows are informative for what they do not show. Returning the antiscalant to 2 mg/L scales the membrane, with 0.677 % of permeability lost irreversibly and a peak saturation index of 1.025, and every other column unchanged. Removing the decoupler or demoting the override changes the table only in the in-band share, by 0.1 point without the decoupler. The decoupler still cuts the recovery integral error from 12.56 to 6.78 % min, by 46 %, but it leaves the peak deviation at 0.79 points (0.78 without it), because the level cascade moves the production setpoint far more slowly than the setpoint steps of scenarios 4 and 6. The override binds once, from 193.5 to 213.1 min, after the conductivity cell crosses the 380 mg/L limit, and lowers the recovery target to 45.35 %. The product then spends 14.2 min above the limit with a peak of 382.6 mg/L, against 19.1 min and 383.4 mg/L with the override demoted.

Two of the four structural elements carry this particular shift and two are insurance against disturbances larger than the shift contains. Structure that returns nothing on the scorecard of a favorable shift is not thereby unnecessary, and a scorecard that cannot separate the two is a poor guide to what to build [8].

2.6.8 Predictive Controller Against the Decoupled Loops

Table 2.8 compares the MPC with the decoupled loops on the cases in which a limit binds or production steps. Each case runs for its own duration with the noise off, as it ships. The loops are the decoupled pairing with the recovery loop closed and the case's supervisory layer, which is level and quality supervision in scenarios 7, 8, 11 and 12 and none in scenarios 4 and 9, whose events are operator setpoint steps that the level cascade would overwrite. The MPC carries the settings of scenario 11 into every case without retuning. Three comparators separate what the plan contributes. The pressure override is a PI controller on the measured feed pressure, with its setpoint 0.5 bar under the rating, that low-selects the recovery target with external reset ($K_c = 4\,\%/\text{bar}$, $\tau_I = 8$ min, chosen on scenario 11 after the fact), and it exists only in the measurement scripts because the application ships none. Hindsight is the best fixed tuning of the loops for that case on a grid searched after the fact, and the last comparator hands the loops the plan's own targets every 0.1 min. Flow error integrates the distance between the permeate flow and the flow that the tank or the operator requests.

Scenario 11 runs a membrane that fouls at its normal rate from $t = 0$ at 50 % recovery with level and quality supervision, and the distribution demand rises from 450 to 540 m^3/h between 90 and 200 min. Both structures deliver 2415 m^3. The loops hold the recovery at 50 % and let the feed pressure follow the fouling and the demand, so it spends 108.6 min over the 70 bar rating and reaches 74.00 bar. The MPC takes its recovery target below 49.9 % at 97.5 min and keeps the pressure at or below 69.81 bar, as shown in Figure 2.4. At 200 min the recovery is 33.6 % with the pump at full speed, while the target, blind to that speed limit, has fallen to 31.6 %, and by 300 min both are back at 49.21 %. Holding the rating costs 2.7569 against 2.6755 kWh/m^3 over the run, 3.0 % more, and 3.009 against 2.811 kWh/m^3 over the peak, because a lower recovery raises more seawater to pressure for the same permeate.

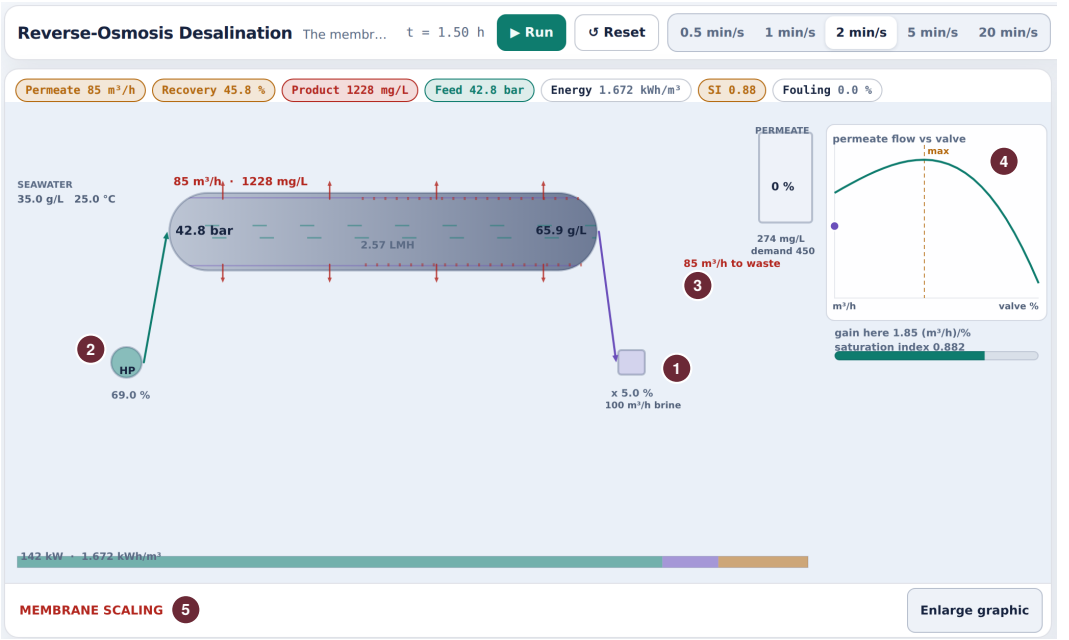

Figure 2.3 The reversed pairing of scenario 5 at 1.50 h, eighty minutes after a production setpoint step from 450 to 500 m^3/h was applied to a controller wired to the concentrate valve, cropped to the application header and the plant view. (1) The concentrate valve has integrated to its 5 % stop. (2) The recovery loop has dragged the pump down to 69.0 %. (3) The permeate, at 1228 mg/L, is diverted to waste, and the permeate tank has emptied. (4) The operating map plots permeate flow against valve position and marks the maximum on which the controller began. (5) The scaling alarm records the recovery excursion that followed the step.

The ablations place the gain in the target adjustment rather than in the optimization. Without it the same plan bends its pressure slack and spends 96.1 min over the rating, because one slack for the whole horizon costs less than a recovery target held over 90 steps. The model frozen at its first plan still holds 69.80 bar, so relinearizing matters little on this case. The classical override also holds the rating, at 69.87 bar, with a flow error of 500.8 at the shipped tuning and 69.1 at the hindsight tuning against the plan's 107.3 m^3/h min. What the optimization adds is coordination while the limit binds. The loops handed the plan's targets every 0.1 min swing the valve through 134.8 % of travel, take the recovery down to 25.07 % and still spend 11.7 min over the rating with a flow error of 1410.4 m^3/h min, while the plan, moving both actuators together, holds the rating with a flow error of 107.3. With instrument noise on the flow meter (2 m^3/h), the conductivity cell (5 mg/L) and the tank level (0.3 %) at four seeds, the MPC spends 2.0 to 4.0 min over the rating against 109.6 to 113.9 min for the loops, and its highest pressure, 70.30 bar, lies 0.5 bar above its band. Halving the integration step moves only the flow errors of the loops, the MPC and the MPC without its target step, by at most 0.3.

In the other cases the plan holds the rating wherever the loops break it. With the recovery raised in steps to 58 % in scenario 9 the loops spend 73.5 min over the rating while the MPC and the override hold it, and the lower recovery they choose also cuts the permeability lost to scale from 20.354 to 13.163 and 12.852 %, although the plan carries no saturation limit. Without its target step the plan still holds the rating in scenario 9, at 69.80 bar. Run on to 600 min, scenario 8 needs more than the pump can give. The MPC holds 69.88 bar with the recovery down to 26.78 % and the pump at full

Table 2.8 The MPC Against the Decoupled Loops, With the Classical Pressure Override and the Plan's Information Handed to the Loops

Case	Structure	$P_{f,\max}$ (bar)	Over rating min (bar min)	Over $C_{\lim}$ min (mg/L min)	Flow error	SEC
11	loops	74.00	108.6 (302.7)	0	301.1	2.6755
	MPC	69.81	0	0	107.3	2.7569
	MPC, no target step	70.75	96.1 (35.8)	0	106.8	2.7129
	MPC, model frozen	69.80	0	0	147.9	2.7535
	loops + override	69.87	0	0	500.8	2.7549
	hindsight + override	69.87	0	0	69.1	2.7506
	hindsight	74.00	108.3 (298.9)	0	39.3	2.6752
	loops, plan's targets	70.47	11.7 (3.4)	0	1410.4	2.7840
9	loops	73.11	73.5 (111.8)	0	135.3	2.5256
	MPC	69.72	0	0	8.9	2.5130
	loops + override	69.57	0	0	96.3	2.5099
8[a]	loops	73.45	159.6 (271.1)	0	138.5	2.7109
	MPC	69.88	0	0	96.0	2.8441
	hindsight + override	70.19	7.7 (0.65)	0	2248.6[b]	2.8295
7	loops	62.00	0	274.7 (742.6)	72.4	2.5629
	MPC	62.00	0	14.9 (42.7)	44.1	2.6506
	loops, plan's target	62.00	0	19.3 (122.7)	88.6	2.6468
4	loops	66.03	0	0	157.4	2.5307
	MPC	66.13	0	0	15.4	2.5308
12[c]	loops, met	66.58	0	14.2 (21.2)	270.5	2.5372
	MPC, met	66.59	0	0	27.9	2.5374

Flow error in m^3/h min, SEC in kWh/m^3 over the run. The salinity limit $C_{\lim}$ is 340 mg/L in scenario 7 and 380 mg/L in scenarios 11 and 12, and the permeate stays under 291 mg/L in the others. Every MPC run planned every 0.1 min (3000 plans on scenario 11) and rejected none, at 0.39 to 0.44 ms per plan in Node in the measured runs. [a]Run on to 600 min. [b]The level cascade's request winding up against a pump at full speed, not lost water (4491.9 against 4490.4 m^3 for the MPC). [c]The challenge with its suggested settings.

speed for the last 61.1 min, at 4.9 % more specific energy than the loops, which spend 159.6 min over the rating. Without its target step the plan spends 105.7 min over it there. In scenario 7 the plan treats the 340 mg/L quality limit as its own constraint and is over it for 14.9 min against 274.7 min under the shipped quality override, at 3.4 % more energy, and the loops handed the plan's recovery target come to 19.3 min.

Two results are losses. Flow tracking is mostly tuning. The hindsight loops of Table 2.8 reach a flow error of 39.3 on scenario 11 while spending 108.3 min over the rating, and a wider grid finds loops that reach 4.6 there, still 108.2 min over the rating, and 14.0 against the plan's 27.9 on the challenge, with 41.7 against 13.8 % of pump travel. On the production step of scenario 4 the plan lets the recovery swing 1.83 points, where the decoupler of Equation (2.15) holds it to 0.21. The plan weighs 1 per $(m^3/h)^2$ of flow error against 4 per $\%^2$ of recovery error, so the 70 m^3/h step outweighs the recovery in its first decisions, and with the recovery weight raised to 4 the swing falls to 0.72 points.

The challenge is met under both structures with its suggested settings, at 2343 m^3 with production within 0.1 m^3/h of its setpoint over the last third under the MPC and 0.7 m^3/h under the loops. The MPC also keeps the permeate under the unscored 380 mg/L quality limit, peaking at 378.2 mg/L, where the loops exceed it for 14.2 min with a peak of 382.6 mg/L. As loaded, with the weak tuning, the recovery loop open and both supervisors off, neither structure meets it, at 2238 and 2250 m^3, and the MPC loses 0.025 % of the permeability to scale.

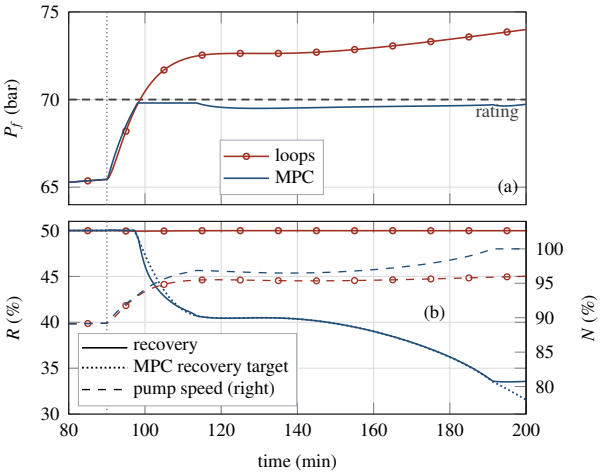

Figure 2.4 Scenario 11 from 80 to 200 min with the MPC and with the decoupled loops. Curves with open circles are the loops and curves without markers the MPC. (a) Feed pressure (solid) against the 70 bar element rating (dashed). (b) Recovery (solid), the recovery target the MPC holds (dotted) and the pump speed command on the right axis (dashed). The demand steps from 450 to 540 m^3/h at 90 min. At 200 min the loops are at 74.00 bar and the MPC at 69.73 bar with 33.6 % recovery and the pump at 100 %.

2.7 Technology in Practice

What Measures It

Recovery is not read by one instrument. It is the permeate flow divided by the feed flow [11], so the recovery loop rests on two flow measurements. A magnetic flowmeter suits such lines: a conductive liquid crossing a magnetic field induces a voltage across the pipe that is linear in the volumetric flow, provided the liquid conducts well enough, the pipe runs full and the flowtube is grounded [25]. Electrodes are set across the pipe horizontally so that a passing gas bubble cannot break their contact with the liquid, and an ungrounded flowtube reads stray currents as error [25]. The chapter gives its flow, recovery and level transmitters a 0.05 min filter and no other error, and the feed pressure transmitter is what the predictive controller's bias compares with its model.

> **TECHNOLOGY** **Sensor: The Permeate Conductivity Cell**
>
> Permeate quality is read as electrical conductivity, an approximate measure of the salt in the water, because dissolved salts carry current and pure water hardly does. The reading is taken in a fraction of a minute and reported at 25 °C, and it is turned into dissolved solids by a relation derived for each water, since conductivity changes with temperature and with the mix of ions [11]. A cell of two electrodes fouls as ions bond to the metal and then reads the water as less conductive than it is, so the permeate looks fresher on the trend than it is. A four-electrode cell is immune to that error, and an electrodeless toroidal cell resists fouling [25]. The chapter's cell adds a 0.45 min transport delay and a 0.35 min response.

What Moves It

The production loop moves the pump through its drive, and the recovery loop moves the concentrate valve. An equal-percentage trim is decidedly nonlinear in its flow capacity, and its formula can never reach zero, so it cannot be trusted at small stem positions [25]. The chapter's valve has a rangeability of 30 and a 0.15 min travel time, and its lower clamp of 5 % is the stop that the reversed pairing of scenario 5 reaches.

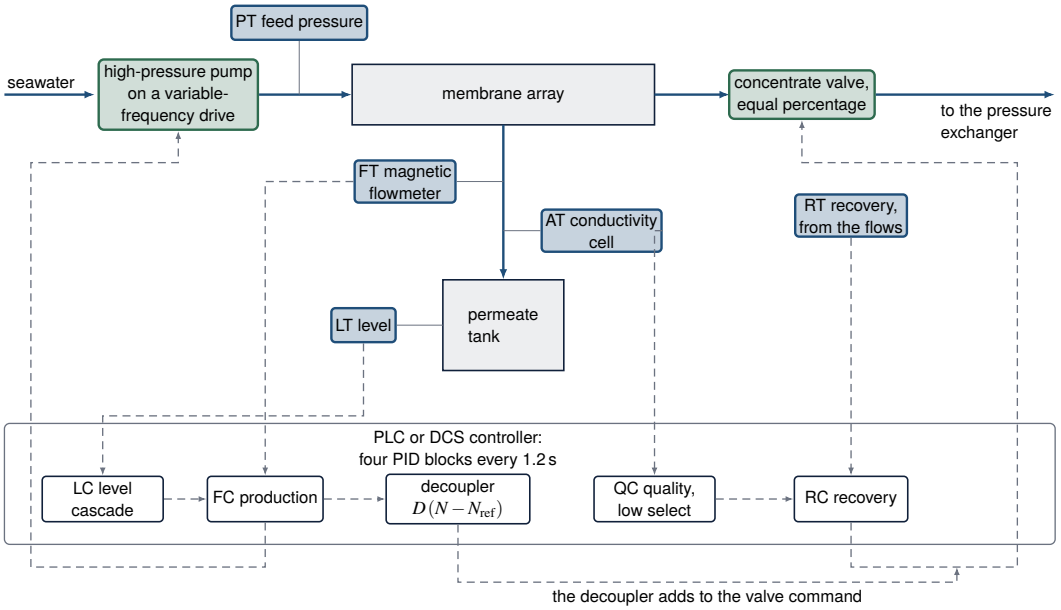

Figure 2.5 The two regulatory loops on plant hardware. The magnetic flowmeter on the permeate line feeds the production controller, which drives the pump through its variable-frequency drive. The recovery, taken from the flows, feeds the recovery controller, which moves the equal-percentage concentrate valve. The level cascade writes the production setpoint, and the conductivity cell's quality controller can only lower the recovery setpoint. Notice the decoupler path: the production controller's own output reaches the valve before the recovery controller sees an error.

> **TECHNOLOGY** **Final Element: The Pump Drive**
>
> The high-pressure pump is a centrifugal machine, and its discharge pressure falls as its flow rises. A variable-frequency drive sets its speed: a rectifier turns line power into direct current, a filter smooths it, and an inverter rebuilds alternating current at the frequency asked for. An induction motor turns slightly slower than the field that frequency sets. Running only as fast as needed saves power, cuts vibration and lets the drive ramp the speed to spare the machine [25]. Such drives were already commonplace on brackish-water trains when a federal handbook for planners was written [11]. In the chapter the drive is a 0.10 min lag on the speed command, and the pump curve droops with flow as Equation (2.5) gives.

Where the Controller Runs

Distributed control systems serve water and wastewater treatment, with programmable logic controllers as local controllers under a supervisory scheme, and supervisory control and data acquisition systems monitor and control water distribution. Loops at both levels run continuously with cycle times from milliseconds to minutes, and some are nested so that one loop's setpoint comes from another loop's process variable [26]. In a research roadmap for water treatment these systems are described as the result of piecemeal development over decades [27]. The chapter's structure maps onto standard function blocks: four PID blocks scanned every 1.2 s, a low selector with external reset on the recovery setpoint, a cascade from the tank level onto the flow setpoint, and one multiplication and one addition for the decoupler. The optional predictive controller plans every 0.1 min, a supervisory period. Figure 2.5 draws the two regulatory loops on their devices.

IN PRACTICE **Recovery Comes First**

A federal handbook for planners calls recovery the most important process variable of a reverse osmosis train, the one the others largely depend on, because it sets the concentration factor $1/(1-R)$. Most trains dose a scale inhibitor, sometimes with acid, because at least one salt is likely to be supersaturated in the concentrate. The handbook lists the ASTM practices for membrane systems, among them D4516 for standardizing performance data, D4472 for record keeping and D3739 for the Langelier saturation index [11]. A later research roadmap asks for inline sensors that would trigger cleaning instead of delaying it until a reduction of water permeability is observed [7].

What Is Modernizing It

Varying the speed of the machine that drives a liquid is more energy-efficient than running it at full speed and throttling its discharge, and it removes a control valve from the line [25]. Energy recovery had become fairly commonplace on seawater trains when a federal handbook put the recovered energy at 25 to 35 % of the input energy and gave the pressure exchanger, then used on smaller systems, an efficiency above 90 % [11]. The chapter's exchanger is a constant efficiency of 0.96. A national research roadmap asks for model-based control and data-driven models that optimize set-points for energy, quality and recovery, and for control systems secured against manipulation by unauthorized users [27]. The chapter's predictive controller, which lowers the recovery target to keep the membrane under its rating, is one small instance of that program.

2.8 Limitations

The lumped array is the largest simplification. A real train is six or seven elements in series in each vessel, the lead element produces two to three times the flux of the tail, and it fouls and fails first. A single well-mixed channel with a log-mean bulk salinity reproduces the concentration factor and the average driving force and cannot show that distribution, so nothing here speaks to element-by-element flux balancing or to the staging decisions that follow from it.

Concentration polarization uses a fixed mass-transfer coefficient. In practice that coefficient depends on the crossflow velocity, which falls along the array as permeate is removed, so polarization is worst at the tail where the salinity is already highest. The application therefore understates the gradient in the place where scale actually starts.

Fouling is one scalar with one rate law. Real fouling is biological, colloidal, organic and inorganic at once, with different kinetics, different symptoms and different cleaning chemistries, and identifying which mechanism is active is a substantial part of the operating skill. The fouling clock here is also compressed by roughly one hundred and seventy times so that six weeks fits in six hours, and the oxidation rate is compressed similarly, so the exposure in scenario 10 is far smaller than the exposure that would do that much damage to a real polyamide film.

OPEN CHALLENGE **No Sensor Yet Reports Fouling Early**

A Department of Energy research roadmap calls membrane fouling the weak point of membrane technology. It lowers productivity, degrades water quality and raises the energy and chemical cost of cleaning, and frequent cleaning shortens membrane life. Practical sensors for fouling have not been adopted commercially by the desalination industry. The roadmap asks for inline sensors that measure several scale-forming solutes at once, and others that report the organic carbon and biomass that feed biofilms, so that pretreatment can change before severe fouling [7]. In this chapter the scaling margin rests on a dose and a feed chemistry that no instrument measures online. Progress would be a measured scaling margin that the recovery loop could defend.

Pre-treatment, intake and post-treatment are boundary conditions. There is no cartridge filter differential pressure, no coagulant dose, no silt density index and no remineralization, so the application cannot show the most common industrial cause of a fouling event, which is a failure in the

pre-treatment train. Salinity is a single equivalent salt, which removes boron, the ion seawater plants most often struggle to reject and the one that decides whether a second pass is needed. The pressure exchanger is a constant efficiency rather than a device with its own mixing, leakage and speed. The high-pressure trip is an alarm and the train keeps running, where a real interlock stops the pump and the lost production usually exceeds the excursion that caused it.

OPEN CHALLENGE **Control That Sees the Whole Train**

A national water research roadmap notes that most treatment systems still need operators to adapt to changing water quality and to worn components, and that the data of their many sensors serve mainly to detect failures. It asks for control networks coupled with models of the interdependencies between unit processes, and for supervisory systems that reach the data of pretreatment and of grab samples, which existing ones often cannot. It adds that the savings of predictive control are hard to quantify without grounding in actual facilities [27]. This chapter's train has no pretreatment, so its supervisory layer sees only the membrane. Progress would be a plantwide structure that carries intake and pretreatment measurements forward to the recovery and production targets.

The predictive controller's model is the plant's own structure at the design feed. Fouling, scale, temperature, feed salinity and pump wear reach it only through two permeability multipliers and an output bias, the feed pressure is taken as settled, and the membrane is held as it is now over the whole 9 min horizon. Each soft limit has one slack for the whole horizon, so the target adjustment rather than the slack keeps the plan off the rating, and that adjustment ignores the pump's top speed, which leaves the recovery at 33.6 % against a target of 31.6 % at the end of the peak in scenario 11. The 0.2 bar band is sized for quiet instruments, and the plan carries neither the saturation index nor a price on energy. Switching from the loops to the MPC at 150 min of scenario 11, with the pressure 2.77 bar over the rating, opens the valve by 21.95 points in one decision and costs up to 63.6 m^3/h of production for 4.56 min. The same switch at 110 min, with the feed pressure at 72.28 bar, opens the valve from 48.12 to 67.53 % in the first plan and brings the pressure under the rating within 0.40 min, at the cost of up to 57.4 m^3/h of production for 3.19 min. The plan serves the limit before either target.

Pause. Try Exercise 2.3 before reading on.

2.9 Discussion

Three results generalize past this plant.

The first is that a relative gain of one is a statement about a product of two numbers. It cannot distinguish a diagonal plant from a triangular one, and triangular is the common case whenever one actuator is much more influential than the other. Chapter 10 has another triangular plant, in which added oxygen raises the total pressure and the nitrogen does not move the oxygen. A large pump against a trim valve, a main heater against a side stream and a feed rate against a reflux ratio all produce nearly triangular gain matrices and a pairing tool that reports no interaction. The remedy is not to abandon the array but to read the gain matrix that produced it, which the single summary number makes easy to skip [2, 18].

The second is that one-way coupling admits a one-way remedy at low cost. Equation (2.15) is one multiplication and one addition per scan, it requires no model of the dynamics, and Table 2.3 shows it removing most of an interaction that survives good tuning. It does require a calibration, and the calibration ages. That is a smaller liability than a full decoupling matrix, which is why triangular plants are worth identifying as such rather than escalating to a multivariable controller [12].

The third result concerns constraints and optima. Table 2.5 has a genuine interior energy minimum and a scaling limit 4.3 points above it, and the distance between them is set by a chemical dose rather than by anything the control system measures. Moving recovery toward the constraint buys a fraction of a percent of energy and risks a permeability loss that no rinse recovers, a trade that is

clear when it is written down and less clear during operation, when the specific energy is displayed continuously and the saturation index is not. Volume 9, Chapter 8 has the same arrangement, a lowest-cost operating point that sits on a constraint that moves with the weather. Desalination supplies water to a large fraction of the world's coastal population and consumes electricity in proportion [1], so the distance between an optimum and a limit is worth stating precisely [5].

2.10 Conclusion

This chapter studied a seawater reverse osmosis train, a two-by-two plant whose pairing tool reports no interaction and whose loops interact anyway. The concentrate valve has a gain of –0.001 m^3/h/% on permeate flow at the design point and gains of opposite sign 8 % to either side, so the reversed pairing drives the valve to its stop and leaves the plant at 82.5 m^3/h against a 500 m^3/h setpoint. The correct pairing has a relative gain of 1.000 05 and still deviates 0.937 points in recovery on a production step, because the gain matrix is triangular rather than diagonal and the array is structurally unable to report that. One static feedforward from the pump command to the valve, calibrated from a gain ratio of 0.914, reduces the peak to 0.160 points. Behind the pairing question sit a shallow energy optimum at 49 % recovery, a crystallization limit 4.3 points above it, a fouling drift that adds 12.0 % to the specific energy in six hours without moving a single controlled variable, and an oxidation event that permanently consumes just under a third of the product quality margin while the permeate flow reads exactly on setpoint. The optional model predictive controller holds the feed pressure of scenario 11 at 69.81 bar where the decoupled loops spend 108.6 min above the 70 bar rating, at 3.0 % more specific energy, and a classical pressure override on the recovery target also holds it, at 69.87 bar, so the rating is held by giving up recovery at the limit rather than by the optimization itself. The app lets a reader measure loop pairing, decoupling and constraint-aware operation on one plant.

KEY POINTS

- Which handle holds which target is decided before any tuning: production belongs on the pump and recovery on the concentrate valve.
- The valve's gain on permeate flow passes through zero at the design point, so the reversed pairing drives the valve to its stop and leaves the plant at 82.5 m^3/h against a 500 m^3/h setpoint.
- The correct pairing has a relative gain of 1.000 05 and still deviates 0.937 points in recovery, because the gain matrix is triangular rather than diagonal.
- One static feedforward from the pump command to the valve reduces the peak to 0.160 points.

Where else this applies. Chapter 10 has another triangular plant, and Chapter 3 a plant coupled both ways, with a decoupler of both terms.

References

[1] M. Elimelech and W. A. Phillip, "The future of seawater desalination: Energy, technology, and the environment," *Science*, vol. 333, no. 6043, pp. 712–717, 2011.

[2] 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.

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

[4] J. G. Wijmans and R. W. Baker, "The solution-diffusion model: a review," *Journal of Membrane Science*, vol. 107, no. 1–2, pp. 1–21, 1995.

[5] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[6] N. Voutchkov, "Energy use for membrane seawater desalination: current status and trends," *Desalination*, vol. 431, pp. 2–14, 2018.

[7] D. Giammar, S. Jiang, P. Xu, R. Breckenridge, T. Edirisooriya, W. Jiang, L. Lin, J. Macknick, N. Rao, D. Sedlak, J. Stokes-Draut, and X. Xu, "National alliance for water innovation (NAWI) municipal sector technology roadmap 2021," U.S. Department of Energy, Report DOE/GO-102021-5565, 2021.

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

[9] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[10] G. Migliorini and E. Luzzo, "Seawater reverse osmosis plant using the pressure exchanger for energy recovery: a calculation model," *Desalination*, vol. 165, pp. 289–298, 2004.

[11] I. C. Watson, O. J. Morin, Jr., and L. Henthorne, "Desalting handbook for planners," U.S. Department of the Interior, Bureau of Reclamation, Denver, CO, Desalination and Water Purification Research and Development Program Report 72, 2003. [Online]. Available: https://www.usbr.gov/research/dwpr/reportpdfs/report072.pdf

[12] I. M. Alatiqi, A. H. Ghabris, and S. Ebrahim, "System identification and control of reverse osmosis desalination," *Desalination*, vol. 75, pp. 119–140, 1989.

[13] S. Sobana and R. C. Panda, "Review on modelling and control of desalination system using reverse osmosis," *Reviews in Environmental Science and Bio/Technology*, vol. 10, no. 2, pp. 139–150, 2011.

[14] A. R. Bartman, P. D. Christofides, and Y. Cohen, "Nonlinear model-based control of an experimental reverse-osmosis water desalination system," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6126–6136, 2009.

[15] A. Zhu, P. D. Christofides, and Y. Cohen, "Effect of thermodynamic restriction on energy cost optimization of RO membrane water desalination," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6010–6021, 2008.

[16] A. Ruiz-García and I. Nuez, "Long-term performance decline in a brackish water reverse osmosis desalination plant. predictive model for the water permeability coefficient," *Desalination*, vol. 397, pp. 101–107, 2016.

[17] A. Antony, J. H. Low, S. Gray, A. E. Childress, P. Le-Clech, and G. Leslie, "Scale formation and control in high pressure membrane water treatment systems: A review," *Journal of Membrane Science*, vol. 383, no. 1–2, pp. 1–16, 2011.

[18] 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.

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

[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] 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.

[22] A. Abbas, "Model predictive control of a reverse osmosis desalination unit," *Desalination*, vol. 194, no. 1–3, pp. 268–280, 2006.

[23] A. R. Bartman, A. Zhu, P. D. Christofides, and Y. Cohen, "Minimizing energy consumption in reverse osmosis membrane desalination using optimization-based control," *Journal of Process Control*, vol. 20, no. 10, pp. 1261–1269, 2010.

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

[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] 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.

[27] D. Sedlak, M. Mauter, J. Macknick, J. Stokes-Draut, P. Fiske, D. Agarwal, T. Borch, R. Breckenridge, T. Cath, S. Chellam, A. Childress, D. Dionysiou, D. Giammar, E. Hoek, S. Jiang, L. Katz, J. Kim, R. Kostecki, J. McCutcheon, Y. Polsky, Z. Stoll, and P. Xu, "National alliance for water innovation (NAWI) master technology roadmap," U.S. Department of Energy, Report DOE/GO-102021-5617, 2021.

2.11 Exercises

2.1 Concept check · 1 min

Which balance does Equation (2.8) express?

(a) The water and salt fluxes that cross the membrane wall
(b) The pressure balance on the high-pressure feed node
(c) The growth of fouling with the permeate flux
(d) The salt inventory of the feed and concentrate channel

2.2 Concept check · 2 min

In Table 2.4 the peak recovery deviation on the scenario 6 production step falls from 0.937 points with no decoupler to 0.160 at $D = 0.95$, and rises again to 0.864 at $D = 1.6$. Which explanation is right?

(a) At 1.6 the valve moves 1.75 times the cancelling move, so the error returns with the opposite sign
(b) At 1.6 the decoupler drives the valve onto its stop, so the recovery loop loses its actuator
(c) At 1.6 the decoupler closes a second feedback path through the recovery, which then oscillates
(d) At 1.6 the conductivity cell delay lets the recovery drift before the decoupler acts

2.3 Concept check · 2 min

A plant engineer wants to use the app before a site study on a seawater train. Which question can the app answer within the limitations stated in Section 2.8?

(a) How much boron the permeate carries, and whether a second pass is needed
(b) How far the recovery moves when production is stepped on the pump
(c) How the cartridge filter pressure drop rises before a fouling event
(d) How the flux of the lead element compares with that of the tail element

2.4 Theory and calculation · 7 min

Use the gain matrix of Equation (2.13), with $K_{11} = 17.147$, $K_{12} = -0.001$, $K_{21} = 0.7029$ and $K_{22} = -0.7687$.

(a) Compute the relative gain λ_{11} of the direct pairing, production on the pump and recovery on the valve. *Checkpoint:* 1.00005.

(b) Compute the Niederlinski index of the swapped pairing, production on the valve and recovery on the pump, and say what it predicts for loops with integral action that are each stable alone. *Checkpoint:* About -1.9×10^4. Unstable for any tuning.

(c) Find the steady gain from the valve to the permeate flow when the recovery loop on the pump is closed perfectly. *Checkpoint:* $+18.75$ m^3/h/%.

2.5 Theory and calculation · 6 min

Use the design-point gains of Equation (2.13).

(a) Find the decoupler gain D that cancels the steady effect of the pump on the recovery, and the steady gain from the pump to the recovery that the shipped value $D = 0.95$ leaves. *Checkpoint:* 0.914 and -0.027 %/%.

(b) In scenario 3 the recovery loop is open with the valve at 50 %, and production steps from 450 to 510 m^3/h. Estimate the pump move and the change in recovery. *Checkpoint:* 3.50 % of speed and 2.46 points.

(c) In scenario 4 both loops are closed and production steps from 450 to 520 m^3/h. Estimate the steady valve move the recovery loop must make, and compare it with the valve travel reported in Section 2.6. *Checkpoint:* 3.73 %, against 50.00 to 53.44 %.

2.6 Theory and calculation · 8 min

At the design point the feed is 35.0 g/L at 25 °C, the recovery is 45.00 %, the flux is 13.636 LMH and the permeate is 250 mg/L. Use Table 2.1 and Equations (2.1) to (2.4).

(a) Find the concentrate salinity from the steady salt balance of Equation (2.8), and the log-mean bulk salinity. *Checkpoint:* 63.43 g/L and 47.82 g/L.

(b) Find the polarization factor, the wall salinity, the wall osmotic pressure and the net driving pressure. *Checkpoint:* 1.100, 52.60 g/L, 40.55 bar and 19.70 bar.

(c) Confirm the permeate salinity and the rejection. Then raise B by 30 %, as the oxidant of scenario 10 does, at the same flux and wall salinity, and find the new permeate salinity. *Checkpoint:* 250 mg/L, 99.286 %, then about 325 mg/L.

2.7 App challenge · 8 min

Use the default seed and the default speed. Read the recovery from its trend or from the `rec_pct` column of the exported file.

(a) Load scenario 4, Both loops, and a relative gain of one, and run to 60 min. Production steps from 450 to 520 m^3/h at 45 min. Read the highest recovery after the step and the concentrate valve position (`valve_pct`) at 60 min. *Checkpoint:* 45.94 % and 53.43 %.

(b) Load scenario 6, A one-way decoupler, and run to 25 min. The same step comes at 10 min. Read the lowest and the highest recovery after the step. *Checkpoint:* 44.84 and 45.00 %.

(c) Set the Pump-to-valve decoupling gain to 1.6, press Reset and run to 25 min again. Read the lowest recovery after the step. *Checkpoint:* 44.14 %.

Accept recoveries within 0.01 point and a valve position within 0.05 %. Compare each result with Tables 2.3 and 2.4.

2.8 App challenge · 6 min

Load scenario 5, The pairing you would have guessed, and run to 120 min at the default seed and a speed of 20. Both loops go to Auto at 5 min and the production setpoint steps from 450 to 500 m^3/h at 10 min.

(a) Read the highest recovery and the time at which it occurs. *Checkpoint:* 76.3 % at about 13.5 min.

(b) Read the pump speed, the permeate flow and the permeate salinity at 120 min from the exported file (`pump_pct`, `flow_m3h`, `tds_mgL`). *Checkpoint:* 68.39 %, 82.5 m^3/h and 1241 mg/L.

Accept the recovery within 0.1 point, the pump within 0.05 %, the flow within 0.5 m^3/h and the salinity within 2 mg/L. Explain the end state with the swapped pairing's relative gain.

2.9 App challenge · 8 min

Load scenario 11, Constraints, not a pairing: MPC against the decoupled loops. Use the default seed and a speed of 20, and run each case for 300 min. From each exported file read the highest feed pressure (`feed_bar`), the minutes above the rating (`metric_overpressure_min`) and the delivered water (`metric_water_m3`) in the last row, and compute the specific energy over the run as `metric_energy_kwh` divided by `metric_water_m3`.

(a) Run the scenario as shipped, with Multivariable control on MPC. *Checkpoint:* 69.81 bar, 0 min, 2415 m^3, 2.757 kWh/m^3.

(b) Set Multivariable control to Loops, press Reset and run again. *Checkpoint:* 74.00 bar, 108.6 min, 2415 m^3, 2.675 kWh/m^3.

Accept pressures within 0.02 bar, minutes within 0.2 and specific energies within 0.002 kWh/m^3. Say what holds the rating, and what it costs.

2.10 Application area · 4 min

At the design point the train makes 450 m^3/h of permeate from 1000 m^3/h of feed at 62.00 bar, with an array drop of 2.50 bar. Take the other constants from Table 2.1.

(a) Find the three terms of Equation (2.9) and the specific energy. *Checkpoint:* 907.0, 114.4 and 86.8 kW, and 2.463 kWh/m^3.

(b) Find the specific energy without the pressure exchanger ($\eta_{PX} = 0$), and the electricity the exchanger saves in a year of continuous operation at this duty, and its value at the app's default electricity price of $0.085 per kWh. *Checkpoint:* 4.907 kWh/m^3, about 9.6 GW h and $0.82 million.

2.11 Application area · 4 min

In the graded challenge the feed salinity steps to 36.5 g/L at 130 min and the scaling tendency h rises from 1.00 to 1.10 at 230 min. The saturation index is C_b/C_{sat}, the ratio in Equation (2.11), with C_{sat} as given in Section 2.3.

(a) Find C_{sat} after the hardness step at an antiscalant dose of 5 mg/L and at 2 mg/L. *Checkpoint:* 70.24 g/L and 65.60 g/L.

(b) Estimate the concentrate salinity at 46 % recovery, neglecting the salt in the permeate, and the saturation index at each dose. Compare with Table 2.7. *Checkpoint:* 67.6 g/L, 0.962 and 1.030.

(c) Find the highest recovery that keeps the index below one at 5 mg/L after the hardness step. *Checkpoint:* About 48.0 %.

2.12 Review the agent · 5 min

An AI assistant is asked to simplify the controls of the train and replies as follows. "The relative gain of the direct pairing is 1.00005, so the production and recovery loops do not interact. The static decoupler adds nothing and can be removed. To save valve wear, slow the recovery loop to an integral time of 8 min. A production step will then move the pump and leave the recovery alone." Find the flaw in the proposal and say how it could be checked.

CHAPTER 3

Paper Machine Dryer Simulator for Control Through a Scanning Quality Gauge

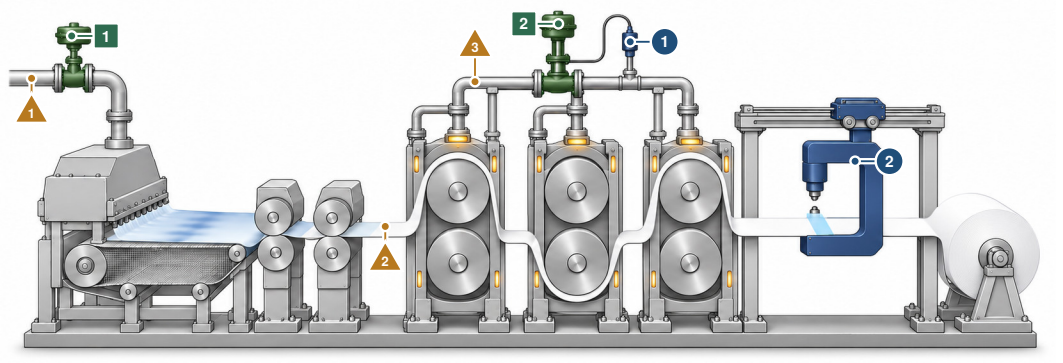

The controller moves 1 thick stock valve, sets the fiber reaching the sheet 2 steam valve, sets the dryer pressure

The controller reads 1 steam pressure, read all the time

2 scanning gauge, posts basis weight and moisture once per pass

What disturbs it 1 stock consistency 2 press dryness (a plugged felt) 3 mill steam header pressure

Notice that the only report on the sheet comes from the scanning frame at the dry end, long after both valves have acted. Illustration, not to scale.

IN PLAIN TERMS **Two Knobs and a Slow Report Card**

The main idea. A paper machine turns watery pulp into a moving sheet and dries it on hot rollers. Two knobs set how heavy and how wet the sheet is. One is the pulp valve and one is the steam.

The challenge. The only report on the sheet comes from a gauge that crosses it back and forth. It gives one average per crossing, so every answer is already old. Each knob also moves both results, because the gauge weighs the water along with the fiber.

Why it matters. Which knob holds which target, how the cross effects are canceled and how often the gauge reports decide how much sheet is made out of specification.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 3.2, the fact on drying energy, the history of the first experiment (Section 3.4), the two try-its and the fact on the scan period (Section 3.6) and the key points
Study	65 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 3.6), Technology in Practice (Section 3.7), Limitations (Section 3.8) and the two open challenges

The control question A paper machine sells how heavy and how wet its sheet is, and measures both with a gauge that averages a whole traverse and holds the answer until the next pass. How much dead time does that instrument add, and how should two interacting quality loops be paired, decoupled and timed around it?

Learning objectives

- Compute the effective dead time of a scanning gauge from the transport delay and the scan period.
- Explain why the steam pressure moves the conditioned basis weight, and compute the relative gains and decoupler gains of the two-by-two plant.
- Show on the app that a shorter scan period settles a ringing loop with no gain changed, and say what it costs the cross direction.
- Judge which half of a static decoupler carries its benefit, and what it costs the other loop.
- Explain why a predictive controller with the scanner in its model loses to retuned loops when its model holds the machine speed.

Prerequisites Sections 1.4.5 (relative gain and pairing), 1.4.6 (simplified decoupling), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/paper-machine/

See also Volume 5, Chapter 3 for composition control through a slow analyzer, Volume 5, Chapter 4 for a quality known hours late from sampled instruments, and Volume 5, Chapter 10 for a product moisture inferred between laboratory samples.

3.1 Introduction

The quality measurements a process plant actually controls on are seldom continuous. A composition arrives once per analyzer cycle, an assay once per shift, a coating weight once per traverse of a gauge, and in each case the instrument reports a property of product already made [1]. A paper machine is the extreme case. One head carrying a beta transmission gauge and an infrared moisture sensor traverses the full width of the sheet, and the two numbers it hands the control system are the only description of the product that exists. In the simulator of this chapter the head takes 20 s to cross an 8 m trim while 333 m of sheet passes it at 1000 m/min, so the reported basis weight is the average of a diagonal stripe rather than the value of anything at any instant.

What a scanning gauge costs a loop is arithmetic that a block diagram hides. Three lags live in the one instrument. The sheet takes L/v seconds to reach the frame, the boxcar average over one traverse is worth a further half period, and holding the answer until the next pass is worth half a period again. The effective dead time is therefore the transport delay plus one whole scan period, 21.9 s for moisture and 30.7 s for basis weight against transport delays of 1.92 s and 10.68 s. What the simulator sets right is the habit of treating a sampling interval as an implementation detail settled after the tuning. It is part of the plant, it is the larger part of it here, and a controller given the transport delay alone rings at a period of 100 s with a swing that grows. Tuning rules for dead-time-dominant plants already carry half a sample period added to the delay [2], and the application makes the whole term visible by drawing the posted staircase against the sheet it is supposed to describe.

The second structural fact is that the two handles are not two. Basis weight is sold conditioned, water included, so a beta gauge weighs the fiber and the water in it together. Evaporating water therefore removes saleable mass, the steam pressure becomes a basis weight actuator, and the plant is a two-by-two with four non-zero connections. The relative gain of the natural pairing is 1.52 and that of the other pairing is -0.52, which is Bristol's formal statement that one pairing is workable and severe while the other cannot be tuned [3]. Both facts generalize well past paper. Sheet and film processes of every kind, plastic film, coating, board and tissue, are measured the same way and pose the same machine-direction and cross-direction problem [4], and any plant whose controlled property is measured on a cycle inherits the dead-time accounting [1]. Volume 5, Chapter 3 meets the same accounting with an online analyzer, and Chapter 5 with a coat weight gauge whose delay is fixed in meters of web.

Simulators of paper machines are usually built to answer thermal questions. Multi-cylinder drying models resolve the heat and mass transfer through the felt, the sheet and the cylinder shell to predict a dryness profile down the machine [5, 6], and control studies on those models concentrate on the steam and condensate system [7]. However, a student handed a first-order plus dead-time transfer function and asked for proportional-integral settings never encounters the instrument that produced the dead time, and on a paper machine the instrument is most of it. This chapter works with a fine-paper machine whose quality controllers execute only when a traversing gauge posts a reading, so that the sample interval, the boxcar average and the zero-order hold appear as measured seconds of dead time and as an oscillation that no gain change removes. The drying model, the relative gain array, the decoupler and the predictive controller are all established methods. Section 3.2 traces how the scanning gauge shaped quality control on paper machines. Section 3.3 gives the fourteen-state machine with its instrument chain, every balance, constant and initial value. Section 3.4 derives the effective dead time of the gauge, measures the two-by-two gain matrix and designs a static decoupler from it. Section 3.5 describes the event-driven control layer, its loss-of-measurement logic, the production cascade and an optional predictive controller whose model contains the scanner. Section 3.6 orders the remedies on one measured plant, in scenarios that run from two handles moved by hand to a grade change with the steam header already sagging. Shortening the scan period buys dead time directly and outperforms any gain change. One half of the static decoupler carries almost all of its benefit, and the measurement says which half. With the production loop in service, decoupling trades a faster moisture loop for a slower weight loop. The same cases are measured under the predictive controller, against the shipped loops and against loops retuned by the controller's own tuning rule. Section 3.7 describes the real gauges, valves and controllers. Sections 3.8 and 3.9 state what the model cannot show and price moisture and press dewatering, and Section 3.10 sums up.

3.2 Quality Control on Paper Machines

Paper was made and sold long before either of its graded properties could be measured on the machine. Basis weight is defined by cutting a sample of known area from the reel, conditioning it and weighing it, and moisture by weighing the same sample before and after an oven. Both are laboratory tests on a finished reel, so the control loop of the 1940s closed through an operator who was told what the last reel had been and adjusted the thick stock valve or the steam pressure for the next one. The correction arrived hours after the sheet it described, and the machine ran with a generous margin above the customer's minimum weight because paper below specification cannot be shipped.

The instrument that changed the problem was the traversing beta gauge. A beta source on one side of the web and an ionization detector on the other give a signal proportional to the mass in the beam, and a frame that carries the pair across the sheet turns a point measurement into a description of the whole web. The reading arrives once per traverse rather than once per reel, which is three orders of magnitude faster than a laboratory test and still slower than every other measurement on the machine. Computer control followed, and the history segment below tells how its first application to a paper machine treated the machine-direction problem as regulation against a noisy sampled measurement rather than setpoint tracking.

HISTORY 1965: A Computer Takes the Basis Weight Loop

A feasibility study for a process computer at a kraft paper mill in Sweden ran from May to September 1963. It found basis weight varying with a standard deviation of 1.3 g/m^2 and set 0.7 g/m^2 as the target under computer control. The computer was installed late in December 1964, and the first on-line basis weight control ran on April 28, 1965, for ten hours. Its control law came from identification experiments on the running machine and was designed to make the variance of the output as small as possible. From the beginning of 1966 the loop ran continuously and held standard deviations of 0.5 g/m^2 in wet basis weight and 0.3 g/m^2 in dry [8].

The two decades that followed separated the problem in two. The machine-direction loops hold the average of each traverse with the thick stock valve and the dryer steam pressure, and the cross-direction loops hold the profile across the web with a row of slice or induction actuators driven from the same head. Adaptive and self-tuning regulators reached the machine-direction loops early, because the plant gain moves with the grade and the speed [9]. The cross-direction problem became a subject of its own, with a hundred or more actuators, an ill-conditioned interaction matrix and a spatial bandwidth limit that no controller can exceed [4]. The two problems share one instrument and pull it in opposite directions, because a faster traverse shortens the machine-direction dead time and reduces the number of measurement points across the sheet.

Assessment of what the resulting loops achieve came later and was uncomfortable. The minimum-variance benchmark, which compares a loop's output variance against the variance no feedback controller could remove given the delay, was introduced for this class of plant [10], and mill-wide audits built on it made control loop performance a routine engineering deliverable [11]. The recurring finding is structural rather than numerical. Loops are tuned against the plant and not against the instrument, pairings chosen at commissioning are never revisited, and interaction between the two quality loops is treated as noise. The scenarios of the app are built around those three failures.

The economic weight behind them is drying energy. The dryer section is the largest single steam consumer in an integrated mill, and reducing its load, by pressing more water out mechanically or by giving away less moisture, is consistently identified as one of the cheapest energy savings available in the industry [12]. Section 3.9 prices both on the simulated machine.

DID YOU KNOW? Drying in the Pulp and Paper Sector

Drying takes 22 % of the energy consumed by the pulp and paper sector of the United States, and running its thermal dryers cost an estimated $919 million in 2020. Steam-heated multicylinder dryers make up 85 to 90 % of the dryers in the sector, in a study the same assessment cites [13].

3.3 Process Model

The simulated machine runs 1000 m/min on an 8.0 m trim. Thick stock is metered by valve NCV-101 into an approach system, drained on the wire, pressed to 43 % dryness and carried through three banks of steam-heated cylinders with sheet path lengths of 28, 58 and 26 m. Scanning frame QT-401 stands 32 m past the last cylinder and 178 m past the slice. At the design line-up the sheet leaves at 80.000 g/m^2 and 5.500 % moisture, which is 38.40 t/h of saleable paper, 36.29 t/h of it fiber and 2.11 t/h of it water that the customer pays the paper price for. Figure 3.1 carries the machine, its instrument tags and each symbol the equations below use.

Table 3.1 sorts what the four loops reach. Mill units are kept throughout, so sheet properties are in g/m^2 and percent, flows in kg/min, speed in m/min, pressures in kPa gauge and temperatures in °C, with the clock in seconds because the scanner sets the time scale. Each conversion between a per-minute flow and a per-second balance is made once and at the point of use across the fourteen integrated states.

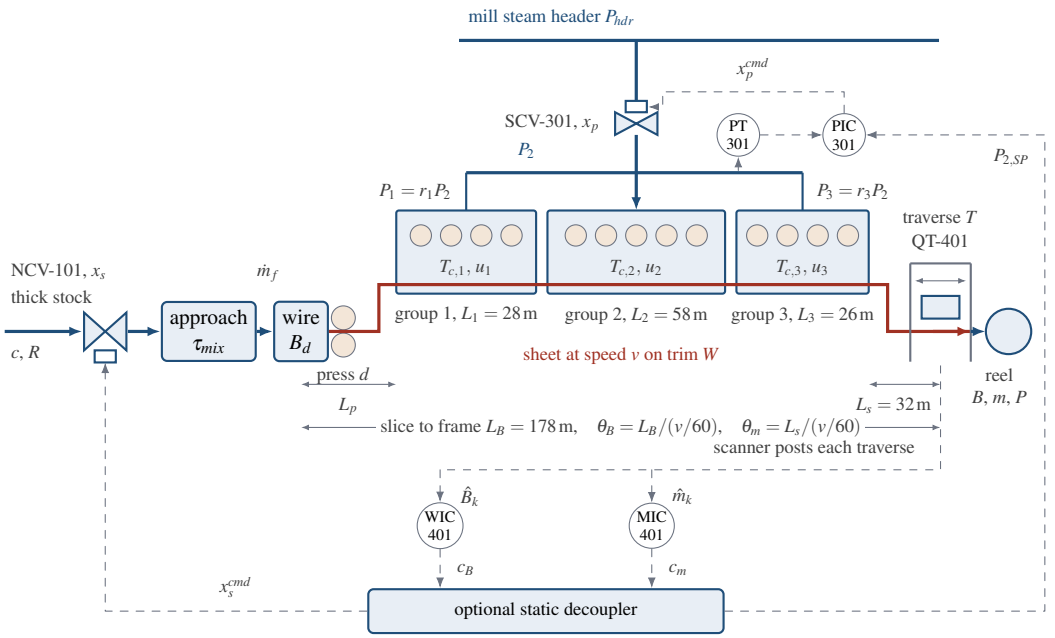

Figure 3.1 Fine-paper machine with the quality loops closed on one traversing gauge. Thick stock of consistency c and retention R passes valve NCV-101 and the approach system to lay down dry basis weight B_d on the wire, the press leaves the sheet at dryness d, and three groups of steam-heated cylinders at shell temperatures $T_{c,j}$ carry the sheet moisture ratio from u_P to u_3. Steam valve SCV-301 sets the main group pressure P_2 against mill header P_{hdr}, with the outer groups at fixed ratios r_1 and r_3. Scanning frame QT-401 posts the boxcar averages $\hat{B}_k$ and $\hat{m}_k$ once per traverse of period T, delayed by θ_B and θ_m. WIC-401 writes the stock valve and MIC-401 the group pressure setpoint that PIC-301 holds, and only PIC-301 reads a continuous transmitter.

3.3.1 Stock System and the Laid-Down Weight

The thick stock valve meters volume and the sheet is made of fiber, so the consistency of the stock is a disturbance and not a detail. The valve position, the approach system inventory and the forming response are three first-order lags in series, as shown in Equations (3.1) to (3.3).

$$\tau_v \frac{\mathrm{d}x_s}{\mathrm{d}t} = x_s^{cmd} - x_s \tag{3.1}$$

$$\tau_{mix} \frac{\mathrm{d}\dot{m}_f}{\mathrm{d}t} = Q_{th} \frac{x_s}{100} \frac{c}{100} - \dot{m}_f \tag{3.2}$$

$$\tau_{form} \frac{\mathrm{d}B_d}{\mathrm{d}t} = \frac{1000 R \dot{m}_f}{v W} - B_d \tag{3.3}$$

where x_s is the valve position in percent, Q_{th} is the thick stock flow at full lift in L/min, c is the consistency in percent, $\dot{m}_f$ is the oven-dry fiber reaching the headbox in kg/min, R is the overall retention, v is the machine speed in m/min, W is the trim in m and B_d is the dry basis weight in g/m^2. The driving term of (3.3) carries kg/min divided by m^2/min and multiplied by 1000 g/kg, so the two per-minute rates cancel and no factor of sixty belongs in it. The machine drive follows its own command with a lag of τ_{drive}, as shown in Equation (3.4).

$$\tau_{drive} \frac{\mathrm{d}v}{\mathrm{d}t} = v^{cmd} - v \tag{3.4}$$

Table 3.1 Variables of the Machine Direction Problem, With the Line-Up Every Scenario Opens On.

Symbol	Role	Nominal	Range
$\hat{B}$ (QT-401)	CV, sampled	80.00 g/m^2	45–130 g/m^2 SP
$\hat{m}$ (QT-401)	CV, sampled	5.50 %	3.5–9 % SP
P_2 (PT-301)	CV, continuous	306 kPa	60–480 kPa SP
P production	CV, $\hat{B}vW$	38.40 t/h	16–50 t/h SP
x_s (NCV-101)	MV	55.00 %	15–95 %
$P_{2,\mathrm{SP}}$	MV of MIC-401	306 kPa	60–480 kPa
x_p (SCV-301)	MV of PIC-301	60.00 %	0–100 %
v	MV of SIC-401	1000 m/min	400–1250 m/min
c consistency	DV, unmeasured	3.50 %	2.5–4.5 %
d press dryness	DV, unmeasured	43 %	36–48 %
P_{hdr}	DV, measured	550 kPa	350–700 kPa
R retention	DV, unmeasured	0.985	0.90–1.00
H hood humidity	DV, unmeasured	140 g/kg	60–260 g/kg
σ_B	DV, gauge noise	0.25 g/m^2	0–1.5 g/m^2
σ_m	DV, gauge noise	0.12 %	0–0.8 %
x_s state	state	55.00 %	0–100 %
$\dot{m}_f$	state	614.0 kg/min	≥ 0
B_d	state	75.600 g/m^2	≥ 1
v state	state	1000 m/min	20–1400 m/min
x_p state	state	60.00 %	0–100 %
P_2 state	state	306 kPa	0–900 kPa
$T_{c,1..3}$	states	112.9 / 138.0 / 125.5	°C, 10–260
$u_{1..3}$	states	1.029 / 0.139 / 0.058	kg/kg, 0.0005–6
u_P	state	1.326	kg/kg, 0.0005–6
ϕ threading	state	1	0–1

Equation (3.3) is the reason the fourth loop behaves as it does, because raising v at fixed $\dot{m}_f$ divides B_d by the same factor and delivers the same tonnage to the reel. Machine speed makes no production except through the basis weight loop.

3.3.2 Press and the Drying Stack

The press is algebraic. Its outgoing dryness is set by the operator through a disturbance slider with a small penalty for a heavier sheet, and the moisture ratio it delivers is approached with a lag, as shown in Equations (3.5) and (3.6).

$$d = \mathrm{clip}_{[0.20,0.70]}\left[\frac{D_p}{100}\left(1 - k_B\left(\frac{B_d}{B_{d,0}} - 1\right)\right)\right] \tag{3.5}$$

$$\tau_{press}\,\frac{\mathrm{d}u_P}{\mathrm{d}t} = \frac{1-d}{d} - u_P \tag{3.6}$$

where D_p is the press dryness setting in percent, k_B is the dryness lost per unit of relative basis weight, $B_{d,0}$ is the reference dry weight and u_P is the moisture ratio leaving the press in kilograms of water per kilogram of fiber. Every moisture state in the model is a dry-basis ratio, and the wet-basis moisture the industry quotes follows from Equation (3.7).

$$m = \frac{u}{1+u} \tag{3.7}$$

Confusing the two puts a water balance out by several percent of itself.

Each dryer group is a residence time and an evaporation flux, as shown in Equations (3.8) to (3.10).

$$t_j = \frac{L_j}{v/60} \tag{3.8}$$

$$e_j = h_d\,\mu_j\,\max(0,\ T_{c,j} - T_{ev})\left(1 - \mathrm{e}^{-\bar{u}_j/u_c}\right)\phi \tag{3.9}$$

$$\frac{\mathrm{d}u_j}{\mathrm{d}t} = \frac{u_{j-1} - u_j}{t_j} - \frac{e_j}{B_d} \tag{3.10}$$

where L_j is the sheet path length of group j in m, t_j is the residence time in s, e_j is the evaporation flux in grams of water per square meter of sheet path per second, h_d is the specific drying conductance, μ_j is a per-group multiplier, $T_{c,j}$ is the shell temperature, T_{ev} is the sheet evaporation temperature, u_c is the falling-rate scale, $\bar{u}_j = \frac{1}{2}(u_{j-1} + u_j)$ is the mean moisture ratio across the group, ϕ is the threading fraction and $u_0 = u_P$. Equation (3.10) is plug flow with a distributed sink, so the first term carries sheet in and out over the residence time and the second is the water boiled off. Evaluating the falling-rate factor at $\bar{u}_j$ rather than at the outlet matters, because a group that dries a great deal otherwise reports a conductance a third too low. The rate factor is 0.965 through group 1, 0.812 through the main bank and 0.246 through the after-dryer, which is why the last group owns the final moisture. The sheet does not sit at 100 °C, and its evaporation temperature is carried by the pocket air, as shown in Equation (3.11).

$$T_{ev} = T_{ev,0} + k_H\,(H - H_0) \tag{3.11}$$

where H is the hood exhaust humidity in g/kg and H_0 is its design value. A hood that is not carrying water away therefore removes drying capacity while nothing in the steam system looks wrong.

3.3.3 Cylinders and the Steam Side

The cylinder shells are the slow part of the plant. Each bank condenses steam on the inside, boils water off the outside and loses heat to the hood, as shown in Equations (3.12) and (3.13).

$$Q_j = e_j\,L_j\,W\,\Delta h_{ev} \tag{3.12}$$

$$C_j\frac{\mathrm{d}T_{c,j}}{\mathrm{d}t} = UA_j\,(T_{sat}(P_j) - T_{c,j}) - Q_j - UA_{amb}\,(T_{c,j} - T_{hood}) \tag{3.13}$$

where Q_j is the group duty in kW, Δh_{ev} is the heat per gram of water removed, C_j is the bank thermal mass in kJ/°C, UA_j is its condensing conductance and UA_{amb} is the loss to the hood. The flux of (3.9) is already per second, so (3.12) carries no time conversion. Each UA_j is set so that the group's design duty alone costs exactly 6 °C of steam-to-shell drop rather than chosen as a round number (the loss to the hood adds 0.3, 0.1 and 1.1 °C), which makes the shell time constants C_j/UA_j a consequence of the machine at 25, 19 and 84 s for the three groups, the after-dryer slowest because it is the most lightly loaded.

The three group pressures are recipe ratios of the one that is manipulated, the steam demand follows from the duties, and the valve delivers a pressure that droops with that demand, as shown in Equations (3.14) to (3.17).

$$P_1 = r_1 P_2, \qquad P_3 = r_3 P_2 \tag{3.14}$$

$$\dot{m}_{st} = \frac{60\sum_j Q_j}{h_{fg}} \tag{3.15}$$

$$\tau_{vp}\frac{\mathrm{d}x_p}{\mathrm{d}t} = x_p^{cmd} - x_p \tag{3.16}$$

$$\tau_{st}\frac{\mathrm{d}P_2}{\mathrm{d}t} = \mathrm{clip}_{[0,P_{hdr}]}\left[P_{hdr}\frac{x_p}{100} - k_d\,\dot{m}_{st}\right] - P_2 \tag{3.17}$$

where $\dot{m}_{st}$ is the steam condensed in kg/min, h_{fg} is the latent heat of the supply steam and k_d is the droop coefficient. The duties of (3.12) are in kJ/s and the operator reads steam in kg/min, which is where the factor of sixty in (3.15) comes from. The droop term in (3.17) is worth 24 kPa at the design 969 kg/min, and it is what turns a wetter sheet into a load disturbance that the pressure loop removes in seconds instead of something the moisture loop discovers through the gauge. Saturation temperature is a two-parameter fit to the steam tables, as shown in Equation (3.18).

$$T_{sat}(P) = 100.75 \left(\frac{P + 101.325}{100} \right)^{0.2549} \tag{3.18}$$

with P in kPa gauge and T_{sat} in °C. The fit is within 0.4 °C between 1.5 and 6 bar absolute, which is the whole range a dryer section uses.

The web itself is a state, so that a break and a rethread are events the rest of the model can see, as shown in Equation (3.19).

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{\phi_{tgt} - \phi}{\tau}, \qquad \tau = \begin{cases} \tau_{break} & \phi_{tgt} < \phi \\ \tau_{thread} & \text{otherwise} \end{cases} \tag{3.19}$$

The asymmetry is deliberate, because a web parts in seconds and takes a quarter of a minute to get back across the machine.

3.3.4 What the Sheet Is Worth and What the Gauge Reports

Basis weight is sold conditioned, so what the beta gauge weighs and what the invoice says is the fiber plus the water still in it, as shown in Equation (3.20).

$$B = \frac{B_d}{1 - m}, \qquad P = \frac{B\, v\, W\, 60}{10^6}\, \phi \tag{3.20}$$

where B is the conditioned basis weight in g/m^2, m is the reel moisture from (3.7) evaluated on u_3, and P is the saleable production in t/h. The dry weight B_d is fixed on the wire and nothing in the dryer section changes it, so (3.20) is the entire mechanism by which the steam valve becomes a basis weight actuator.

The gauge reads that sheet twice over. The two properties are set at different places on the machine, so they arrive at the frame with different transport delays, both computed from the live speed, as shown in Equation (3.21).

$$\theta_B = \frac{L_B}{v/60}, \qquad \theta_m = \frac{L_s}{v/60} \tag{3.21}$$

with $L_B = 178\,\mathrm{m}$ from the slice and $L_s = 32\,\mathrm{m}$ from the last cylinder, giving 10.68 s and 1.92 s at design speed. Slowing the machine makes both longer, which is the reverse of what an operator slowing down to sort out a problem expects. The head then averages one traverse, posts the average and holds it, as shown in Equation (3.22).

$$\hat{B}_k = \frac{1}{T} \int_{t_k - T}^{t_k} B(\tau - \theta_B)\, d\tau + \varepsilon_k \sigma_B, \qquad \hat{B}(t) = \hat{B}_k \tag{3.22}$$

for $t_k \le t < t_{k+1}$, with $t_{k+1} - t_k = T$ the scan period, ε_k drawn from the unit normal distribution once per completed pass and σ_B the gauge repeatability. The moisture reading follows the same expression with θ_m and σ_m. Per-pass noise stands in for cross-direction profile variation aliasing into the machine-direction reading, which is the dominant source of scatter on a real gauge [14]. A pass that accumulated sheet for less than half a period posts nothing at all and the previous value stands, which makes a sheet break a loss of measurement rather than a reading of zero.

Table 3.2 Model Parameters. Adjustable Constants Are Displayed in the Model Panel and May Be Changed While the Machine Runs.

Symbol	Meaning	Value	Unit
Adjustable, displayed in the model panel			
W	trim width at the reel	8.0	m
Q_{th}	thick stock flow at full lift	31896.6	L/min
τ_v	stock valve travel	3.0	s
τ_{mix}	approach system mixing	22	s
τ_{form}	forming and headbox	6.0	s
τ_{drive}	machine drive ramp	8.0	s
τ_{press}	press section response	6.0	s
k_B	dryness lost per unit weight	0.05	—
h_d	specific drying conductance	0.39671	$\mathrm{g/m^2 s\,^\circ C}$
μ_3	after-dryer multiplier	0.85	—
u_c	falling-rate scale	0.35	kg/kg
$T_{ev,0}$	evaporation temp. at H_0	78.0	°C
k_H	evaporation temp. per H	0.05	$^\circ\mathrm{C}/(\mathrm{g/kg})$
Δh_{ev}	heat per gram of water	2.70	kJ/g
C_1	group 1 bank thermal mass	34000	kJ/°C
C_2	group 2 bank thermal mass	78000	kJ/°C
C_3	group 3 bank thermal mass	31000	kJ/°C
UA_{amb}	bank loss to the hood	6.0	kW/°C
T_{hood}	hood temperature	55	°C
r_1	group 1 pressure ratio	0.30	—
r_3	group 3 pressure ratio	0.63	—
τ_{st}	group pressure response	14	s
τ_{vp}	steam valve travel	4.0	s
k_d	pressure droop with demand	0.024752	$\mathrm{kPa}/(\mathrm{kg/min})$
h_{fg}	latent heat of supply steam	2135	kJ/kg
τ_{Pm}	pressure transmitter lag	1.5	s
τ_{thread}	time to thread the tail	15	s
—	basis weight tolerance	0.8	$\mathrm{g/m^2}$
—	moisture tolerance	0.30	%
—	runnability limit m_{wet}	8.0	%
—	brittle limit m_{dry}	3.2	%
—	shell picking limit	142	°C
—	paper, fiber and steam price	820 / 540 / 22	$/t
—	off-specification discount	0.55	—
Hidden, never displayed			
$B_{d,0}$	reference dry basis weight	75.6	$\mathrm{g/m^2}$
μ_1, μ_2	group 1 and 2 multipliers	1.00	—
H_0	design hood humidity	140	g/kg
UA_1	group 1 shell conductance	1347.10	kW/°C
UA_2	group 2 shell conductance	4032.91	kW/°C
UA_3	group 3 shell conductance	368.84	kW/°C
τ_{break}	time for the web to part	2.0	s

Sheet path lengths are $L_p = 34\,\mathrm{m}$ from the slice to the first cylinder, $L_{1,2,3} = 28$, 58 and 26 m through the three groups and $L_s = 32\,\mathrm{m}$ to the frame, for $L_B = 178\,\mathrm{m}$ in total.

3.3.5 Constants, Startup and Numerics

Every constant is in Table 3.2. Thirty-six sit in the model panel and can be moved while the machine runs. The remaining eight are never shown, because they are the calibration constants a student is meant to infer rather than read off the screen.

No scenario begins on a drifting machine. The initial condition solves Equations (3.2) to (3.17) for zero derivatives at the operator's three handles by damped successive substitution, which converges because the plant contains one dominant negative feedback. A hotter cylinder dries faster, which draws more heat, cools the shell and drops the group pressure through the droop term of (3.17). At the design line-up that solution returns thick stock valve 55.00 %, steam valve 60.00 %, group pressures 92, 306 and 193 kPa, shells at 112.9, 138.0 and 125.5 °C, and dryness running 43.0, 49.3, 87.8 and 94.5 % down the machine. The section removes 95.8 g/m^2 of water, which is 766 kg/min for 34.5 MW of duty split 8.1, 24.2 and 2.2 MW, and consumes 58.2 t/h of live steam for 38.40 t/h of paper. The application's verification suite checks that line-up against the model, and it checks separately that the residual of every state derivative is below 10^{-9} relative at the stored initial condition.

Time stepping is Runge-Kutta of fourth order at a fixed 0.25 s interval, short enough that several steps fall inside the 1.25 s pass of the sheet through the after-dryer at top speed. The regulatory scan is 0.5 s, the historian records two samples per simulated second, and the gauge scatter is seeded with 20260913. The scan period of the gauge is a setting, adjustable from 4 to 40 s. Section 1.6.1 explains the engine step, the seed and the speed multiplier, which here runs from 1× to 30× and changes only the pace of the wall clock.

3.3.6 Assumptions

Detail is concentrated at the instrument and cut back everywhere else. Each of the eight simplifications below stops being true somewhere, named with it.

The cross direction is absent. This machine has one line of sheet, so the application shows exactly what a traversing measurement costs in the machine direction and cannot show the profile control the same head exists to serve [4]. The absence also removes the real reason a mill hesitates to speed the traverse up, which survives as an argument and not as a measurement.

Per-pass noise stands in for profile variation. Aliasing of a real cross-direction profile into the machine-direction reading is systematic rather than Gaussian, so a normal draw per pass has the right magnitude with the wrong structure. The gauge is otherwise perfect, with no standardization cycle, no sheet flutter, no cross-sensitivity between the beta and infrared measurements and no drift.

Three dryer groups and no condensate system. A machine of this size carries forty to sixty cylinders in five or six groups, each with its own pressure, syphons and differential controller. Here the outer two groups follow the main bank by the fixed ratios of (3.14), and a flooded group cannot be produced.

Drying is one conductance and a falling-rate factor. There is no sheet temperature state, no pocket ventilation balance and no split between the contact and open-draw parts of a group, and h_d was fitted so the design line-up gives a plausible dryness profile rather than a measured one [5, 6].

The wire and press are algebraic. Equation (3.5) is a slider with a small basis weight correction, so the application shows that a plugged felt is the largest routine disturbance on the machine and not why it plugged. The press and dryer balances also divide by the dry weight on the wire at that instant rather than the weight formed 34 to 146 m upstream of each group, so the extra water a stock move brings reaches the frame through the last 32 m about 8 s before the fiber that carries it. This lumped approximation is retained for the scanner and interaction lessons, so stock-step water arrival is not a physical transport prediction.

Consistency, retention and hood humidity are sliders. The white water loop, the save-all, the broke system and the wet-end chemistry that set them are absent, which matters because their true dynamics recirculate over tens of minutes and a consistency upset on a real machine comes back around.

A sheet break is a threading fraction. Equation (3.19) falls in 2 s and recovers in 15 s, while the real event is violent and threading is manual and takes minutes. Faithful is the part the control system experiences, that the measurement stops, the controllers hold and the cylinders heat up with nothing to dry.

The arithmetic is too reliable. A frozen step, a frozen scan and one seed return the identical trajectory to any student on any computer, which no real console and no real network will do.

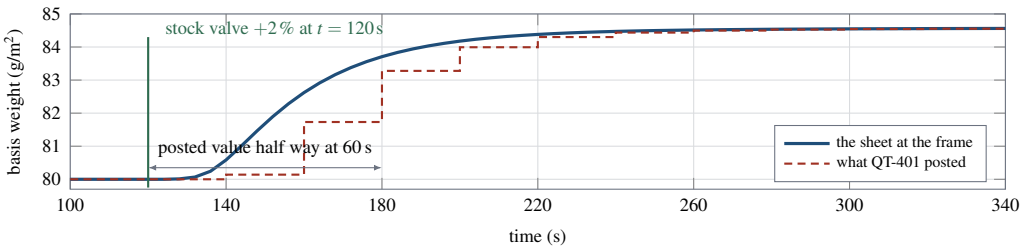

Figure 3.2 One open-loop stock step read through the scanning gauge, recorded in scenario 2. The very thick solid line is the sheet at the frame and the dashed staircase is the value QT-401 posted. The sheet at the frame reaches 63 % of its 4.58 g/m^2 move 43.5 s after the step and 95 % at 95 s, while the posted staircase is half way at 60 s.

3.4 Dynamic Analysis

3.4.1 The Effective Dead Time of a Scanning Gauge

Equation (3.22) contains three sources of phase lag and they add. The transport delay θ is a pure delay. The boxcar average over one traverse has its centroid half a period behind the instant of posting, so on any smooth trajectory the posted value is the sheet as it was $T/2$ before the pass ended. Between passes the number does not move, so the age of the value the controller acts on is distributed uniformly over $[0, T]$ and averages $T/2$ again. The effective dead time is therefore Equation (3.23).

$$\theta_{\text{eff}} = \theta + \frac{T}{2} + \frac{T}{2} = \theta + T \tag{3.23}$$

At the design line-up (3.23) gives $1.92 + 20 = 21.92$ s for moisture and $10.68 + 20 = 30.68$ s for basis weight. For moisture that is eleven times the transport delay by itself, and the instrument rather than the machine owns the number. The application's verification suite checks both halves of (3.23) on a ramped sheet. The posted value lags the truth by 9.74 s at the instant of posting, a quarter second short of $T/2$ because the boxcar sums the 0.5 s scans, and the mean age of the held value over a 300 s window is 10.13 s.

🕮 **Pause.** Try Exercise 3.1 before reading on.

The measured behavior of a single step is shown in Figure 3.2. The thick stock valve is stepped from 55.00 to 57.00 % at $t = 120$ s, which is worth 4.58 g/m^2 at the reel once settled. The dry weight at the frame begins to move 11.5 s later, because the fiber being weighed was formed 178 m upstream and the machine is doing 16.67 m/s. The conditioned weight starts about 5 s after the step, because the wetter sheet arrives first (Section 3.3.6), but only 0.06 g/m^2 of it comes before the fiber. The posted number does not move until the pass that ends at $t = 140$ s, 20 s after the step, and then it moves only to 80.140 g/m^2, the average of a traverse that was mostly old sheet.

Reading the two traces of Figure 3.2 against each other separates the plant from the instrument. The sheet is half way 36 s after the step and the posted number at 60 s, the sheet is 90 % of the way at 76 s and the posted number at 100 s, and both gaps of 24 s are within a few seconds of the 20 s that (3.23) attributes to the instrument. The posted trace is a sequence of averages rather than a sampled version of the true one. The 81.73 g/m^2 posted at $t = 160$ s is the sheet as it passed the frame 10 s earlier, while the paper under the head at that instant was already 82.66 g/m^2.

3.4.2 The Two-by-Two Gain Matrix

The second manipulated variable of the quality layer is the group pressure setpoint rather than the steam valve, because PIC-301 holds that setpoint against the header and the droop. Perturbing the model's own steady state by ± 0.25 % of stock valve at constant group pressure, and by ± 4 kPa of

group pressure at constant stock valve, gives the gain matrix of Equation (3.24).

$$\mathbf{K} = \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix} = \begin{bmatrix} +2.207 & -0.0514 \\ +0.889 & -0.0607 \end{bmatrix} \tag{3.24}$$

where the first row is basis weight in g/m^2 and the second is moisture in percentage points, the first column is per percent of stock valve and the second is per kPa of group pressure. Inverted, 19.5 kPa of group pressure buys a gram per square meter and 16.5 kPa buys a point of moisture. None of the four connections is small.

The signs are the interesting part. More stock is heavier because more fiber arrives per unit area, and wetter because the same dryer section has more water to remove, and by (3.20) the extra water is weighed as well, so the weight rises twice. More steam pressure is drier and also lighter, with not one gram of fiber different. Both effects are visible from the keyboard. A 2 % increase in stock valve at a fixed steam valve makes the sheet 4.58 g/m^2 heavier and 1.86 points wetter, of which only 2.75 g/m^2 is fiber, and a 6 % increase in steam valve makes it 1.79 points drier and 1.48 g/m^2 lighter. The manual figure of 4.58 g/m^2 exceeds twice K_{11}, which is 4.41 g/m^2, because a fixed valve is not a fixed pressure. The wetter sheet draws more steam and the droop of (3.17) lets the group pressure fall, which makes the sheet wetter and heavier again.

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

The relative gain of the natural pairing follows from (3.24), as shown in Equation (3.25).

$$\Lambda_{11} = \left[1 - \frac{K_{12}K_{21}}{K_{11}K_{22}}\right]^{-1} = 1.517 \tag{3.25}$$

By the reading of Section 1.4.5, a relative gain between one and two says that pairing basis weight with the stock valve and moisture with the group pressure is correct and that the interaction is severe, because each loop slows the other down [3]. The other pairing has $\Lambda_{12} = 1 - 1.517 = -0.517$, so the gain of one loop changes sign between the other loop being in manual and in automatic, and the pair cannot be stabilized by any tuning of the two controllers. The application's verification suite recomputes Λ_{11} from the model rather than quoting it and requires the value to fall within 0.02 of 1.517.

3.4.3 Static Decoupling From a Signal That Already Exists

The two controllers reach each other only through the gauge, which is 21.9 s behind, although the correction each one needs is already available from the other controller's output at the instant it is computed. A simplified static decoupler passes each demand straight through and adds the other through the steady-state cross gain, as shown in Equation (3.26).

$$\begin{aligned} x_s^{cmd} &= \text{clip}_{[15,95]}\left[c_B + D_{12}\left(c_m - P^{ref}\right)\right] \\ P_{2,\text{SP}} &= \text{clip}_{[60,480]}\left[c_m + D_{21}\left(c_B - x_s^{ref}\right)\right] \end{aligned} \tag{3.26}$$

where c_B and c_m are the two controller outputs, $x_s^{ref} = 55\,\%$ and $P^{ref} = 306\,\text{kPa}$ are the design line-up the bias is taken about, the limits are in percent of stock valve and kPa, and the two gains are the ratios of Equation (3.27).

$$D_{12} = -\frac{K_{12}}{K_{11}} = 0.0233, \qquad D_{21} = -\frac{K_{21}}{K_{22}} = 14.64 \tag{3.27}$$

in percent of stock valve per kPa and in kPa per percent of stock valve respectively. These are the gains of Equation (1.8), and the simplified form is chosen for the reasons of Section 1.4.6: it needs only the two ratios and cannot introduce a pole of its own [15]. Chapter 2 uses one half of the same

structure on a plant whose gain matrix is triangular. An error in x_s^{ref} or P^{ref} is a constant bias that the two integrators remove, so both can be constants rather than values latched at commissioning. Equation (3.26) is static and therefore puts the right correction on at the wrong time, because the stock reaches the sheet through a 44 s path and the steam through a can bank whose slowest shell is 84 s. Section 3.6 measures what survives.

> ⧖ **HISTORY** **1965: The Coupling Found by the First Experiment**
>
> The engineers of that first computer-controlled machine did not know at the start that its two qualities were coupled. The first identification experiment showed it clearly. More thick stock flow raised both the moisture content and the dry basis weight, while a change of steam pressure moved the moisture but not the dry weight. They therefore controlled dry basis weight, estimated from the beta-ray and moisture readings, as a loop with one input, the thick stock flow. The moisture loop was given two inputs, with the stock flow used to remove the disturbance that the weight loop put into the moisture [8]. This chapter controls the conditioned weight instead, which couples the plant both ways.

3.5 Controller Design

3.5.1 The Discrete Algorithm

One shared proportional-integral-derivative implementation serves all four loops. It is the velocity form of Section 1.6.2, with the derivative on the measurement, and Equations (3.28) to (3.31) restate it in this chapter's symbols.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{3.28}$$

$$P_k = K_c e_k, \qquad D_k = -K_c \tau_D\, a\, \dot{\mathrm{PV}}_k \tag{3.29}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{3.30}$$

$$\begin{aligned} u_k^u &= u_{bias} + S_k + I_{k-1} \\ u_k &= \mathrm{clip}_{[u_{\min}, u_{\max}]}\,(u_k^u) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - u_k^u}{T_t}\right] \Delta t_k \end{aligned} \tag{3.31}$$

where $a = \pm 1$ is the action, S_k accumulates the proportional and derivative terms, u_k^u is the unsaturated demand, I_k is the integral state, T_t is the tracking constant of the unwinding path and Δt_k is the interval since that loop last executed. They are Equations (1.11) to (1.14), with S_k the running sum of the fully formed terms. Windup is prevented by back-calculation, the second term of the integral update in (3.31), with $T_t = \tau_I$ on every loop, and the transfer between manual and automatic is bumpless as described in Section 1.6.2. The shared algorithm carries derivative action, with an optional first-order filter on the measurement, but no loop here exposes a derivative time and τ_D is zero on all four, because the only directly measured continuous variable on the machine is a steam pressure whose plant is a 14 s lag behind a 4 s valve and the other three act on a staircase.

The four loops are configured as listed in Table 3.3. Output limits are hard clips and there are no controller rate limits, because the actuator dynamics of (3.1), (3.4) and (3.16) already supply the travel time. The clips of (3.26) act after the sum and outside both controllers, which see only their own outputs, so a decoupled command held at a limit by the other loop's term would not slow either integrator. No shipped case reaches those limits, and the decoupled group pressure setpoint peaks at 455 kPa in the challenge run with the suggested settings.

3.5.2 Two Clocks and an Event

PIC-301 and SIC-401 execute every 0.5 s. PIC-301 reads a transmitter that never stops reading, while the tonnage SIC-401 reads changes with the live speed and, through $\hat{B}$, only when a pass posts. WIC-401 and MIC-401 execute on the arrival of a posted reading and not otherwise, so between

Table 3.3 The Four Implemented Loops and Their Default Settings. Action $a = +1$ Is Reverse Acting and $a = -1$ Is Direct Acting in the ISA Sense.

Loop	CV	MV	K_c	τ_I	a
WIC-401	$\hat{B}$	x_s	0.32	70	$+1$
MIC-401	$\hat{m}$	$P_{2,\mathrm{SP}}$	18	90	-1
PIC-301	P_2	x_p	0.45	20	$+1$
SIC-401	P	v	16	240	$+1$
Swapped pairing, kept as a separate controller pair					
WIC-401s	$\hat{B}$	$P_{2,\mathrm{SP}}$	6	200	-1
MIC-401s	$\hat{m}$	x_s	0.35	200	$+1$

Gains carry the units of the pairing: % of stock valve per g/m^2, kPa per point of moisture, % of steam valve per kPa and m/min per t/h. Integral times are in s. Output limits are 15 to 95 % on x_s, 60 to 480 kPa on $P_{2,\mathrm{SP}}$, 0 to 100 % on x_p and 400 to 1250 m/min on v.

passes their outputs are frozen and the machine runs open loop for 19.5 s out of every twenty. Writing them instead as continuous controllers fed a staircase is the standard mistake, and it produces a loop that integrates one stale error forty times over.

The integration interval each quality controller is handed is the elapsed time since the previous posting, clamped as shown in Equation (3.32).

$$\Delta t_k = \min\left(t - t_{k-1},\ 2T\right) \tag{3.32}$$

The clamp exists for one specific failure. During a break the gauge posts nothing, so $t - t_{k-1}$ grows without limit, and a controller handed a five-minute interval when the first reading after threading arrives adds $K_c \tau_I^{-1}$ times the error times five minutes to its integral in one step and throws the valve to its stop. Two periods also keeps the explicit back-calculation in (3.31) stable, because the worst ratio $\Delta t_k / T_t$ is then 2×40 s over the 60 s floor on τ_I, which is 1.33 and below the value of two at which an explicit unwinding step diverges. The 60 s lower limit on the integral time of both quality loops exists for that reason and not because a faster reset would be badly tuned.

3.5.3 Response to a Lost Measurement

Three behaviors are commissioned around the loss of the gauge and each prevents a specific failure. The quality controllers hold their last output, which is what every industrial quality control system does, even though the held stock valve position of 55.13 % and group pressure setpoint of 311 kPa were chosen for a sheet that no longer exists. The production setpoint is parked on the measurement while the web is down, because a 38 t/h target against no sheet is a 38 t/h error that no actuator can act on. The production loop is re-closed at 95 % of a threaded web rather than at 50 %, because against half a sheet the inferred tonnage is half of nothing. In a counterfactual run of scenario 10 with the gate moved to 50 % in a copy of the controller, the speed command goes to its 1250 m/min stop at $t = 460.5$ s, 24 s before the web is 90 % across, against a peak of 1032 m/min with the shipped gate.

3.5.4 The Production Cascade

SIC-401 moves the machine speed to hold a tonnage it can only change through the basis weight loop, which follows from Equation (3.3). It measures that tonnage as a mill production computer does, from the posted weight, the live speed and the trim, as shown in Equation (3.33).

$$\hat{P} = \frac{\hat{B}\, v\, W\, 60}{10^6}\, \phi \tag{3.33}$$

Neither the consistency nor the retention enters it, because no instrument on the machine reads them, and a speed move shows in $\hat{P}$ at once while the lighter sheet it makes reaches $\hat{B}$ one effective dead

time later. Ten percent more speed at a fixed stock flow delivers exactly the same fiber to the reel and a sheet 9.3 % lighter, and holding the grade at that speed needs 1.10 times the stock valve position. The loop is therefore gated on the weight loop being in service, it is disabled under the swapped pairing where the stock valve carries moisture, and its integral time of 240 s is an order of magnitude longer than the scan period of the loop nested inside it.

3.5.5 Model Predictive Control

A fifth quality structure, listed last, never the application's default and pinned only in scenario 14, gives the thick stock valve and the group pressure setpoint to one model predictive controller (MPC). It acts only in automatic and decides once when a run is initialized at $t = 0$ and then at each posting of the gauge, so its period is the scan period, 20 s as shipped. Each decision plans $N = 40$ passes ahead (800 s at that scan) in 11 move blocks of $[1,1,1,1,2,2,3,4,6,9,10]$ passes, 22 decision variables in all, and only the first move is sent, held until the next posting. A break posts nothing and so triggers no decision, which holds both commands as the loops hold theirs. The receding-horizon structure is the one industrial packages share [16, 17], as set out in Section 1.6.3.

Its model comes from the balances of Section 3.3, differentiated numerically about the design line-up with the declared values of Table 3.2, the nominal disturbances of Table 3.1, the web threaded and the speed fixed at 1000 m/min, and discretized with a zero-order hold at the 0.5 s scan. Its sheet and steam states are x_s, $\dot{m}_f$, B_d, u_P, x_p, P_2, $T_{c,1..3}$ and $u_{1..3}$. PIC-301 is inside it as the shipped discrete PI on its own tuning behind the 1.5 s transmitter, so the second input is $P_{2,\mathrm{SP}}$ and not the valve. The transport delays of (3.21) become chains of 21 and 4 scans (10.5 s and 2.0 s), the weight path being a copy of the stock line driven by the stock command 21 scans earlier, and the gauge of (3.20) is linearized with sensitivities $\partial B/\partial B_d = 1 + u_3 = 1.0582$, $\partial B/\partial u_3 = B_d = 75.600\,\mathrm{g/m^2}$ and $\partial m/\partial u_3 = 100/(1+u_3)^2 = 89.30$ points. The pass average and hold of (3.22) are composed exactly over one pass, as shown in Equation (3.34).

$$\begin{aligned} \mathbf{x}_{j+1} &= \mathbf{A}_h \mathbf{x}_j + \mathbf{B}_h \mathbf{v}_k + \mathbf{E}_h \mathbf{d}, \qquad j = kn, \ldots, kn+n-1 \\ \hat{\mathbf{y}}_{k+1} &= \frac{1}{n} \sum_{j=kn+1}^{kn+n} \mathbf{C}_g \mathbf{x}_j, \qquad n = T/h \end{aligned} \tag{3.34}$$

where $h = 0.5\,\mathrm{s}$, $\mathbf{x}$ collects the sheet and steam states, the PI integral, the transmitter and the two delay chains, $\mathbf{v}_k = [x_s^{cmd}, P_{2,\mathrm{SP}}]^\top$ is the decision held over pass k, $\mathbf{d}$ holds the estimated offsets of consistency, press dryness and header pressure, $\mathbf{C}_g$ is the linearized gauge and $\hat{\mathbf{y}} = [\hat{B}, \hat{m}]^\top$ is the next posting, all in deviations from the design line-up. With the stock decisions still in the weight chain carried as states, the plan's model has 28 states at the shipped scan, and the same composition serves any period from 4 to 40 s. Its steady gains, 2.2070, −0.0514, 0.8889 and −0.0607, round to those of (3.24), and steps on the running plant give 2.2067, −0.0514, 0.8877 and −0.0607. Started at the design point, it predicts the next 30 postings after a 1 % stock step within 1.5 % of the posted move and after a 10 kPa pressure step within 3.8 %, at scan periods of 4, 8, 20 and 40 s.

The quadratic program minimizes Equation (3.35) over the 22 block values.

$$\begin{aligned} J = &\sum_{k=1}^{N} \left[\left(\frac{\hat{B}_k - B^{\mathrm{SP}}}{0.8} \right)^2 + \left(\frac{\hat{m}_k - m^{\mathrm{SP}}}{0.30} \right)^2 \right] \\ &+ \sum_{j=1}^{11} \left[\lambda_s \Delta x_{s,j}^2 + \lambda_P \left(\frac{\Delta P_{2,\mathrm{SP},j}}{10} \right)^2 \right] \\ &+ \sum_{i=1}^{3} \left(\rho_i s_i^2 + \sigma_i s_i \right) \end{aligned} \tag{3.35}$$

Here $\hat{B}_k$ in g/m² and $\hat{m}_k$ in percent are predicted postings, Δ is the change between successive blocks with the first taken from the command in force, $\lambda_s = 1$ per squared percent of stock valve and $\lambda_P = 3$

per $(10\,\mathrm{kPa})^2$. Each slack $s_i \geq 0$ widens one soft limit over the whole horizon. The postings are held within the declared tolerances of $0.8\,\mathrm{g/m^2}$ and 0.30 points of their setpoints, widened by $0.8 s_1$ and $0.30 s_2$, and the steam valve that PIC-301 is predicted to command is held at or under $100\,\%$, widened by $10 s_3$ percent, at prices $\rho = (10, 10, 1000)$ and $\sigma = (1, 1, 100)$. The hard limits are the loops' own clamps, 15 to 95 % of stock valve and 60 to 480 kPa of pressure setpoint, and PIC-301, SIC-401, the hold on a lost measurement and every clamp stay in service beneath the plan. Neither structure receives a preview of a setpoint change.

Offset is removed by estimating disturbances [18]. A steady-state Kalman filter, with its gain from the discrete Riccati equation of the model sampled at the pass and augmented with the pass average, carries the consistency and press dryness offsets as random walks of 0.01 and 0.05 percentage points per pass, with the declared repeatabilities σ_B and σ_m as measurement noise. It advances every 0.5 s on the commands sent and corrects at each posting under every structure, so the MPC takes over warm. A second filter at 0.5 s reads PT-301 and estimates x_p, P_2 and a header offset, so that each plan starts from a steam side consistent with the transmitter. Any plan that fails the checks of optimality, finiteness and limits is discarded, and the weight and moisture controllers, which track every command the MPC sends, then act on that pass as two loops until a valid plan returns. The plant runs on the run's own parameters and disturbances, the gauge and the transmitter impose their averaging, hold, lag and noise, and the MPC reads only the two posted values, whether a pass has posted, the scan period, the PT-301 reading, the commands in force, the setpoints, its own settings, the tuning of PIC-301 and the declared constants. It never reads the speed or the tonnage.

The two move suppressions and the two random-walk sizes were chosen once, on scenario 14, from a grid of 369 settings, as the one making the least off-specification paper among the 62 that pass no more gauge noise into the sheet over the last ten minutes than the shipped decoupler. The rule was first applied before the correction that gave SIC-401 the measurement of (3.33) in place of the true fiber flow, and then once more on the corrected application. The earlier choice ($\lambda_s = 100$, $\lambda_P = 30$, 0.05 and 0.8 percentage points) makes 6.09 t of off-specification paper on scenario 14 against 2.68 t for the settings shipped. A decision takes 0.56 to 0.61 ms on average in Node.js.

3.6 Results

The fifteen scenarios of the app are ordered the way a laboratory would work through them, from the machine by hand (scenario 1) through the gauge, the scan period, the interaction, the pairing and the decoupler (scenarios 2 to 8) to a grade change, a sheet break, three routine upsets, the predictive controller and a graded challenge (scenarios 9 to 15). Each carries its own mode, structure, tuning, disturbances, timed events and completion rule, as Section 1.6.4 describes for every app, and every number in Tables 3.4 to 3.7 was taken from the running app. The screen during the grade move of scenario 6 is reproduced in Figure 3.4.

3.6.1 The Machine With a Clean Gauge and a Noisy One

Scenario 3 closes both quality loops on an undisturbed machine with the gauge repeatability set to zero. Nothing happens. Over 18 min and 53 posted traverses the stock valve does not move by a thousandth of a percent and the sheet holds $80.000\,\mathrm{g/m^2}$ and $5.500\,\%$, which is the only time in the application that it is true. With the default basis weight repeatability of $0.25\,\mathrm{g/m^2}$ set before Run, the stock valve command accumulates $4.79\,\%$ of travel over the same window and the sheet, which was flat, wanders between 79.77 and $80.17\,\mathrm{g/m^2}$. Doubling the repeatability to $0.5\,\mathrm{g/m^2}$ doubles that travel to $9.58\,\%$ and spreads the sheet from 79.54 to $80.33\,\mathrm{g/m^2}$. Over five other seeds the travel at $0.25\,\mathrm{g/m^2}$ runs from 3.48 to $5.02\,\%$ and always doubles with the noise. The variation in the paper is of the same order as the noise in the gauge, because one diagonal stripe across an eight-meter sheet is not the sheet and the loop transmits the difference into the product.

Table 3.4 Scan Period at a Fixed Moisture Tuning of 44 kPa per Point and 60 s. Peak-to-Peak Is the Moisture Swing Read Over 700 to 1400 s, With the Scan Period Set Before Run. Sheet Averaged Is the Paper One Traverse Integrates Over at 1000 m/min.

Scan period T (s)	30	20	10	8
θ_{eff}, moisture (s)	31.9	21.9	11.9	9.9
Peak-to-peak, scenario 3[a] (points)	4.49	2.62	0.18	0.07
Peak-to-peak, scenario 4[b] (points)	12.6[c]	8.72	0.33	0.12
Sheet averaged (m)	500	333	167	133

[a] Weight loop at 0.32 and 70 s, moisture setpoint 5.5 to 5.0 % at 200 s, the application's own verification run. [b] Weight loop at 0.50 and 60 s, setpoint move at 120 s. [c] Unstable, the steam valve running between its stops.

3.6.2 Scan Period Against Loop Stability

Scenario 4 loads the settings an internal model control rule returns when the dead time entered is the 1.92 s of transport rather than the 21.9 s the instrument costs, which is a moisture gain of 44 kPa per point against the tuned 18 kPa. A half-point moisture setpoint move at $t = 120$ s then produces an oscillation that grows, from 1.35 points peak to peak in the first cycle to 8.26 points in the eleventh, so that over the next twenty minutes the moisture spans 1.18 to 9.90 % about a 5.00 % setpoint, with a period of 100 s. That period is 4.5 times the 21.9 s of Equation (3.23), where a dead-time-dominant loop at the edge of stability oscillates, and the growth places this one beyond the edge. The basis weight goes around with it between 76.70 and 83.97 g/m^2 although nobody moved that setpoint. Each pass the controller sees an error that is largely its own last correction not having arrived, and corrects again on the strength of it. The weight loop of scenario 4 is tuned on the transport delay as well (0.50 % per g/m^2 and 60 s), and through the interaction it triples the ring. With that loop at its tuned 0.32 and 70 s the same moisture tuning swings 2.84 points over 700 to 1400 s, against 8.72 points.

Scenario 5 leaves that tuning alone and speeds the head up from twenty seconds a pass to eight at $t = 700$ s. The growth stops and the oscillation dies away, falling from 3.24 points peak to peak over 400 to 700 s to 1.24 points over the three minutes that begin three minutes after the change and roughly halving every three minutes after that, and the moisture stays within 0.05 points of 5.00 % from 937.5 s after it, with no gain touched. Run the other way, with 8 s passes slowed to thirty within a minute of the setpoint move, the moisture is 2.24 points from setpoint by 700 s. Table 3.4 holds the moisture tuning at 44 kPa per point and 60 s and sweeps the scan period alone.

> **TRY IT IN THE APP** **5 · Scan Period Against Loop Stability**
>
> Load the scenario, press Run in Auto and watch the Moisture strip. The loop rings with the tuning of scenario 4 and the swing grows. At 700 s the head is speeded up from twenty seconds a pass to eight, no gain is touched, and the oscillation dies away until the moisture stays within 0.05 points of 5.00 %.

Reading Table 3.4 across is unambiguous. Halving the traverse from 20 to 10 s turns a growing swing of 2.62 points into a decaying 0.18 points, a factor of fifteen, and in scenario 4 cuts it from 8.72 to 0.33 points, on a plant that did not change and a controller that was not retuned. At thirty seconds a pass the scenario 4 tuning does not hold the sheet at all. This is the highest-value change available on a quality control system and it is a genuine trade, because the same head also measures the cross-direction profile and the sheet it averages over shrinks with it, from 500 m per pass at 30 s to 133 m at eight. Machine-direction control wants the head to run and cross-direction control wants it to dawdle [4].

Table 3.5 An Eight Gram Grade Move With Four Decoupler Configurations. Times Are Measured From the Setpoint Move and Off-Specification Tonnage Is Read at $t = 900$ s, by Which Point It Has Stopped Accruing in Every Column.

	Neither	D_{12} only	D_{21} only	Both
D_{12} (%/kPa)	0	0.023	0	0.023
D_{21} (kPa/%)	0	0	14.5	14.5
Peak moisture error (points)	+1.80	+1.92	+0.24	+0.23
Time of that peak (s)	62	62.5	102.5	103
Moisture inside 0.3 points (s)	214.5	177.5	0	0
Weight first in band (s)	87.5	75	108.5	111
Weight settled in band (s)	87.5	202.5	477.5	496
Off-specification paper (t)	2.38	2.20	3.14	4.54

✦ DID YOU KNOW? A Faster Head Beats a Retuned Gain

From the simulator: with the moisture tuning held at 44 kPa per point, the setpoint move of scenario 3 leaves a swing of 4.49 points at a 30 s pass and 0.07 points at 8 s. Each reading then averages 133 m of sheet instead of 500 m (Table 3.4).

3.6.3 Interaction and What Half a Decoupler Is Worth

Scenario 6 raises the basis weight setpoint from 80 to 88 g/m^2 with two independent loops in service and the moisture setpoint untouched, and scenario 8 repeats it with the decoupler of Equation (3.26) switched on. Table 3.5 reports all four combinations of the two decoupler gains, every figure read on the sheet at the frame by $t = 900$ s against a setpoint move at $t = 120$ s. The production loop stays in service, so the heavier sheet also slows the machine toward the speed that holds 38.4 t/h, 921 m/min by 900 s.

Two independent loops make an 8 g/m^2 grade move into a 1.80 points moisture excursion that peaks 62 s after the setpoint moved, about three passes of the head, and that is not back inside three tenths of a point for good until 214.5 s. The full decoupler cuts the excursion to 0.23 points, and the moisture never leaves that band. The two gains are not equally worth having. Turning D_{21} off returns the excursion to 1.92 points while turning D_{12} off leaves it at 0.24 points, so the stock-to-steam half carries essentially all of the benefit to the moisture and this measurement says which way to run a one-way decoupler. The D_{12} only column makes 2.20 t of off-specification paper against 2.38 t with no decoupling, with the same moisture excursion within 0.12 points.

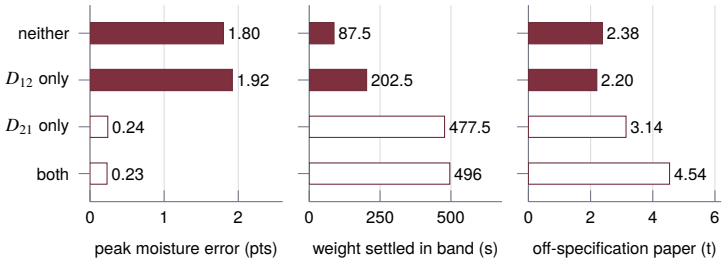

Figure 3.3 The eight gram grade move of scenarios 6 and 8 under the four decoupler configurations of Table 3.5. Dark bars run without the stock-to-steam gain D_{21} and white bars with it. Notice that D_{21} alone cuts the moisture peak from 1.80 to 0.24 points, while the same white bars show the weight settling later and more off-specification paper with the production loop in service.

> **TRY IT IN THE APP** **8 · A Decoupler**
>
> Load the scenario, press Run in Auto and watch the Moisture strip after the weight step at 120 s: the peak error is 0.23 points. Switch to Advanced, press Reset, set Stock-to-steam decoupling to 0 and run again. The peak returns to 1.92 points, more than with no decoupler at all (Table 3.5).

The weight loop pays for it. Decoupling makes the weight slower to reach its band, at 111 s against 87.5 s, because the stock valve now carries a steam-driven bias as well as its own controller's output, and the weight then overshoots to 88.94 g/m^2 and stays just above its band from 313.5 to 616 s. Over that interval the production loop is still slowing the machine, from 954 m/min at 300 s to 931 m/min at 600 s, and by (3.3) a slower machine lays a heavier sheet on the wire at the same stock flow. The weight is settled in its band at 496 s against 87.5 s, and the off-specification tonnage rises from 2.38 to 4.54 t. With the production loop in service, decoupling on this plant is the familiar trade between the two loops, a faster moisture loop bought with a slower weight loop. Doubling D_{21} to 29.25 overshoots the other way, to 1.77 points dry, and has made 8.64 t of off-specification paper by 900 s and is still making it. Figure 3.3 draws the four configurations side by side.

The residual 0.23 points is not a tuning error. Equation (3.26) is static, so it puts the right amount of steam on at the wrong time, and the mismatch between a 44 s stock path and an 84 s dryer path is what remains. The correction is nonetheless made from a signal the control system already holds instead of one that will not exist for another 21.9 s.

3.6.4 A Pairing That Cannot Be Tuned

Both connections in (3.24) are large, so both pairings are physically available. Basis weight responds to steam pressure at 19.5 kPa per gram per square meter, moisture responds to the stock valve at 1.12 % of valve per point, and step testing either loop with the other in manual returns a clean first-order plus dead-time fit and a sensible tuning. Scenario 7 puts both in service and dilutes the stock chest from 3.50 to 3.35 % consistency at $t = 150$ s. Twenty-two minutes after the dilution the natural pairing has the sheet back at 80.0 g/m^2 and 5.49 %, and the swapped pairing is at 72.8 g/m^2 and 13.12 % and still walking, with the stock valve down at 47.5 % and the group pressure setpoint down at 117 kPa. Integrated over the same 1500 s run, with a point of moisture weighted as four grams, the absolute error of the swapped pairing is 15.5 times that of the natural one, and it has made 13.6 t of off-specification paper against 2.1 t. That is what $\Lambda_{12} = -0.52$ means in operation. Each swapped loop has its own Auto and Manual switch, and with the weight loop alone in Manual the group pressure setpoint holds at 306 kPa while the moisture loop on the stock valve is stable on its own. The moisture falls to 3.87 % and climbs back through 5.00 % at 1200 s, slowly because that loop is tuned at 0.35 % per point and 200 s.

Table 3.6 Three Routine Upsets Open Loop and Closed Loop. Open Loop Is the Steady State With Both Valves Held. Recovery Is the Time From the Upset Until the Moisture at the Frame Is Back Inside 0.3 points for Good, Except in the Dilution Column Where It Is the Basis Weight Returning to Its 0.8 g/m^2 Band.

	Header dip 550 to 380 kPa	Felt plug 43 to 41 %	Chest dilution 3.50 to 3.30 %
Open-loop moisture (points)	8.06	3.43	−2.7
Open-loop weight (g/m^2)	—	—	−6.6
Closed-loop peak moisture (%)	6.60	8.23	3.59
Closed-loop peak weight (g/m^2)	—	—	73.2
Recovery (s)	117	327	328
Actuator that moved	x_p	$P_{2,\mathrm{SP}}$	x_s
from	61 %	310 kPa	55 %
to	85.6 %	354 kPa	58.9 %

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

3.6.5 Disturbances the Loops Can and Cannot See

Table 3.6 collects the three routine upsets, each with its open-loop effect on the sheet and its effect with the full architecture of Table 3.3 in service. These scenarios run with the default gauge noise, and over three other seeds the closed-loop peaks move by at most 0.18 points and 0.14 g/m^2 and the recovery times by up to 57 s.

The header dip of scenario 12 is the one the machine sees. Left alone it is worth eight points of moisture, because 206 kPa in the main group instead of 306 gives a sheet leaving at 13.56 % that would not survive the calender. Instead PT-301 sees the sag within a second, the steam valve runs from 61 to 93.7 % before settling at 85.6 %, and the sheet peaks 1.1 points high, is back inside 0.3 points for good in 117 s and within 0.15 points of setpoint in 285.5 s. Eight points of upset became one, and the one that got through did so only while the valve was traveling. That is the argument for a regulatory layer underneath a quality layer, because the same disturbance reaching the sheet unmeasured would be discovered a scan period later.

The felt plug of scenario 13 is the one nothing measures. Two points of press dryness is 3.43 points of reel moisture open loop, larger than any setpoint move anybody makes deliberately, and closed loop the sheet still peaks at 8.23 % thirty-three seconds after the felt goes, above the 8 % runnability limit. The chest dilution of scenario 11 is faster than the instrument, because the dry weight on the wire has made 62 % of its fall by the time the first pass that shows it posts, 30 s after it starts, so the sheet undershoots past its open-loop value to 73.2 g/m^2 91 s after the upset and the moisture, which nobody touched, falls to 3.59 %, four tenths of a point above the brittle limit.

3.6.6 Loss of Measurement and the First Reel Back

Scenario 10 parts the web at $t = 150$ s and threads it again five minutes later with both quality controllers in automatic throughout. During the break the gauge posts nothing, both controllers hold, and the cylinders keep condensing steam with nothing to evaporate into. The group 2 shell climbs from 138.2 to 144.5 °C, which is its saturation temperature, and groups 1 and 3 rise six degrees each, while five minutes of break puts 3.1 t of fiber in the broke pit. The first paper through after threading is the wet slug that was sitting in the dryer section, at 9.70 % moisture 36.5 s after the tail is sent across (9.70 to 10.00 % over four seeds), and 56.5 s later the same sheet is over-dried at 4.49 % against a 5.5 % target, because it is meeting cylinders six degrees hotter than they were. The control system finds out one scan period later and has acted a second one after that. No tuning shortens the sequence, because the whole of it happens inside two samples of the only instrument that can see it.

3.6.7 Grade Change and the Challenge

Scenario 9 moves three setpoints at once, from 80 to 64 g/m^2, from 5.5 to 6.0 % moisture and from 38.4 to 34.0 t/h. The moisture overshoots upward first, to 6.63 %, which is the wrong direction for a lighter sheet. Twenty percent less fiber is twenty percent less mass to weigh, so the conditioned weight of (3.20) collapses, WIC-401 cuts the stock valve from 55 to 45 %, and MIC-401 takes the group pressure from 306 down to about 170 kPa to stop the sheet going bone dry. The weight first reaches its band 257 s after the change and the moisture is in its band to stay 239 s after it, although the gauge noise moves the moisture time between 239 and 1037 s over six seeds. The production rate has not arrived at half an hour, reading 33.6 t/h against 34.0 t/h with the machine still winding up through 1103 m/min, because a cascade whose inner loop reports three times a minute has to be an order of magnitude slower again. Half an hour of grade change leaves 4.32 t of the 16.67 t made as neither grade, and five other seeds range from 3.05 to 5.08 t.

The graded challenge asks for 88 g/m^2 at 5.4 % on a machine whose press is half a point wet, whose hood is already humid, whose header is 50 kPa low and whose head is taking thirty seconds a pass, with a consistency dilution and a hood failure arriving later in the run. The answer has four parts and only one is tuning. The pairing has to be the natural one. The decoupler has to be on, because an eight gram weight move is a one and a half point moisture disturbance and there is no reason to make the moisture controller discover it through a gauge 32 s behind. The scan period has to come down, because at thirty seconds a pass (3.23) gives 31.9 s of effective dead time in the moisture loop and 40.7 s in the weight loop, against 11.9 and 20.7 s at ten seconds. Only then does the tuning follow. The constraint has to be checked before any of it, because once the dilution and the hood failure are in, the new grade needs the stock valve at 62.3 % and the group pressure at 450 kPa, which is a steam valve at 96 % of a 500 kPa header. Configured that way, which is the suggested line-up (decoupled, the tuning of Table 3.3, a 10 s scan), the last third of the run holds the sheet inside 0.35 g/m^2 and 0.21 points and the posted values the challenge scores inside 0.56 and 0.30 points, 44 and 33 % inside the limits of $\pm 1.0\,\mathrm{g/m^2}$ and ± 0.45 points that apply from 1188 s to the end of the 1800 s run. Left as it starts (two loops, slack tuning, a 30 s scan) the machine never reaches the grade at all, finishing seven grams light and 1.2 points wet with 17.2 t of off-specification paper against 8.5 t, and its worst posted deviations in the last third are 8.44 g/m^2 and 1.25 points. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt. Settings chosen before Run define the starting line-up and remain allowed, and the shipped tuning still fails from such a line-up, with a 2.99 g/m^2 maximum late weight error.

3.6.8 Model Predictive Control Against the Decoupled Loops

Scenario 14 is the grade change of scenario 9, 80 to 64 g/m^2, 5.5 to 6.0 % and 38.4 to 34.0 t/h at $t = 120$ s, with both quality handles given to the MPC of Section 3.5.5, and the challenge is now scenario 15. Table 3.7 runs it and eight other cases under four structures on each case's own seed, noise, events and duration, with the same instruments, no preview for any structure, the shipped structures at their shipped tuning (the challenge at its suggested line-up) and the MPC held at scenario 14's settings in every case. The third column is the like-for-like comparator, the decoupler retuned once on the grade change by the rule that chose the MPC's settings over a grid of 400 tunings ($K_c = 0.45$ and $\tau_I = 70$ s on weight, 13 and 60 s on moisture) and then held on every case.

On scenario 14 the MPC makes 2.68 t of off-specification paper against 4.32 t for the shipped decoupler and 6.53 t for two loops, while the retuned decoupler makes 1.51 t. On five further noise seeds it makes less than both shipped structures and more than the retuned decoupler every time, as listed in Table 3.7. The MPC moves at once and hard. At the posting that carries the change it takes the stock valve from 55.17 to 46.25 % and the pressure setpoint from 305.0 to 226.0 kPa, where the decoupler takes the stock only to 49.73 % with nearly the same pressure, and its weight enters the band 93.5 s after the change against 257.5 s.

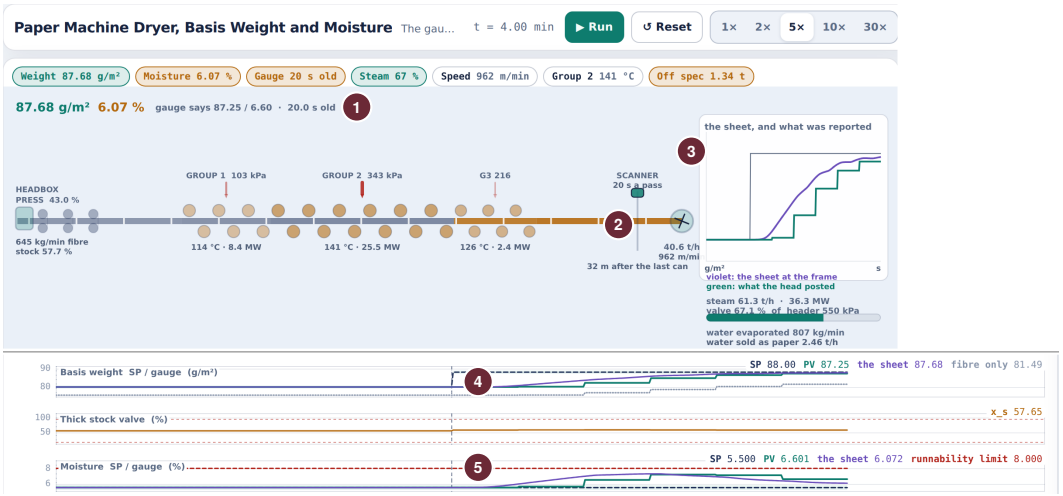

Figure 3.4 The machine at $t = 240$ s in scenario 6, two minutes after the basis weight setpoint moved from 80 to 88 g/m^2 with two independent quality loops in service and the moisture setpoint left at 5.5 %. The picture is a composite of two parts of the screen: the header and plant view above, and the first three trend strips below the gray rule. (1) The readout of the sheet beside what the gauge says. (2) The scanning head, drawn on its frame at the position reached in the current traverse. (3) The sheet at the frame against the staircase the head posted. (4) The basis weight strip, which carries the posted staircase the controller acts on, the sheet at the frame and the dry weight the steam valve cannot touch. (5) The moisture strip, with the excursion into a loop that was asked to do nothing, whose 1.80 points peak 62 s after the step Table 3.5 reports. The gauge still posts 6.601 % against a 5.500 % setpoint while the sheet at the frame is down to 6.072 %, and the posted 87.25 g/m^2 against an 87.68 g/m^2 sheet is the same lag on the weight strip.

Most of the gap to the retuned decoupler is the machine speed. SIC-401 reads the tonnage of (3.33), so the lighter grade first slows the machine to 955 m/min at 160 s and then drives it to 1054 m/min by 280 s, while the MPC's model holds 1000 m/min. By (3.3) a faster machine lays a lighter sheet. The MPC's sheet, at 63.30 g/m^2 and 6.21 % at 220 s when the decoupled loops' was still 71.56 g/m^2, falls to 59.99 g/m^2 with the moisture at 7.11 % by 280 s, as shown in Figure 3.5, and its moisture is in band to stay only 665.5 s after the change. With the gain of SIC-401 at zero, so that the speed stays at 1000 m/min, the MPC makes 1.95 t against 3.27 t for the retuned decoupler. The retuned loops gain from the production loop, and the MPC loses to it. The shipped decoupler is nearly indifferent to the fixed speed, at 4.37 t against 4.32, while the MPC falls from 2.68 to 1.95 t.

The tuning rule's noise condition holds on the tuning seed alone. There the MPC's 2σ of weight over the last ten minutes is 0.260 g/m^2 against 0.271 g/m^2 for the shipped decoupler, but on four of the five further seeds it is the larger. On the quiet machine of scenario 3 with the default gauge noise the MPC spreads the sheet over a 2σ of 0.305 g/m^2 against 0.214 and moves the stock valve 15.4 % against 7.9, where the earlier settings gave 0.211 g/m^2 and 6.25 %. The settings that suit the grade change pass more of the gauge noise into the paper.

Elsewhere the MPC makes the least off-specification paper of the four structures on the chest dilution (3.92 t against 5.59, 7.55 and 6.14) and on the header dip (0.51 against 1.93, 1.36 and 1.51), and on the challenge with its suggested line-up it makes 3.73 against 8.50 for the shipped decoupler and 7.11 for the retuned one, with 0.16 points of worst posted moisture against 0.30. On the clean moisture step it makes the least and is in band soonest at each of the scan periods 8, 10, 20 and 30 s. It loses three cases. The eight gram weight move costs it 3.33 t against 2.38 for two loops and 1.84

Table 3.7 Two Loops (TL), the Shipped Decoupler (Dec), the Decoupler Retuned Once by the MPC's Rule (Ret) and the MPC on Nine Cases. Each Case Runs on Its Own Seed, Noise and Events for 1800 s Unless Stated.

Case and measure	TL	Dec	Ret	MPC
Grade change, sc. 14, off spec (t)	6.53	4.32	1.51	2.68
seeds 1–5, least (t)	3.68	3.05	1.10	2.28
seeds 1–5, most (t)	5.08	5.08	1.33	2.39
both in band to stay (s)	1617.5	753.5	654.5	665.5
moisture in band to stay (s)	1617.5	239.5	113.5	665.5
2σ weight, last 10 min (g/m^2)	0.300	0.271	0.246	0.260
2σ moisture, last 10 min (points)	0.315	0.181	0.151	0.177
SIC-401 gain 0, off spec (t)	8.73	4.37	3.27	1.95
Weight move, sc. 6, to 900 s (t)	2.38	4.54	1.84	3.33
SIC-401 gain 0 (t)	5.65	3.36	2.26	1.13
Gauge noise, sc. 3, 2σ weight (g/m^2)	0.238	0.214	0.272	0.305
2σ moisture (points)	0.201	0.113	0.094	0.162
stock valve travel (%)	7.0	7.9	10.0	15.4
Moisture step, sc. 3, to 1400 s (t)	0.22	0.28	0.40	0.20
in band to stay (s)	23.5	29.0	40.0	21.5
Chest dilution, sc. 11 (t)	5.59	7.55	6.14	3.92
Header dip, sc. 12 (t)	1.93	1.36	1.51	0.51
Felt plug, sc. 13 (t)	4.81	7.40	5.76	5.04
Sheet break, sc. 10 (t)	1.15	1.35	1.75	1.62
Challenge, suggested (t)	—	8.50	7.11	3.73
worst posted moisture (points)	—	0.30	0.29	0.16

Off spec is the application's off-specification tonnage and band times count from the event on the true sheet at the frame. The noise case sets 0.25 g/m^2 and 0.12 points per pass before Run, with 2σ over 60 to 1800 s. The moisture step is 5.5 to 5.0 % at 200 s on a clean gauge at the 20 s scan. The suggested challenge line-up is decoupled, so two loops have no entry. The MPC decided at every posting, 90 times in an 1800 s run at the 20 s scan (74 in the break, 180 in the challenge at its 10 s scan), in 0.56 to 0.61 ms on average in Node.js, and rejected no plan.

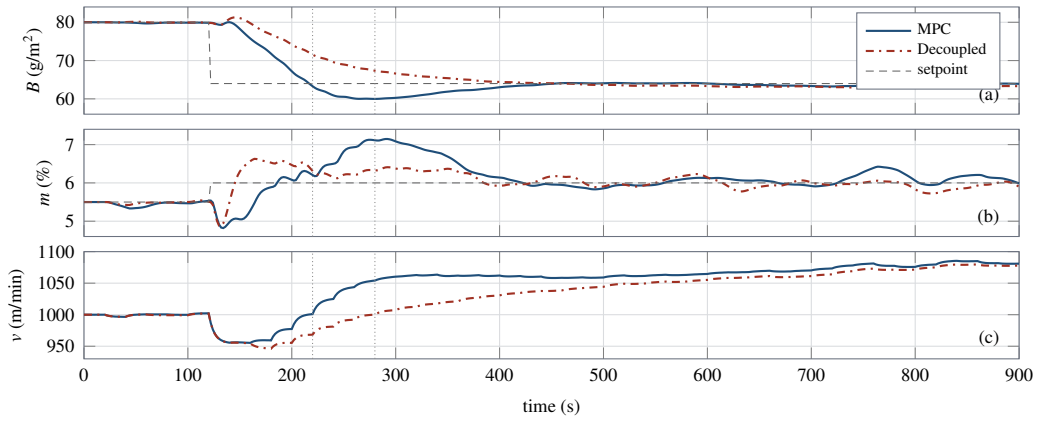

Figure 3.5 Scenario 14 under the MPC and under the decoupled loops of scenario 9, read on the true sheet at the frame. (a) Basis weight. (b) Moisture. (c) Machine speed set by SIC-401. The MPC is solid, the decoupled loops are dash-dotted and the setpoint is thin and dashed. The dotted lines mark 220 and 280 s.

for the retuned decoupler, and with the gain of SIC-401 at zero the same move costs it 1.13 t against 2.26, 3.36 and 5.65, so this loss is also the speed. The felt plug costs 5.04 t against 4.81 for two loops, and 4.29 against 4.94 with the speed held. The sheet break costs 1.62 t against 1.15 and 1.35, and

still 1.45 against 1.26 for the decoupler with the speed held. Its model has no break, and between the rethread at 450 s and 900 s it moves the stock valve over 51.3 to 56.9 % while the decoupled loops stay within 54.9 to 55.3 %. Switched to MPC on the challenge as loaded, the machine still fails, at 0.53 points of moisture against the 0.45 allowed, but with 5.40 t of off-specification paper against 17.22 for the shipped loops.

Ablations locate what the model contributes. Reading each posting as the sheet now, with the scanner and the transports left out, costs the most on the grade change, the weight move, the dilution and the felt plug (3.37, 4.36, 5.71 and 6.36 t against 2.68, 3.33, 3.92 and 5.04). Replacing the steam observer by the design header costs the most on the header dip (6.49 against 0.51 t) and the challenge (10.58 against 3.73). Two single-loop plans in one program, each handle seeing only its own quality, make 2.43 t on the grade change, less than the full plan and less on four of the five further seeds, and 3.08 on the weight move, so the interaction in the model does not pay on the two setpoint moves, although it does on the dilution, the felt plug, the header dip and the challenge (4.84, 5.32, 0.59 and 4.38 t without it). The steam valve ceiling never binds in a shipped scenario, and every scenario and the suggested challenge run identically with and without it apart from the solver's iteration count. It binds only when the header of scenario 12 is lowered to 320 kPa, where it is worth 3.88 t against 5.64 without it and 8.56 for the shipped decoupler.

3.7 Technology in Practice

What Measures It

Basis weight is read by attenuation of beta radiation. The first survey of nucleonic gauges by the International Atomic Energy Agency, in 1962 and 1963, counted around 20 000 gauges in 21 countries, and in Finland 72 % of them were used in the paper industry [19]. On the first computer-controlled paper machine the beta-ray gauge sat at a fixed position at the dry end and moisture was read by a capacitance gauge that could traverse the web but was normally held still. The machine tender and the test laboratories also measured basis weight, and the weight and size of each reel, a very accurate average, calibrated the other gauges [8]. Full-sheet web scanning systems now read moisture, thickness, ash and many other properties of the web [20]. The chapter's frame QT-401 carries a beta gauge and an infrared moisture head, averages one traverse and holds it, with per-pass repeatabilities of 0.25 g/m^2 and 0.12 %.

> **TECHNOLOGY Sensor: The Beta Transmission Gauge**
>
> A transmission gauge places the material between a radioactive source and a detector, and the material attenuates the beta particles before they reach the detector. Krypton-85, strontium-90, promethium-147 and thallium-204 are the sources listed for gauging paper [19]. Beta rays are absorbed about equally by fiber and by water, so the reading is proportional to the mass of both per unit area, the wet basis weight [8]. A drifting gauge leaves no mark on a trend, because the loop holds the drifted reading on its setpoint while the reel moves off grade. The first computer-controlled machine therefore used the weight and size of each finished reel to compensate the drift of its beta-ray gauge [8]. The chapter's gauge is perfect apart from a repeatability of 0.25 g/m^2 per pass.

What Moves It

On the first computer-controlled machine the computer moved basis weight through the setpoints of analog regulators on the thick stock flow valve and the thick stock dilution valve [8]. The chapter keeps the first of these as NCV-101, a valve with a 3 s travel feeding an approach system that adds a 22 s mixing lag. The steam side is a valve on the main dryer group, and the machine drive is the third handle, moved only by the production loop.

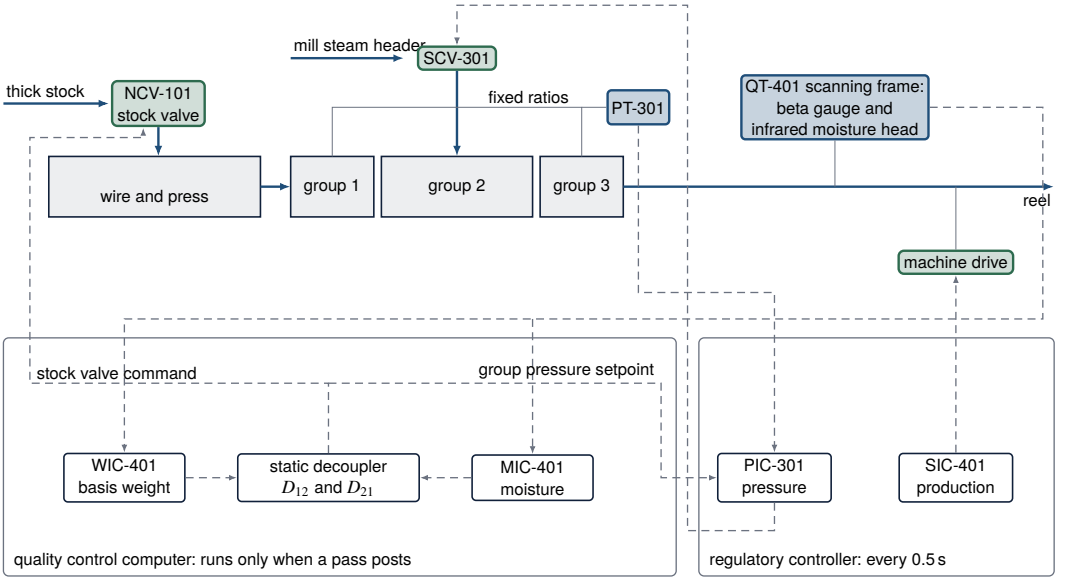

Figure 3.6 The quality loops on plant hardware. The scanning frame carries the beta gauge and the infrared moisture head across the sheet and posts one average of each per pass. Only then does the quality computer run WIC-401 and MIC-401, whose outputs pass through the static decoupler to the stock valve and to the setpoint of PIC-301, which holds the main group pressure through SCV-301 every half second. SIC-401 moves the machine drive. Notice that the steam valve is the only final element under a loop with a continuous measurement.

TECHNOLOGY **Final Element: Steam to the Dryer Groups**

A multi-cylinder dryer section carries 40 to 90 steam-heated cylinders, and groups of them run at different steam pressures. The first cylinders the sheet meets run at the lowest pressure and temperature and later groups run hotter, which keeps the paper from sticking to the shells. The drying rate depends on the temperature and volume of the steam entering the cylinders and on how well condensate is removed from inside them, as well as on the felts, the hood air and the pocket ventilation [20]. A group that floods therefore dries less at the same pressure, and the moisture loop sees only a wetter sheet. In the chapter one valve, SCV-301, sets the main group pressure with a 4 s travel, the outer groups follow at ratios of 0.30 and 0.63, and no condensate system is modeled.

Where the Controller Runs

Industrial control loops are often nested or cascaded, so that one loop's setpoint comes from another loop, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [21]. A distributed control system adds a supervisory level that oversees the controllers of the local processes [21]. The chapter's structure sits on two such levels with different clocks. PIC-301 and SIC-401 execute every 0.5 s as ordinary PID blocks. WIC-401 and MIC-401 execute only when the head posts a pass, every 20 s as shipped, and the decoupler between them is two gains and two sums. The optional predictive controller decides once per posting as well, in 0.56 to 0.61 ms on average. Figure 3.6 draws the loops on their devices.

IN PRACTICE **Steam per Kilogram of Water**

Energy use in papermaking is set most by the operation of the steam and condensate system and of the air handling around the dryers. A well-maintained drying system averages about 1.3 kg of steam per kilogram of water evaporated, and many run with higher steam use [20]. The same check can be made on the simulated machine, which at its design line-up consumes 58.2 t/h of live steam while it removes 766 kg/min of water.

What Is Modernizing It

A 2025 assessment of industrial drying notes that conventional dryers often dry unevenly, leaving overdried regions that raise energy use and degrade paper quality, while microwave drying targets wet spots. It reports machine learning applied to drying that reduced steam use by 1.5 %, and it lists the life of existing dryers, 30 to 60 years, as a barrier that makes their replacement slow [13]. Sensors that read properties beyond weight and moisture on the moving web are a second direction, taken up in the open challenges below. For the chapter's structure, a dryer section that keeps its steam-heated cylinders for decades means that the pairing and the decoupler of Section 3.4 will run on the same plant for a long time, and are worth getting right.

3.8 Limitations

The model cannot be used to size a dryer section. Three lumped groups, a fitted conductance, no condensate system, no sheet temperature state and no pocket ventilation balance put it well short of what a thermal design calculation requires [5, 6], and no number in it should be read as a prediction of a real machine's dryness profile. The cross direction is absent altogether, which is the largest omission, because the traversing head serves profile control as much as machine-direction control and the trade in Table 3.4 therefore has only one of its two sides measured here [4]. The gauge is perfect apart from its per-pass noise, so standardization cycles, sheet flutter, cross-sensitivity between the two measurements and slow drift are all outside the model, and a drifting gauge is a more common cause of an off-specification reel than any upset in Table 3.6.

OPEN CHALLENGE **Controlling the Property the Customer Buys**

Stiffness is one of the most important properties of paper, yet it is measured offline on samples cut from the finished reel, manually and slowly, too late to adjust the machine. Mills therefore use more pulp than necessary so that the product meets its stiffness standard, and as much as 5 % of the paper made each year is judged substandard and remanufactured. A laser ultrasonic sensor that reads stiffness on the moving web was at the demonstration stage in a 2012 review, which noted that a mill could then run to stiffness targets rather than basis weight targets [22]. Here basis weight stands in for what the customer buys. Progress would be an online property measurement trusted as a setpoint.

Dry weight control is deliberately not offered. Mills also control $B(1-m)$ from the same two readings, and that pairing interacts much less than the conditioned pairing of Equation (3.20), so working out why it helps is worth more than being handed it. Apart from the predictive controller, the four loops implemented are the only ones available, and consistency ratio control, dry line control, a steam box, a condensate differential loop and cross-direction actuators are absent even though each changes the machine-direction problem. The disturbances arrive as steps and timed events instead of the slow wandering a real plant does, and consistency, retention and hood humidity are open sliders whose real dynamics recirculate through the white water and broke systems over tens of minutes. Section 3.3.6 names the point at which each simplification stops holding.

The predictive controller's model omits whatever moves with the operating point. It is linearized once at the design grade, where the moisture gain to the group pressure is −0.0607 points per kPa, while the plant's is −0.0905 near the new grade of scenario 14 (stock 46 %, 190 kPa) and −0.0452 at the challenge's line-up. It fixes the speed at 1000 m/min and the transports at their design values, although at the 1103 m/min that ends scenario 14 they are 9.68 and 1.74 s, and it has no break. Its

Kalman filter assigns every such error to the consistency and the press dryness, so with neither changed it ends scenario 14 reading 3.21 % and 39.90 %. Its settings come from one case and one noise sequence, its bands are the declared tolerances whatever the run's are, and each soft limit has a single slack for the whole horizon.

> ⚠ **OPEN CHALLENGE** **Control That Keeps Performing Through Grade Changes**
>
> An industry research plan of 1999 counted among its priorities techniques that sustain the performance of control systems without deterioration over time. Its pathway for control system effectiveness listed multivariable and nonlinear control and control systems that automate start-up and transitions, with product loss during transitions minimized as the result sought [23]. In this chapter half an hour of grade change still leaves 4.32 t as neither grade, and the predictive controller, linearized once at the design grade, loses to retuned loops. Progress would be a multivariable controller whose model follows the grade and the speed.

3.9 Discussion

Moisture is sold at the price of paper. At the design line-up the machine ships 38.40 t/h of which 2.11 t/h is water, so one reel in eighteen is water that the customer pays for at $820 per tonne. At a fixed fiber rate every tenth of a point of moisture setpoint is worth about 41 kg/h more of it, which is roughly 33 $/h, and the economic case for a tighter moisture loop is that the setpoint can then move closer to the specification limit without crossing it. The variability the loop cannot remove sets that distance, and Table 3.4 prices it, because the same controller holding 4.49 points of swing at thirty seconds a pass and 0.18 points at ten seconds needs entirely different margins.

The other half of the economics runs the other way. Water removed in the press is much cheaper than water removed in the dryers, and the exchange rate depends on what is being held. Leaving the steam valve where it is, two points of press dryness add 2.0 MW of dryer duty while the sheet drifts to 8.9 % moisture, which is 1.0 MW per point of dryness. Holding the reel moisture on 5.5 % instead costs 3.1 MW for the same two points, which is 1.5 MW per point and 40 kW for every tonne an hour of production, because the extra water now has to be boiled off rather than shipped. The fixed-valve calculation understates the cost by half, and it is the calculation a felt-conditioning budget is usually argued from. Dryer energy is the largest single steam demand in a paper mill and one of the cheapest places to save it [12].

Nothing restricts the plant of this chapter to the four loops built on it. The fourteen balances of Section 3.3 form a complete nonlinear state-space model with hard actuator limits, and the controller of Section 3.5.5 uses their linearization to coordinate the two qualities as one two-input problem. Its steam valve ceiling never binds in the challenge, where the plan predicts at most 96.1 % and the valve reaches 96.3 %, nor in any other shipped scenario (Section 3.6.8). A sampled, averaged and held measurement is also where state estimation earns its keep, because a moving horizon estimator running on Equations (3.2) to (3.10) reconstructs the current sheet between passes and recovers part of the phase that (3.23) removes, which is the reasoning that puts a soft sensor in front of a slow laboratory measurement [1]. Volume 5, Chapter 10 builds such a soft sensor for powder moisture and corrects its bias from a sampled, delayed laboratory result. The unmeasured consistency of scenario 11 and the unmeasured press dryness of scenario 13 are both estimable from the two quality readings and the steam demand, and the predictive controller's filter estimates both, although it confuses them with its own linearization error (Section 3.8). Rebuilding the balances in the Gekko Optimization Suite under Python takes the eight hidden constants as the quantities to be estimated [24].

3.10 Conclusion

This chapter studied a fine-paper machine as a laboratory for control through an instrument that is late, intermittent and spatially incomplete at once. The simulator integrates fourteen states, with the scanning gauge represented as a boxcar average over one traverse, a zero-order hold between passes and a transport delay computed from the live speed. The effective dead time is the transport delay plus

one whole scan period, 21.9 s for moisture and 30.7 s for basis weight, and a loop tuned against the transport delay alone rings at 100 s with a growing swing, 4.5 times the dead time it faces. Shortening the traverse from 20 to 10 s at an unchanged gain reduces that swing from 8.72 to 0.33 points and halves the sheet each reading averages over, from 333 to 167 m.

Basis weight sold conditioned makes the plant a two-by-two with a relative gain of 1.52 on the natural pairing and −0.52 on the other, and the swapped pairing left in service for twenty-two minutes after a consistency dilution reaches 72.8 g/m^2 and 13.12 % while still walking. A static decoupler cuts the moisture excursion of an eight gram grade move from 1.80 to 0.23 points, while the weight, held just above its band as the production loop slows the machine, takes the off-specification tonnage from 2.38 to 4.54 t.

An optional predictive controller with the gauge, the transports and PIC-301 in its model makes 2.68 t of off-specification paper on a grade change against 4.32 t for the shipped decoupler, and it loses to the decoupler retuned by the same rule (1.51 t) chiefly because its model holds the speed that the production loop moves.

A fuller model would add a cross-direction profile with slice actuators, a condensate system with a differential loop per group, dry weight as a selectable controlled variable, and a consistency transmitter so that feedforward on the disturbance can be compared against feedback on the quality. As it stands, the plant's limiting element is the instrument, and the reader learns it by operating the machine.

KEY POINTS

- The gauge is late by the time the sheet takes to reach it plus one pass of the head, and for moisture the gauge owns most of that delay.
- Shortening the pass from 20 to 10 s at an unchanged gain cuts the scenario 4 swing from 8.72 to 0.33 points.
- Conditioned basis weight couples the plant both ways: the relative gain is 1.52 on the natural pairing and −0.52 on the swapped one, which cannot be tuned.
- A static decoupler cuts the moisture excursion of an eight gram grade move from 1.80 to 0.23 points, almost all of it from the stock-to-steam half, while the weight loop pays.

Where else this applies. Chapter 2 runs one half of the same decoupler on a triangular plant, and Volume 5, Chapter 3 meets the same dead-time accounting with an online analyzer.

References

[1] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

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

[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] J. G. VanAntwerp, A. P. Featherstone, R. D. Braatz, and B. A. Ogunnaike, "Cross-directional control of sheet and film processes," *Automatica*, vol. 43, no. 2, pp. 191–211, 2007.

[5] B. Wilhelmsson and S. Stenstrom, "Heat and mass transfer coefficients in computer simulation of paper drying," *Drying Technology*, vol. 13, no. 4, pp. 959–975, 1995.

[6] M. Karlsson and S. Stenström, "Static and dynamic modeling of cardboard drying part 1: Theoretical model," *Drying Technology*, vol. 23, no. 1-2, pp. 143–163, 2005.

[7] M. Karlsson, O. Slätteke, B. Wittenmark, and S. Stenström, "Reducing moisture transients in the paper-machine drying section with the mid-ranging control technique," *Nordic Pulp & Paper Research Journal*, vol. 20, no. 2, pp. 150–155, 2005.

[8] K. J. Åström, "Computer control of a paper machine—an application of linear stochastic control theory," *IBM Journal of Research and Development*, vol. 11, no. 4, pp. 389–405, 1967.

[9] G. A. Dumont, “Application of advanced control methods in the pulp and paper industry—a survey,” *Automatica*, vol. 22, no. 2, pp. 143–153, 1986.

[10] T. J. Harris, “Assessment of control loop performance,” *The Canadian Journal of Chemical Engineering*, vol. 67, no. 5, pp. 856–861, 1989.

[11] W. L. Bialkowski, “Control of the pulp and paper making process,” in *Control System Applications*. Boca Raton, FL: CRC Press, 2018, pp. 43–66.

[12] L. Kong, A. Hasanbeigi, and L. Price, “Assessment of emerging energy-efficiency technologies for the pulp and paper industry: a technical review,” *Journal of Cleaner Production*, vol. 122, pp. 5–28, 2016.

[13] R. P. P. da Silva, I. J. Okeke, K. Thirumaran, M. Mosalpuri, F. Samadi, S. U. Nimbalkar, and J. Cresko, “Current state and future projections of drying processes in the US food and pulp and paper sectors: Energy, economic, and environmental assessment,” *Energy Conversion and Management: X*, vol. 27, p. 101181, 2025.

[14] A. P. Featherstone and R. D. Braatz, “Control-oriented modeling of sheet and film processes,” *AIChE Journal*, vol. 43, no. 8, pp. 1989–2001, 1997.

[15] E. Gagnon, A. Pomerleau, and A. Desbiens, “Simplified, ideal or inverted decoupling?” *ISA Transactions*, vol. 37, no. 4, pp. 265–276, 1998.

[16] 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.

[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] K. R. Muske and T. A. Badgwell, “Disturbance modeling for offset-free linear model predictive control,” *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

[19] International Atomic Energy Agency, “Technical data on nucleonic gauges,” International Atomic Energy Agency, Vienna, IAEA-TECDOC 1459, Jul. 2005. [Online]. Available: https://www-pub.iaea.org/MTCD/Publications/PDF/te_1459_web.pdf

[20] M. Miller, M. Justiniano, and S. McQueen, “Energy and environmental profile of the U.S. pulp and paper industry,” Energetics, Incorporated, for the U.S. Department of Energy, Industrial Technologies Program, Columbia, MD, Tech. Rep., Dec. 2005.

[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] L. Kong, A. Hasanbeigi, and L. Price, “Emerging energy-efficiency and greenhouse gas mitigation technologies for the pulp and paper industry,” Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep., Nov. 2012.

[23] American Forest & Paper Association, “Agenda 2020: The path forward, an implementation plan,” American Forest & Paper Association, Washington, DC, Tech. Rep., 1999.

[24] 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.

3.11 Exercises

3.1 ○ Concept check · 2 min

At the design line-up the moisture loop faces an effective dead time of 21.9 s, although the sheet takes only 1.92 s to travel from the last cylinder to the frame. What makes up the other 20 s?

(a) The can bank, whose slowest shell has an 84 s time constant

(b) The 0.5 s regulatory scan, added once for each of the forty scans of a pass

(c) The approach system and the forming lags between the stock valve and the wire

(d) Half a scan period for the boxcar average and half a period for the hold

3.2 ○ Concept check · 1 min

In scenario 1, with the machine in Manual, raising the steam valve by 6 % makes the sheet 1.48 g/m^2 lighter, although not one gram of fiber changes. Why?

(a) The beta gauge weighs the water in the sheet, and more steam evaporates more of it
(b) Hotter cylinders shrink the sheet, so less fiber lies on each square meter
(c) The production loop speeds the machine up as the group pressure rises
(d) The press leaves the sheet wetter when the dryer cylinders run hotter

3.3 Concept check · 2 min

Each loop of the swapped pairing, basis weight on the group pressure and moisture on the stock valve, gives a clean first-order plus dead-time fit and a sensible tuning when it is step tested with the other loop in Manual. Why does the pair fail when both are put in service in scenario 7?

(a) The gain of 19.5 kPa per g/m^2 needs more pressure than the group can take
(b) Closing either loop reverses the sign of the gain the other acts on, as $\Lambda_{12} = -0.52$ says
(c) The gauge posts the two qualities at different times, so the loops drift out of step
(d) The swapped loops have no back-calculation, so both integrators wind up together

3.4 Theory and calculation · 6 min

The scanning frame stands $L_B = 178\,\text{m}$ from the slice and $L_s = 32\,\text{m}$ from the last cylinder. Use Equations (3.21) and (3.23).

(a) At 1000 m/min and a 20 s scan, find the two transport delays and the two effective dead times. *Checkpoint:* 10.68 s and 1.92 s; 30.68 s for weight and 21.92 s for moisture.

(b) Repeat the effective dead times for scan periods of 30 and 10 s, and find the length of sheet one traverse averages over at each. *Checkpoint:* 40.68 s and 31.92 s with 500 m; 20.68 s and 11.92 s with 167 m.

(c) The operator slows the machine to 800 m/min at the 20 s scan. Find the transport delays and effective dead times now, and say whether slowing down helps the quality loops. *Checkpoint:* 13.35 s and 2.40 s; 33.35 s and 22.40 s.

3.5 Theory and calculation · 7 min

Use the gain matrix of Equation (3.24), $K_{11} = 2.207$, $K_{12} = -0.0514$, $K_{21} = 0.889$ and $K_{22} = -0.0607$, with the units given there.

(a) Compute the relative gains Λ_{11} and Λ_{12}. *Checkpoint:* 1.52 and -0.52.

(b) Compute the Niederlinski index of Equation (1.7) for the natural pairing and for the swapped pairing. *Checkpoint:* 0.66 and -1.93.

(c) Compute the decoupler gains of Equation (3.27). With the shipped $D_{21} = 14.5$, find the steady moisture change per percent of weight controller output, and repeat for $D_{21} = 29.25$. *Checkpoint:* $D_{12} = 0.0233$, $D_{21} = 14.6$; $+0.009$ and -0.887 points per percent.

3.6 Theory and calculation · 5 min

At the design line-up the machine runs at 1000 m/min on an 8.0 m trim, making 80.00 g/m^2 at 5.50 % moisture. Use Equations (3.7) and (3.20).

(a) Find the dry weight, the saleable production and the water and fiber in it, in t/h. *Checkpoint:* 75.60 g/m^2; 38.40 t/h, of which 2.11 t/h is water and 36.29 t/h fiber.

(b) In scenario 1 a 2 % stock step adds 2.75 g/m^2 of fiber and makes the sheet 1.86 points wetter. Find the new conditioned weight. *Checkpoint:* About 84.57 g/m^2, against the 84.58 g/m^2 measured.

(c) At a fixed fiber rate, find the extra water sold for each tenth of a point of moisture setpoint, and its value at the paper price of $820 per tonne. *Checkpoint:* About 41 kg/h, worth about $33 an hour.

3.7 App challenge · 8 min

Load scenario 2, The gauge: step the stock and time the answer. Run it to 300 s at the default seed and a speed of 10. The stock valve steps from 55 to 57 % at 120 s. Read the posted basis weight (`weight_meas_gsm`) and the sheet at the frame (`weight_true_gsm`) from the trend or the exported file.

(a) At what time does the posted basis weight first change after the step, and to what value? *Checkpoint:* At 140 s, to 80.14 g/m^2.

(b) The settled move is 4.58 g/m^2. Find the times after the step at which the sheet and the posted value first pass half of it, 82.29 g/m^2. *Checkpoint:* About 36.5 s and 60 s.

(c) What does the posting made at 160 s report, and what is the sheet at the frame at that instant? In the exported file a posting first appears in the row half a second after the pass ends. *Checkpoint:* 81.73 g/m^2 against 82.66 g/m^2.

Accept weights within 0.02 g/m^2 (0.05 g/m^2 for the sheet in (c)) and times within 1 s. Account for the gap in (b) with Equation (3.23).

3.8 App challenge · 8 min

Load scenario 4, Tuned as if the gauge were continuous, and use the default seed and a speed of 30. The moisture loop carries 44 kPa per point and 60 s, and the moisture setpoint moves from 5.5 to 5.0 % at 120 s. Read the sheet moisture (`moist_true_pct`) from the exported file.

(a) Run to 1400 s. Find the lowest and highest moisture between 700 and 1400 s, and the period of the oscillation. *Checkpoint:* 1.18 and 9.90 %, 8.72 points peak to peak, a period of about 100 s.

(b) Reload the scenario, set Scanner period to 10 s before Run, and run to 1400 s again. Find the same two values. *Checkpoint:* 4.83 and 5.16 %, 0.33 points peak to peak.

(c) Divide the period of (a) by the effective dead time of the moisture loop, and explain the result of (b) without a change of gain. *Checkpoint:* 4.5; the dead time falls from 21.9 to 11.9 s.

Accept moistures within 0.02 points and the period within 3 s.

3.9 App challenge · 8 min

Use the default seed and a speed of 30, and run each case to 900 s. The basis weight setpoint moves from 80 to 88 g/m^2 at 120 s with the moisture setpoint left at 5.5 %. Read the highest sheet moisture after the step (`moist_true_pct`) and the Off spec tile (or `metric_offspec_t`) at 900 s.

(a) Load scenario 6, The interaction, undeclared, which runs two independent loops. *Checkpoint:* 7.30 % and 2.38 t.

(b) Load scenario 8, A decoupler, which runs the full decoupler. *Checkpoint:* 5.73 % and 4.54 t.

(c) Reload scenario 8, switch to Advanced, set Steam-to-stock decoupling (D_{12}) to 0 before Run, and run again. *Checkpoint:* 5.74 % and 3.14 t.

Accept moistures within 0.02 points and tonnages within 0.02 t. Which half of the decoupler carries the benefit to the moisture, and why does the off-specification tonnage rise?

3.10 Application area · 4 min

A plugged press felt drops the press dryness by two points. With the steam valve held, it adds 2.0 MW of dryer duty while the reel moisture drifts to 8.9 %. With the reel moisture held on 5.5 %, it adds 3.1 MW. Use Equation (3.15) with $h_{fg} = 2135\,\mathrm{kJ/kg}$ and the steam price of Table 3.2.

(a) Find the extra steam in each case, in kg/min and t/h. *Checkpoint:* 56.2 kg/min (3.37 t/h) and 87.1 kg/min (5.23 t/h).

(b) Find the hourly steam cost of holding the moisture, and say why holding it costs half as much again as leaving the valve alone. *Checkpoint:* About \$115 an hour.

3.11 Application area · 4 min

A sheet break stops the gauge from posting for five minutes, and the quality controllers hold their last outputs. The moisture loop runs at $K_c = 18\,\text{kPa}$ per point and $\tau_I = 90\,\text{s}$, with the integration interval of Equation (3.32).

(a) The first reading after threading shows a moisture error of one point. Find the integral step with an unclamped interval of 300 s and with the clamp at the 20 s scan. *Checkpoint:* 60 kPa and 8 kPa.

(b) The integral time of both quality loops cannot go below 60 s. Find the largest ratio $\Delta t_k / T_t$ the clamp allows at the longest scan, 40 s, and say why it matters. *Checkpoint:* 1.33, below 2.

(c) Name the two other protections the production loop needs around a break, and what each prevents. *Checkpoint:* A parked setpoint and a reclose at 95 % of a threaded web.

3.12 Review the agent · 5 min

An AI assistant reviews scenario 14 and writes: “The MPC makes 2.68 t of off-specification paper on the grade change, against 1.51 t for the decoupled loops retuned by the same rule. Its model of the scanner and the interaction is therefore not worth having. Raise its move suppressions until it behaves like the decoupler, or replace it with the retuned loops.” Find the flaw in the reasoning and say how you would check it on the app.

CHAPTER 4

Vapor-Chain Concentration Control in a Three-Effect Milk Evaporator Simulator

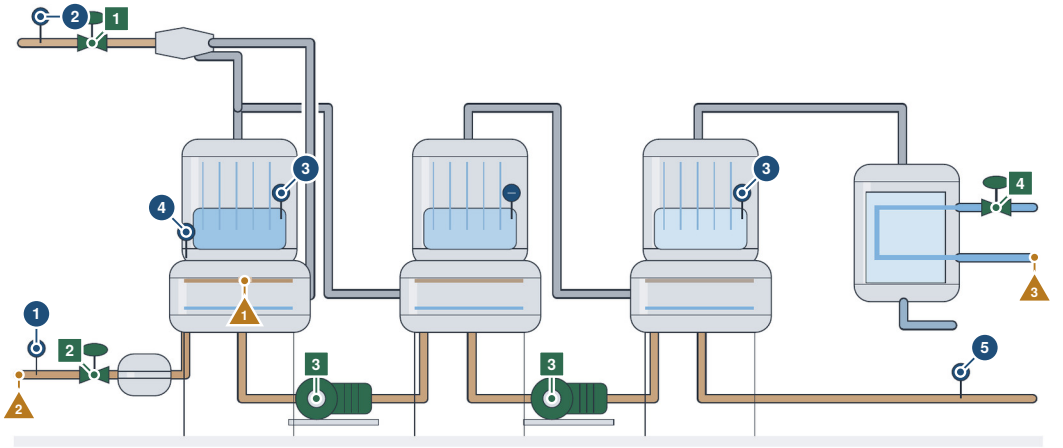

The controller moves 1 motive steam to the steam jet 2 milk fed to the train 3 pumps between the effects 4 cooling water to the condenser

The controller reads 1 feed flow and feed solids 2 motive steam flow 3 sump levels 4 first-effect temperature 5 densitometer at the end of the product line

What disturbs it 1 deposits on the tubes 2 feed solids 3 warmer cooling water

The vapor of each effect heats the next within two minutes, while the concentrate reaches the only product reading at the end of the line minutes later. Illustration, not to scale.

IN PLAIN TERMS Steam Used Three Times, Judged at the End

The main idea. An evaporator thickens skim milk by boiling water out of it under a vacuum. Steam heats the first of three boiling stages, and the vapor from each stage heats the next one.

The challenge. The only gauge of the thick milk sits at the end of a long pipe, so it sees a steam change minutes late. The thickness can be fixed with the steam or with the milk flow, while deposits slowly push the first stage toward its heat limit.

Why it matters. The handle chosen decides whether the plant pays in steam or in milk, and how long it runs before it stops for cleaning.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on steam per pound of water in the Introduction, the history of the sugar pan that held its own concentration (Section 4.2), the two try-its and the fact on level tuning (Section 4.6) and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the two open challenges

The control question A densitometer at the end of a three-effect train reads the concentrate minutes after a steam move has reached every effect. Should the product solids be held with the steam or with the feed, how tightly should the sumps between the effects be held, and what ends a run when fouling spends the temperature margin of the first effect?

Learning objectives

- Explain why a steam step reaches the three vapors within two minutes and the densitometer only after about six.
- Compute the steam a feed-solids change needs from the mass balance, and the share a ratio built on the average economy leaves to the trim.
- Use the temperature budget to predict how the condenser and the deposits of each effect move the first effect's temperature.
- Compare steam and feed pairing at equal loop gain, and state what each costs in steam or in milk processed.
- Choose between tight and averaging level loops from the flow variation they pass to the dryer.

Prerequisites Sections 1.4.3 (throughput manipulator), 1.4.8 (averaging level control), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive supervisor.

The app https://apmonitor.com/apps/multi-effect-evaporator/

See also Volume 4, Chapter 4 for a feed rate cut by engine and rotor overrides when a limit binds, and Volume 4, Chapter 5 for torque and pressure overrides beneath a throughput supervisor on a product judged downstream.

4.1 Introduction

Milk powder is made by removing water twice. A falling-film evaporator first concentrates skim milk by boiling it under vacuum, and a spray dryer then removes most of what remains, so the evaporator decides both how much steam it spends and what concentrate the dryer receives. Energy use and energy efficiency have been analyzed along the whole route from fluid milk to milk powder [1]. The evaporator saves energy by reusing it. Its effects boil at successively lower pressures, the vapor boiled from the milk in one effect heats the next, and a thermal vapor recompressor (TVR), a steam jet, returns part of the first effect's vapor to that effect's own steam chest. Falling-film trains of this kind have been modeled from first principles and identified from plant data [2, 3], and their dynamics and control have been analyzed with mechanical vapor recompression [4].

In the train modeled here, 20 t/h of skim milk at 9.0 % solids leave the third effect as 3.75 t/h of concentrate at 48.0 %. The 3.39 t/h of motive steam that drive the thermocompressor entrain 5.30 t/h of the first effect's vapor, and with 0.18 t/h of steam to the preheater they evaporate 16.25 t/h of water, a steam economy of 4.56 kg/kg. Heat and concentration travel through the train at different speeds. After a 0.10 t/h step of the motive steam, the vapor made in the three effects shows half of its change after 16, 75 and 82 s, while the densitometer at the end of the product line needs 380 s.

The simulator is built against two beliefs. The first is that an evaporator's steam economy is fixed by its design, so that every extra kilogram of steam evaporates as much water as the average one. Here the average kilogram of motive steam evaporates 4.80 kg of water and the next one 2.4 kg, because a hotter steam chest makes the thermocompressor entrain less. The second is that the temperature of the first effect, which must stay below a product limit, is a matter of the first effect. It rests on the condenser and on the temperature drops of the effects below it, so deposits in the second and third effects spend its margin while deposits in the first spend economy. A transfer function from steam to densitometer carries neither. The class of problem is general. A quality measured at the end of a chain of thermally coupled units, a fast handle that costs energy against a slow one that costs throughput, buffers whose level tuning decides how flows propagate, and a hard limit approached by slow degradation are shared by the sugar evaporators where multiple effects began [5], by black

liquor trains [6], by multiple-effect desalination [7], and by any plant that must choose where its throughput is set [8] and whether its inventories are held tight or allowed to average [9].

> ✦ DID YOU KNOW? **Steam per Pound of Water, Effect by Effect**
>
> To evaporate one pound of water, a dairy evaporator with one effect needs 1.2 pounds of steam, one with two effects 0.6, one with four 0.3 and one with seven less than 0.1. Trains with thermal vapor recompression have a typical steam economy of 4 to 8 [10]. The train modeled here runs at 4.56 kg/kg.

Evaporator studies model a particular plant in order to identify or control it [2, 3, 11], and a first course in process control meets feedforward, level control and override control one loop at a time [12]. This chapter puts them on one train, on which the vapor chain can be watched while fouling spends the temperature margin. The effect balances, the thermocompressor energy balance, ratio feedforward, averaging level control, override control and the throughput supervisor are standard practice. Two measured results stand out. The marginal economy of the motive steam is half of its average, and at the deposits of a twenty-hour run the first effect's own deposit costs 1.19 K of temperature margin and 0.38 kg/kg of economy where those of the other two cost 4.20 K and 0.16 kg/kg. The model has 28 plant states and 26 ledger states whose water, solids and energy balances close to round-off, with a densitometer delay keyed to the product volume and a steady state solved by Newton's method, all written out here. The app carries eight cases and one graded shift, which set the vapor path against the liquid path, a steam pairing against a feed pairing and a buffer against a tight level loop, and a further case adds predictive supervision. Section 4.2 follows the domain's history, Section 4.3 builds the model and Section 4.4 reads its vapor chain and temperature budget. Section 4.5 describes the controller, Section 4.6 reports what the cases measure, with the predictive supervisor in Section 4.6.8, Section 4.7 describes the real instruments, final elements and controllers, and Section 4.8 states the limitations.

4.2 From the Sugar House to the Dairy Train

Evaporation in several effects entered industry through the sugar house. Norbert Rillieux's United States patents of 1843 and 1846, for improvements in sugar-works and in sugar-making, placed vacuum pans in series so that the vapor boiled from the juice in one heated the next, and both are reproduced in a documentary history of technology [5, 13]. The principle outgrew sugar. Seawater desalination by multiple-effect evaporation is analyzed effect by effect for its thermal performance [7], and the black liquor of the pulp mill is concentrated in falling-film effects whose plates have been modeled as distributed-parameter systems [6].

Two later developments shaped the dairy train. The first is vapor recompression. A steam jet compresses part of the vapor a train makes and returns it to the first effect, and an investigation of ejector design by analysis and experiment was published in 1950 [14]. The second is the falling film, which transfers heat at small temperature differences with a short residence, at the price of a film that has to cover the whole tube wall. Criteria for the breakup of thin liquid layers flowing over solid surfaces were set out in 1964 [15], and in falling-film evaporators a broken film has been identified as a cause of fouling [16].

Milk brought constraints of its own, a product that must not be overheated and deposits that grow on every heated surface. The fouling of heat exchangers in the dairy industry has been reviewed [17], and a dairy evaporator runs from one cleaning to the next rather than indefinitely. Control of these trains moved from single loops toward models. Process identification and physical modeling were compared for falling-film evaporators in 1994 [2], and the same authors placed a four-effect falling-film evaporator under multivariable supervisory control [11]. The dynamics and control of falling-film evaporators with mechanical vapor recompression were analyzed in 1999 [4], and dynamic models of an evaporation process were compared in 2000 [3]. Two structural decisions recur there and in plantwide control at large, which variable sets the throughput [8] and whether the inventories

Table 4.1 Variables of the Evaporator, Grouped by Role, at the Reset of Case 2

Variable	Role	Design value	Range
w_{AT}	CV, AT-301	48.0 % SP	40–52
F	CV, FT-200	20.0 t/h SP	6–26
$\dot{m}_m$	CV, FT-100	3.39 t/h	0–6
T_1	CV, TT-101, ceiling	69.5 °C SP	60–70
ℓ_1, ℓ_2, ℓ_3	CVs, LT-101 to 103	50 % SP	20–80
T_{pre}	CV, TT-400	67.0 °C SP	55–70
P_3	CV, PT-500	110 mbar SP	80–160
u_F, x_s	MVs, feed pump, spindle	66.7, 56.4 %	0–100
u_{p1}, u_{p2}, u_{p3}	MVs, pumps	50.5, 46.6, 46.9 %	0–100
x_p, x_c	MVs, preheat, cooling water	24.8, 40.9 %	0–100
R_1, R_2, R_3	MVs, recirculation	0 t/h	0–8
w_F	DV, feed solids	9.0 %	7–11
$T_{cw,t}$	DV, tower water	25 °C	12–36
P_m	DV, steam header	8.0 bar	4–10
ϕ	DV, fouling rate	1	0–4
T_{in}	DV, feed to preheater	62 °C	40–66
β_3	DV, LT-103 bias, unmeasured	0 %	±20
M_1, M_2, M_3	states, liquid	600, 400, 250 kg	
S_1, S_2, S_3	states, solids	97.1, 96.6, 120.0 kg	
H_i, as T_i	states, stored heat	65.32, 57.66, 49.52 °C	
$T_{w,i}$	states, tube walls	70.53, 63.53, 56.12 °C	
$R_{f,i}$	states, deposits	0	
E_s, E_{pre}	states, as θ_s, T_{pre}	72.14, 67.00 °C	
T_{cw}	state, water at condenser	25.0 °C	
$F, L_i, R_i, x_{s,p,c}$	states, actuators	as the MVs	

between units are held tightly or allowed to average [9]. The simulator puts both decisions, the thermocompressor and the fouling cycle on one train, and it logs what no plant can, the true solids of each tonne that leaves for the dryer.

> **HISTORY 1843: A Sugar Pan That Held Its Own Concentration**
>
> A United States patent to Norbert Rillieux of New Orleans, dated August 26, 1843, claimed four improvements in evaporating sugar juices. The fourth regulated the concentration of the sirup with a differential thermometer. One metal element sat in the sirup leaving the last evaporator, the other was acted on only by its vapor, and together they turned the spindle of a throttle valve in the juice supply. When the regulator showed the concentration too great, the valve opened and admitted more juice to the first evaporator. A simple thermometer would not do, the patent explains, because a sirup boiled in a vessel exhausted by an air pump changes its temperature with the pressure, while the difference between the sirup and its vapor indicates the concentration under all pressures [18]. The concentration was held with the feed, the pairing of case 3.

4.3 The Evaporator as Modeled

The simulated train is laid out in Figure 4.1. Skim milk passes a live-steam preheater and falls as a film down the tubes of effect 1, whose steam chest takes the thermocompressor's discharge. The liquid of each effect collects in its separator and sump and is pumped to the top of the next calandria, and the vapor of each effect condenses in the shell of the next, except that the thermocompressor draws part of effect 1's vapor back to the chest and effect 3's vapor goes to a condenser on cooling water. The product pump sends the concentrate along a product line, past the densitometer AT-301, to the dryer.

Table 4.1 groups the variables by role. The 54 entries of the state vector are 28 plant states and 26 cumulative ledgers of the water, solids and energy streams, for the whole plant and for each effect. One integration step advances ledgers and plant together, and because every held quantity is linear in the states, each balance closes from the state vector alone. The model works in SI units with absolute pressures, and because the engine's clock is the minute, every rate is multiplied by 60 on its way out.

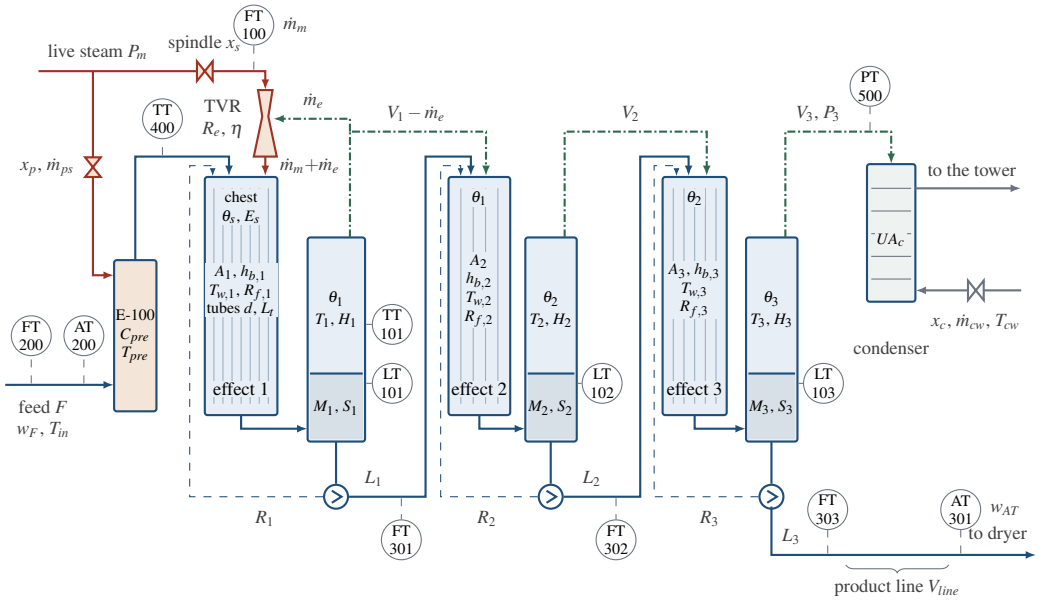

Figure 4.1 Three-effect falling-film evaporator as modeled. Live steam at P_m passes the spindle x_s and FT-100 into the thermocompressor (TVR), which entrains $\dot{m}_e$ of effect 1's vapor and discharges into the steam chest at θ_s, and a branch through x_p heats the preheater E-100. Feed F at w_F and T_{in} passes FT-200 and AT-200 to the top of effect 1. Each calandria of area A_i holds tubes of bore d and length L_t, a wall at $T_{w,i}$ and a deposit $R_{f,i}$, and each separator holds the liquid M_i and solids S_i at T_i under vapor at θ_i. Pumps move L_1 to L_3, thin dashed lines are the recirculation R_i, dash-dotted green lines are vapor, and effect 3's vapor V_3 condenses at P_3 against cooling water $\dot{m}_{cw}$ at T_{cw} through x_c. The product travels the line V_{line} to AT-301.

4.3.1 Mass, Solids and Stored Heat

Each effect is one well-mixed liquid, film, sump and piping together, whose mass and solids follow Equation (4.1).

$$\begin{aligned} \frac{\mathrm{d}M_i}{\mathrm{d}t} &= L_{in,i} - L_i - V_i \\ \frac{\mathrm{d}S_i}{\mathrm{d}t} &= L_{in,i}\, w_{in,i} - L_i\, w_i, \qquad w_i = S_i/M_i \end{aligned} \tag{4.1}$$

where $L_{in,1} = F$ with $w_{in,1} = w_F$, $L_{in,i} = L_{i-1}$ with $w_{in,i} = w_{i-1}$ for the other two, L_i is the flow each pump delivers and V_i the vapor the effect makes. No solids leave with the vapor. The level is $\ell_i = (M_i - M_{min,i})/M_{sump,i}$, and a pump delivers the flow of its first-order actuator times $\min(1, \max(0, \ell_i/0.02))$, so it starves below 2 %. The heat held by the liquid and by the steel at the boiling temperature, the separator, the vapor duct and the next calandria's shell, follows Equation (4.2).

$$\begin{aligned} \frac{\mathrm{d}H_i}{\mathrm{d}t} &= L_{in,i} h_{in,i} + Q_{b,i} - L_i c_p(w_i) T_i \\ &\quad - V_i h_{v,i} - UA_{loss}(T_i - T_{amb}) \\ T_i &= \frac{H_i}{c_w M_i - (c_w - c_s) S_i + C_{sep,i}} \\ \theta_i &= T_i - \mathrm{BPE}(w_i) \\ h_{v,i} &= c_w \theta_i + \lambda(\theta_i) + c_{pv}\mathrm{BPE}(w_i) \end{aligned} \tag{4.2}$$

where $h_{in,1} = c_p(w_F)T_{pre}$, $h_{in,i} = c_p(w_{i-1})T_{i-1}$ for the other two, $c_p(w) = c_w(1-w) + c_s w$, and $\lambda(T) = 2.501 \times 10^6 - 2370\,T$ in J/kg with T in °C, within 0.15 % of the steam tables between 40 and 80 °C. The liquid boils, so its vapor space sits at the saturation temperature θ_i, one boiling-point elevation below T_i, and θ_i is the temperature at which that vapor condenses in the next calandria. Saturation follows the Antoine equation $p_{sat}(T) = 133.322 \times 10^{8.07131 - 1730.63/(233.426+T)}$ in Pa. The elevation is an illustrative correlation, as shown in Equation (4.3).

$$\mathrm{BPE}(w) = 1.2x + 0.9x^4 \text{ K}, \qquad x = \frac{w}{1-w} \tag{4.3}$$

It is 0.23, 0.39, 1.76 and 2.54 K at 16, 24, 48 and 52 %, steep above 40 %.

4.3.2 Heat Through the Walls and the Vapor Chain

Every calandria wall sits between a condensing side on its whole area and a boiling film on its wetted share only, as shown in Equation (4.4).

$$\begin{aligned} C_w A_i \frac{\mathrm{d}T_{w,i}}{\mathrm{d}t} &= \underbrace{h_c A_i(\theta_{h,i} - T_{w,i})}_{Q_{c,i}} - \underbrace{U_{b,i} A_i f_{wet,i}(T_{w,i} - T_i)}_{Q_{b,i}} \\ \frac{1}{U_{b,i}} &= \frac{1}{h_{b,i}} + R_{f,i} \end{aligned} \tag{4.4}$$

where the heating temperature is $\theta_{h,1} = \theta_s$ in the steam chest, $\theta_{h,2} = \theta_1$ and $\theta_{h,3} = \theta_2$. An effect makes exactly the vapor that the next calandria condenses, plus, for effect 1, what the thermocompressor draws, and effect 3's vapor goes to the condenser, as shown in Equation (4.5).

$$\begin{aligned} V_i &= \frac{Q_{c,i+1}}{\lambda(\theta_i) + c_{pv}\mathrm{BPE}_i} + \dot{m}_e\,\delta_{i1}, \qquad i = 1,2 \\ V_3 &= \frac{\varepsilon_c \dot{m}_{cw} c_w(\theta_3 - T_{cw})}{\lambda(\theta_3) + c_{pv}\mathrm{BPE}_3}, \qquad \varepsilon_c = 1 - \mathrm{e}^{-UA_c/(\dot{m}_{cw}c_w)} \end{aligned} \tag{4.5}$$

with $\dot{m}_{cw} = x_c \dot{m}_{cw,max}$, the condenser duty taken as zero if $\theta_3 < T_{cw}$, and $P_3 = p_{sat}(\theta_3)$. The vapor inventories are neglected, but the heat stored at each boiling temperature is a state, so the chain is not an algebraic loop. A steam move reaches effect 3 through three stored heats in series.

4.3.3 Thermocompressor and Steam Chest

The thermocompressor is an illustrative energy balance. A fraction η of the motive steam's isentropic enthalpy drop from P_m to the suction pressure P_s compresses motive and entrained vapor together to the chest's pressure P_d, which gives the entrainment ratio of Equation (4.6).

$$\begin{aligned} R_e &= R_{e,max}\tanh(r/R_{e,max}) \\ r &= \max\left(0,\ \eta\,\Delta h_m/\Delta h_c - 1\right) \\ \Delta h_m &= c_{p,g}(T_m + 273.15)\left[1 - (P_s/P_m)^{\kappa}\right] \\ \Delta h_c &= c_{p,g}(T_1 + 273.15)\left[(P_d/P_s)^{\kappa} - 1\right] \end{aligned} \tag{4.6}$$

where $P_s = p_{sat}(\theta_1)$, $P_d = p_{sat}(\theta_s)$, T_m is the saturation temperature of the header, $\kappa = 0.3/1.3$ and $c_{p,g} = 1.3R_v/0.3$ for an ideal gas with $\gamma = 1.3$. The motive flow is $\dot{m}_m = x_s \dot{m}_{m,max} P_m/P_{m0}$ and the entrained flow $\dot{m}_e = R_e \dot{m}_m$. At the design point the jet compresses 25.05 kPa to 34.08 kPa, a ratio of 1.36, and entrains 1.566 kg/kg. The steam chest holds only its steel's heat, $E_s = C_{sh}\theta_s$, and condenses everything it receives at θ_s, as shown in Equation (4.7).

$$\frac{\mathrm{d}E_s}{\mathrm{d}t} = \dot{m}_m h_g(T_m) + \dot{m}_e h_{v,1} - (\dot{m}_m + \dot{m}_e)c_w\theta_s - Q_{c,1} \tag{4.7}$$

where $h_g(T) = c_w T + \lambda(T)$. A hotter chest raises P_d and lowers R_e, which is the mechanism behind both the marginal economy of Section 4.4 and the loss of economy to fouling.

4.3.4 Fouling and the Wetting Floor

Deposits grow faster in thicker concentrate, on hotter walls and on dry tubes, by the illustrative law of Equation (4.8).

$$\begin{aligned} \frac{\mathrm{d}R_{f,i}}{\mathrm{d}t} &= k_{f,i}\,\phi\left(\frac{w_i}{0.3}\right)^{n_w} \mathrm{e}^{\frac{E_f}{R}\left(\frac{1}{343.15}-\frac{1}{T_{w,i}}\right)} \\ &\quad \times \left[1 + k_{dry}(1 - f_{wet,i})\right] \\ f_{wet,i} &= \min\left(1,\ \frac{\Gamma_i}{\Gamma_{min}}\right), \qquad \Gamma_i = \frac{L_{in,i} + R_i}{A_i / L_t} \end{aligned} \tag{4.8}$$

where $T_{w,i}$ is absolute in the exponent, ϕ multiplies every rate, and A_i/L_t is the wetted perimeter of the tube bundle, 30.0, 12.5 and 17.5 m, the equivalent of 208, 87 and 121 tubes. Below the wetting floor the dry share of the tubes transfers no heat and fouls 31 times as fast. The constants make the deposits grow at 5.1, 5.8 and 8.5×10^{-6} m^2 K/W per hour at the clean design point, and a cleaning in place (CIP) resets them to zero at once. At the design point the three films run at 2.31, 3.09 and 1.48 times the floor, so effect 3, which needs 5.04 t/h to stay wet, is the first to break at turndown.

4.3.5 Preheater, Condenser Water and Actuators

The preheater is one liquid heat capacity heated by live steam, as shown in Equation (4.9).

$$\begin{aligned} C_{pre}\frac{\mathrm{d}T_{pre}}{\mathrm{d}t} &= F c_p(w_F)(T_{in} - T_{pre}) + \dot{m}_{ps}\lambda_{pre} \\ \dot{m}_{ps} &= x_p \dot{m}_{ps,max} P_m / P_{m0} \end{aligned} \tag{4.9}$$

The feed, the three pumps, the three recirculation pumps, the spindle and the preheat and cooling-water valves follow their commands through first-order lags of $\tau_p = 2$ s, and the water reaching the condenser follows the tower through $\tau_{cw}\,\mathrm{d}T_{cw}/\mathrm{d}t = T_{cw,t} - T_{cw}$.

4.3.6 Economy and Its Bound

The steam economy counts the water evaporated per kilogram of live steam, motive and preheat together, and it is bounded by one kilogram per effect plus the thermocompressor's recycle, with every kilogram of steam counted at the heat it gives up, as shown in Equation (4.10).

$$\varepsilon = \frac{V_1 + V_2 + V_3}{\dot{m}_m + \dot{m}_{ps}} < (N + R_e)\,\frac{\dot{m}_m q_m + \dot{m}_{ps}\lambda_{pre}}{\lambda(\theta_1)(\dot{m}_m + \dot{m}_{ps})} \tag{4.10}$$

with $N = 3$ effects, while $q_m = h_g(T_m) - c_w\theta_s$ is the heat a kilogram of motive steam gives up in the chest, 2.51 MJ/kg at the design point. The economy is 4.56 kg/kg there against a bound of 4.85 kg/kg.

4.3.7 Design Point

At the reset of case 2 the flows into, between and out of the effects are 20.00, 11.12, 7.46 and 3.75 t/h, the solids 9.0, 16.19, 24.14 and 48.00 %, and the holdups 600, 400 and 250 kg, residences of 3.2, 3.2 and 4.0 min. The effects boil at 65.32, 57.66 and 49.52 °C under vapor at 65.09, 57.27 and 47.76 °C, the chest is at 72.14 °C, and the temperature differences through the three calandrias are 6.82, 7.43 and 7.75 K. The effects make 8.88, 3.66 and 3.71 t/h of vapor with duties of 5.77, 2.33 and 2.41 MW, the condenser takes 2.46 MW from 147 t/h of cooling water warmed from 25.0 to 39.4 °C, and the preheater lifts the feed from 62 to 67 °C with 0.11 MW. The clean overall coefficients are 2352, 2095 and 1478 W/(m^2 K).

4.3.8 Initialization and Numerical Integration

Table 4.2 lists the parameters. A case starts from the steady state that belongs to its own configuration, with the feed, the levels, the preheat and the pressure of effect 3 at their setpoints and the free handle, the motive steam on steam pairing or the feed on feed pairing, placing the product on its solids

Table 4.2 Parameters of the Evaporator, Marked Visible Where the Show Model Dialog Offers a Slider and Hidden Where It Does Not

Symbol	Meaning	Value	Class
A_1, A_2, A_3	heating areas	360, 150, 210 m^2	visible
$h_{b,i}$	film and wall, W/(m^2 K)	3075, 2650, 1735	visible
h_c	condensing coefficient	10 000 W/(m^2 K)	hidden
C_w	tube wall per area	6 kJ/(m^2 K)	hidden
$C_{sep,i}$	steel at T_i, MJ/K	4.0, 3.0, 3.0	hidden
C_{sh}	steam chest steel	3.0 MJ/K	hidden
UA_{loss}, T_{amb}	loss per effect	500 W/K, 25 °C	hidden
d, L_t	tube bore, length	46 mm, 12 m	hidden
Γ_{min}	minimum wetting rate	0.08 kg/(m s)	visible
$M_{min,i}$	film and pipe holdup	350, 200, 120 kg	hidden
$M_{sump,i}$	sump, 0 to 100 %	500, 400, 260 kg	visible
F_{max}	feed pump	30 t/h	hidden
$L_{max,i}$	transfer, product pumps	22, 16, 8 t/h	hidden
R_{max}	recirculation pump, each	8 t/h	hidden
τ_p	actuator lag	2 s	hidden
$\dot{m}_{m,max}$	motive steam, full spindle	6 t/h	visible
P_{m0}	design header pressure	8 bar	hidden
η	ejector efficiency	0.28	visible
$R_{e,max}$	entrainment ceiling	3 kg/kg	hidden
UA_c	condenser	171 kW/K	visible
$\dot{m}_{cw,max}$	cooling water, full valve	360 t/h	hidden
τ_{cw}	cooling-water circuit	15 min	visible
C_{pre}	preheater liquid	590 kJ/K	hidden
$\dot{m}_{ps,max}$	preheat steam, full valve	0.72 t/h	hidden
λ_{pre}	preheat steam latent heat	2200 kJ/kg	hidden
$k_{f,i}$	fouling rates, 10^{-9} m^2 K/J	3.4, 3.8, 3.8	visible
E_f, n_w	fouling activation, exponent	80 kJ/mol, 1.5	hidden
k_{dry}	fouling on dry tubes	30	hidden
V_{line}	product line to AT-301	73 L	visible
T_{max}	product temperature limit	70 °C	visible
t_0	hours since CIP at reset	0 h	visible
ν	instrument noise scale	1	visible
c_w, c_s, c_{pv}	water, solids, vapor, J/(kg K)	4187, 1400, 1900	fixed
R_v	gas constant of steam	461.5 J/(kg K)	fixed

setpoint. Newton's method solves the nine unknowns the mass balances do not give explicitly, θ_s, θ_1 to θ_3, the three wall temperatures, the free handle and the cooling-water valve, while a fixed point on the mass balances gives the flows and the solids. At the reset of case 2 the residual is 5.5×10^{-13} after 12 iterations, and no plant state other than the deposits moves by more than 1×10^{-9} per minute. Below the design feed a second steady state exists, with effect 3 under its wetting floor, so the solver walks the feed down from 20 t/h in steps of at most 1.44 t/h to stay on the wet branch. A start that does not converge, needs an actuator beyond its range or leaves a film below its floor, as a feed setpoint of 13 t/h chosen before the run would, gives way to the case's own steady state, from which the loops then move the plant. The plant advances by classical fourth-order Runge-Kutta steps of 0.5 s on the engine of Section 1.6.1, the loops execute every 5 s, a row is logged every 5 s (14 400 rows in the 20 h case), and every random sequence is seeded at 20261103, so a noisy run replays identically. The tube walls are the fastest states, with time constants $C_w/(h_c + h_{b,i})$ of 0.46 to 0.51 s, inside the stability limit of the method at this step, and halving the step moves the temperatures and solids of case 1 at 11 min by less than 1×10^{-3}. Over the 480 min challenge and over a 300 min fouling run with a CIP the water, solids and energy ledgers close below 1×10^{-9} relative, one of the 37 checks the application passes. The speed selector spans 0.25 to 300 simulated minutes per second of real time, 5 by default, and leaves the trajectory unchanged.

4.3.9 Assumptions

Each effect is one well-mixed liquid at its boiling point, while a real falling film is plug flow down the tubes and often split into passes, so the model's residence-time distribution is too broad and its three- to four-minute holdups stand for the whole train's residence. The vapor inventories are neglected and the pressure of each vapor space follows the heat stored at its boiling temperature, which holds while the vapor spaces are small against the stored heat and fails for a fast pressure transient. Swell and shrink of the boiling sumps and the flash of condensate between calandrias are not modeled.

The elevation, the fouling law, the wetting floor with its linear loss of wetted area and the thermocompressor characteristic are illustrative. The thermocompressor is an energy balance with ideal-gas enthalpies and a smooth ceiling rather than a manufacturer's map, and real ejectors break down above a critical compression ratio, which the model does not reproduce. The Antoine equation is used beyond the range it was fitted for in the motive steam, where it gives 169.7 °C at 8 bar. Viscosity does not lower the film coefficient, so U falls only by fouling and by dry tubes, and recirculation changes the wetting of a lumped effect but not its composition.

The condenser is an ε-NTU exchanger with a fixed UA_c, no non-condensable gas and no vacuum pump. The densitometer turns density into solids exactly, with no temperature or fat error, CIP is instantaneous, and mechanical vapor recompression is not included, because a fan across effect 1 would only replace the thermocompressor's recycle and one across the whole train would need about 25 K of lift. The 70 °C limit, the 69.5 °C ceiling and all sizes and capacities are teaching values, not a design basis.

4.4 The Vapor Chain and the Temperature Budget

4.4.1 Two Paths Through One Train

Heat and concentration leave the steam chest by different routes. Heat travels as vapor and needs only to raise the temperatures that drive it through each calandria. Concentration travels as liquid, through the holdups of the three effects and then along 73 L of product line, 85 s at the design flow, to a densitometer that reports every 5 s. The step of case 1 separates the two, as shown in Figure 4.2. With the noise off, and read at every 0.5 s step, the vapor made in effect 1 shows half of its tenth-minute change after 16 s, because the thermocompressor draws more of that vapor as soon as its motive flow rises. The second and third effects follow at 75 and 82 s, and AT-301, which first moves by 0.01 % after 150 s, at 380 s. The vapors overshoot their tenth-minute values by 12 to 17 % at about four minutes, because the solids building up in each effect raise its boiling-point elevation and take back part of the temperature difference the heat had gained. With the elevation held at its design values, in a copy of the model outside the application, the overshoot shrinks to 2 to 4 %.

> **TRY IT IN THE APP** **1 · Step the Steam, Time the Chain**
>
> Load the scenario, press Run and watch the stopwatch beside the staircase. At one minute the motive steam rises by 0.10 t/h and nothing corrects it. The vapor of effect 1 arrives first, after about 20 s, effects 2 and 3 after about 75 and 70 s, and the densitometer AT-301 last, after about 400 s.

What delays the second and third effects is storage rather than transport. The condenser holds the bottom of the staircase, so a larger duty through each calandria needs a larger temperature difference, and the liquid and the steel of the chest and the effects must be heated to it. Their heat capacities at the boiling temperature are 3.0, 6.24, 4.41 and 3.71 MJ/K, with 2.16, 0.90 and 1.26 MJ/K more in the tube walls. The step brings 70 kW of extra heat once FIC-100 has moved the motive flow, 95 % of the way after 53 s, and of the 7.11 MJ it has supplied two minutes after the move, 5.35 MJ have gone into raising those temperatures rather than into vapor for the next effect.

Pause. Try Exercise 4.1 before reading on.

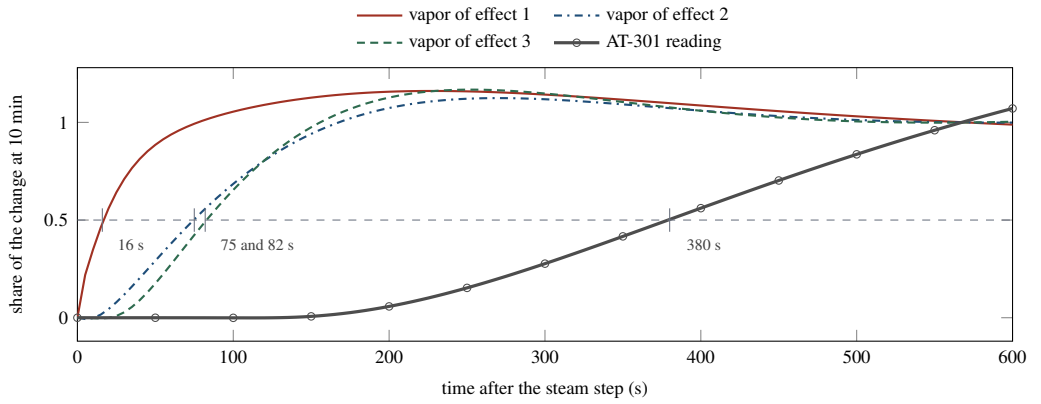

Figure 4.2 Vapor chain after a 0.10 t/h step of the motive steam in case 1, with the solids loop off and the instrument noise off. Each trace is the change since the step as a share of its mean over the tenth minute, 0.096, 0.062 and 0.055 t/h for the vapor of the three effects (solid, dash-dotted and dashed) and 1.80 % for the AT-301 reading (very thick, open circles), and the ticks mark where each crosses one half.

4.4.2 The Marginal Kilogram of Steam

The economy of the train is an average, and the thermocompressor makes the margin dearer than the average. At the design point the live steam evaporates 4.56 kg of water per kilogram and the motive steam alone 4.80 kg. Solved at steady state with the feed and the levels fixed, 0.10 t/h more motive steam evaporates 0.24 t/h more water, a marginal economy of 2.41 kg/kg, and 0.20 t/h more evaporates 0.45 t/h, 2.23 kg/kg. The extra steam heats the chest from 72.14 to 73.15 °C, which raises the pressure the jet must reach and lowers its entrainment from 1.566 to 1.530 kg/kg, so each added kilogram of motive steam brings less recycled vapor with it. The size of the effect belongs to the illustrative characteristic of Equation (4.6), in which the entrainment falls as the compression ratio rises. The bound of Equation (4.10) counts that recycle, and the train reaches 94.0 % of it at the design point, 92.6 % after 20 h of fouling (4.14 against 4.47 kg/kg) and 95.5 % at 13 t/h (5.37 against 5.62 kg/kg).

4.4.3 The Temperature Budget of the First Effect

The product temperature limit applies to the first effect, but the first effect's temperature is set below it. Reading the staircase upward from the condenser gives Equation (4.11).

$$T_1 = \theta_3 + \mathrm{BPE}_3 + \Delta T_3 + \mathrm{BPE}_2 + \Delta T_2 + \mathrm{BPE}_1 \tag{4.11}$$

where $\Delta T_i = \theta_{h,i} - T_i$ is the temperature difference through calandria i. At the design point this is $47.76 + 1.76 + 7.75 + 0.39 + 7.43 + 0.23 = 65.32\,^{\circ}\mathrm{C}$, 4.18 K below the 69.5 °C ceiling. The drop through effect 1 and its deposit do not appear, because they lie above T_1 and lift the steam chest instead. Fouling therefore spends the margin through the second and third effects, as shown in Figure 4.3. When the ceiling binds in case 7 at 14.9 h, $\Delta T_2 + \Delta T_3$ has grown from 15.18 to 19.29 K, while θ_3 is still 47.77 °C and the three elevations total 2.34 K against 2.39.

Table 4.3 separates the two consequences of fouling at the deposits reached after 20 h. The deposit in effect 1 alone lifts T_1 by 1.19 K, because a hotter chest entrains less and sends more of effect 1's vapor down the train, and it costs 0.38 kg/kg of economy. The deposits in effects 2 and 3 lift T_1 by 4.20 K and cost 0.16 kg/kg. With all three in place, holding 48 % at 20 t/h needs $T_1 = 71.10\,^{\circ}\mathrm{C}$, above the 70 °C limit, so late in the cycle the evaporator can cut its feed, run hot or leave the product thin, and nothing else.

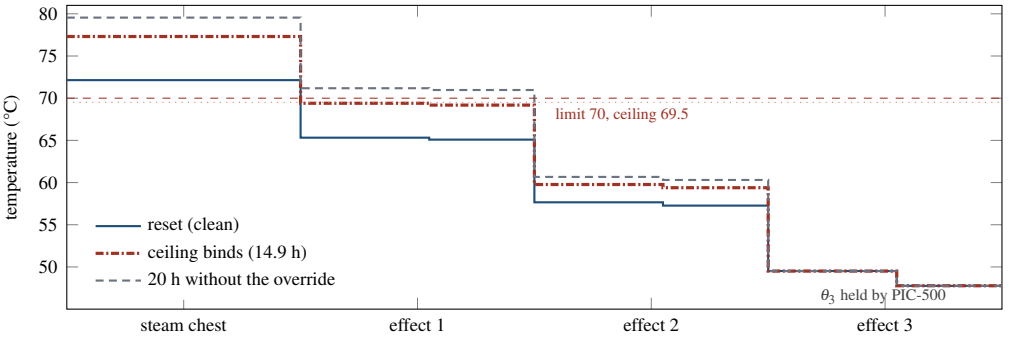

Figure 4.3 Temperature staircase of case 7 read from the engine. Each effect is drawn as its boiling product at T_i and then its vapor at θ_i, one boiling-point elevation lower, at the reset (solid), when the ceiling binds (dash-dotted) and at 20 h without the override (dashed), and the drops between the steps are the temperature differences through the calandrias. From the reset to the moment the ceiling binds, the drops through effects 2 and 3 widen from 15.2 to 19.3 K over a bottom held at 47.8 °C, and without the override the first effect reaches 71.18 °C at 20 h.

Table 4.3 Steady State at 20 t/h and 48 % With the Deposits of 20 h Placed in One Effect at a Time

Deposits in	T_1 (°C)	θ_s (°C)	R_e (kg/kg)	$\dot{m}_m$ (t/h)	ε (kg/kg)
none	65.32	72.14	1.566	3.39	4.56
effect 1	66.51	74.78	1.187	3.71	4.18
effects 2 and 3	69.52	76.40	1.509	3.51	4.40
all three	71.10	79.44	1.130	3.86	4.03

R_f = 1.27, 1.38 and 1.89 × 10^{-4} m^2 K/W, the values of case 7 at 20 h.

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

4.5 Implemented Controller

Every $\Delta t_c = 5$ s, once per ten physics steps and first at 5 s after a priming execution at the reset, the controller executes ten PI loops, a ratio station, a wetting-floor function and a supervisor, and the integrals advance by that interval. Within a physics step a scripted event that is due comes first, then the controller when a scan is due, then the integration, then the checks for dry-out, sump levels and the temperature limit, and last the outputs and the logged row.

4.5.1 Instruments

No loop sees the state directly. Each transmitter follows its true value through an exact first-order lag, and colored noise, itself filtered to first order and refreshed at every control step, is added to the result. FT-200 and FT-100 lag 1 s with 0.3 % and 0.5 % of reading, the effect thermometers lag 5 s with 0.04 K, the level transmitters are damped 5 s with 0.5 % of span, FT-301 to FT-303 lag 1 s with 0.5 %, PT-500 lags 1 s with 0.15 mbar, TT-400 lags 3 s with 0.04 K, and the refractometer AT-200 lags 30 s with 0.02 % solids. LT-103 may carry an unmeasured bias. The densitometer AT-301 is the framework's shared sampled analyzer fed by a transport delay keyed to the cumulative product volume $\int L_3/\rho(w_3)\,\mathrm{d}t$, with $\rho(w) = 1/[(1-w)/1000 + w/1600]$ in kg/m^3, so the concentrate it sees left effect 3 one line volume earlier. It draws a sample every 5 s, publishes it 5 s later with 0.04 % of white noise and holds it, so at the design flow its reading is about 90 s old. One parameter scales all of the noise.

Table 4.4 Pairing, Direction and Shipped Tuning of Each Loop

Tag	CV	MV	a	K_c	τ_I, output
FIC-200	F	u_F	+1	2	0.1 min, 0–100
FIC-100	$\dot{m}_m$	x_s	+1	5	0.1 min, 0–100
AIC-301 steam	w_{AT}	trim a	+1	0.04	10 min, $-ff$ to $C-ff$
AIC-301 feed	w_{AT}	F_{SP}	−1	0.12	12 min, 6–30
TIC-601	T_1	steam limit	+1	0.1	3 min, 0 to D
FY-700	w_{AT}	feed limit	−1	0.12	6 min, 6 to F_{dem}
LIC-101 to 103	ℓ_i	u_{pi}	−1	0.6	20 min, 0–100
TIC-400	T_{pre}	x_p	+1	5	0.5 min, 0–100
PIC-500	P_3	x_c	−1	2	1 min, 0–100

Gains in % per t/h, % per t/h, t/h per %, t/h per %, t/h per K, t/h per %, % per %, % per K and % per mbar, in the order of the rows. Flow outputs in t/h. The tight level alternative is $K_c = 6$ with $\tau_I = 1$ min.

4.5.2 The PI Algorithm

All ten loops use the shared velocity-form PID law of Section 1.6.2, Equations (1.11) to (1.14), with no derivative action, so the running sum equals the proportional term. The action a is -1 for the loops whose actuator must rise with the measurement, the three level loops, PIC-500, the feed-paired AIC-301 and FY-700, and $+1$ for the others. Windup is removed from the integral term alone, by back-calculation whose tracking time equals τ_I, and every loop here has integral action. The clamp limits what is emitted and nothing else, and no loop here uses the shared law's optional slew limit. A loop in MANUAL tracks the operator's value, so the change back to AUTO is bumpless. Table 4.4 gives each loop's pairing, direction and shipped tuning.

4.5.3 The Steam Demand

On steam pairing, AIC-301 trims a steam-to-feed ratio. FFC-302 turns the measured feed into the motive steam the mass balance asks for, as shown in Equation (4.12).

$$\begin{aligned} ff &= F_m\,(1 - w_{F,m}/w_{SP})/\hat{\varepsilon} \\ D &= \text{clamp}(ff + a,\ 0,\ \dot{m}_{m,max}) \end{aligned} \tag{4.12}$$

where F_m is FT-200, $w_{F,m}$ is AT-200 with the feed-solids compensation on and the design 9 % without it, $\hat{\varepsilon} = 4.8$ kg/kg of water per kilogram of motive steam, a is AIC-301's output and D the steam demand, with $ff = 0$ when the ratio is off. The estimate is the design economy of the motive steam, $16.25/3.39 = 4.80$, so at the design point the ratio alone asks for 3.39 t/h and the trim rests at zero. Because it is an average, a drop of the feed solids from 9.0 to 8.4 % makes the ratio add 0.052 t/h where the steady state needs 0.076 t/h. The trim is limited to $-ff$ and the steam cap C less ff, switching the ratio on or off moves its term into the integral so the demand does not jump, and a change of pairing, or of the solids loop's state, makes the loop track for one step. On feed pairing AIC-301's output is the feed setpoint and the motive steam is its own setpoint. The feed then becomes the throughput manipulator of Section 1.4.3, moved by a quality loop.

4.5.4 The Ceiling and the Throughput Supervisor

TIC-601 is a high limit on the steam demand. Its setpoint is the ceiling $T_{lim} = 69.5\,^{\circ}$C, its measurement TT-101, and its output is clamped between zero and D, so while the product is cool it rests on the clamp at the demand and its own back-calculation keeps it from winding up. FIC-100 receives $\min(D, u_O)$. The override counts as limiting once its unclamped output falls 0.002 t/h below the demand, and it stops once that output returns to 0.03 t/h above it. FY-700 then acts on the feed in the same way. While the steam is capped, by the override or at 6 t/h, a second solids loop with its setpoint 0.3 % below w_{SP} computes a feed limit between 6 t/h and the operator's throughput F_{dem}, which

approaches the feed setpoint at 0.5 t/h per minute, and FIC-200 receives $\min(F_{dem}, u_T)$. Once the cap is released the limit walks back up at the loop's own pace. The supervisor declares CIP due while the override is limiting and either the feed cut exceeds 5 % or AT-301 reads more than 1 % below its setpoint, counted by a leaky timer that must reach 300 s, so a brief release does not restart the count. A feed rate cut by a selector when an equipment limit binds is also the structure of the engine and rotor overrides in Volume 4, Chapter 4, and Volume 4, Chapter 5 pairs torque and die-pressure overrides with a throughput supervisor of the same kind.

4.5.5 Levels, Preheat, Condenser and the Wetting Floor

LIC-101 to LIC-103 set the three pumps from the sump levels. The sumps hold 500, 400 and 260 kg between empty and full, 2.7, 3.2 and 4.2 min of outflow at the design point, so proportional action alone would return the levels with time constants of 2.3, 2.5 and 3.3 min at the averaging gain and of 14, 15 and 20 s at the tight one. Section 1.4.8 shows why the PI form of such a loop passes flow variation on, enlarged near its natural frequency. TIC-400 holds the preheated feed at 67 °C with the preheat steam, and PIC-500 holds effect 3 at 110 mbar, a saturation temperature of 47.8 °C, with the cooling water. FY-800 is protection logic, live in both modes when it is switched on, which sets each recirculation to $R_i = \mathrm{clamp}[(1+m)\Gamma_{min}A_i/L_t - L_{in,i,m},\ 0,\ R_{max}]$ with $m = 15\,\%$ and the inflow measured by FT-200, FT-301 or FT-302.

4.5.6 What Each Model Knows

The equations of Section 4.3 form the simulation model. The classical solids loop has no fitted dynamic model. Its process paths were measured here by open-loop steps from the reset state of case 1, with the solids loop off and the noise and the fouling off. By a two-point fit at 28.3 and 63.2 % of the change after 89 min, a 0.10 t/h step of the motive steam moves the true solids at AT-301 by 33.0 % per t/h with a time constant of 6.1 min and an apparent dead time of 4.3 min. For the feed, with its setpoint ramp at the 10 t/h per minute maximum so that the step is complete in 6 s, the gain is −8.3 % per t/h from a 1 t/h cut and −7.1 % per t/h from a 1 t/h increase, with 5.9 and 5.0 min for the cut. The two paths differ far more in gain than in time. The steam raises the evaporation of all three effects within two minutes, while a feed change concentrates effect 1 at once and reaches the others through the liquid, and with tight level loops the feed path's apparent dead time falls to 1.3 min. The shipped tunings give loop gains KK_c of 1.32 on steam and 1.0 on feed. The classical controllers use the mass balance with an average economy in Equation (4.12), the wetting requirement of FY-800 and the 0.3 % offset of FY-700, and none of them contains the falling entrainment of Section 4.4.

4.5.7 Model Predictive Control

A separate supervisory model predicts AT-301 and TT-101 from motive steam and feed requests, within the MPC option that Section 1.6.3 describes in general. Multivariable evaporator supervision motivates this comparison [11], while the receding-horizon form follows industrial MPC [19]. Central steps of $\pm 0.02\,\mathrm{t/h}$ steam and $\pm 0.2\,\mathrm{t/h}$ feed were applied after 120 min with clean calandrias, noise off, averaging levels and the quality and protective supervisors off. Four first-order paths were fit over 60 min. In output order (reported solids, first-effect temperature) and input order (steam, feed), their coefficients are

$$K = \begin{bmatrix} 33.454 & -8.465 \\ 8.239 & -1.121 \end{bmatrix}, \qquad \tau = \begin{bmatrix} 6.078 & 5.463 \\ 6.031 & 5.347 \end{bmatrix},$$
$$\theta = \begin{bmatrix} 4.234 & 5.264 \\ 0 & 5.548 \end{bmatrix}, \qquad G_{ij}(s) = \frac{K_{ij} e^{-\theta_{ij} s}}{1 + \tau_{ij} s}. \tag{4.13}$$

Gains have units of percentage points or kelvin per t/h, and times are in minutes. The fitted RMS errors are 0.0114 % and 0.0467 % for solids, and 0.00281 and 0.006 27 K for temperature. The implemented

model rounds the steam delay to 4 min for solids and the shared feed delay to 5.5 min, giving four lag states and nineteen delay cells. Independent replay gives solids RMS errors of 0.0125 and 0.0481 % for these rounded paths. This approximates the measured paths, including transport and reporting. It does not reconstruct the true composition.

Every 0.5 min, measured flows advance the discrete model about the flows and readings at activation. Fast flow tracking is assumed when future requests enter those same paths. A bias b_i for each output is updated as $b_i^+ = b_i + 0.2(y_{i,m} - \hat{y}_i - b_i)$, with $y_{i,m}$ the current instrument reading and $\hat{y}_i$ its nominal prediction. This corrects output offsets, without the general disturbance-model guarantees discussed in [20]. Changing plant parameters does not change K, τ or the delay cells. Neither feed composition nor a future disturbance enters this predictor.

The plan covers 30 min in sixty steps and twelve blocks of $(1,1,1,1,2,2,3,4,6,9,15,15)$ steps. For steam and feed requests u_s, u_f in t/h, the objective is

$$\begin{aligned} J = \sum_{k=1}^{60} \big[& (\hat{w}_k - w_{SP})^2 + 0.001(u_{s,k} - u_{s,0})^2 \\ & + 0.5(u_{f,k} - F_{SP})^2 \big] \\ + \sum_{j=1}^{12} & \big[2(\Delta u_{s,j})^2 + 0.1(\Delta u_{f,j})^2 \big] \\ + 10^5 \sigma_T^2 & + 10^4 \sigma_T , \end{aligned} \tag{4.14}$$

where $\hat{w}$ is biased predicted solids in percent, $u_{s,0}$ is the activation steam flow, F_{SP} is requested throughput, and $\sigma_T \geq 0$ is one temperature slack shared by the entire horizon. The weights apply to these engineering units. The predicted band is $\hat{T}_{1,k} \leq T_{lim} + \sigma_T$. Steam requests lie between zero and 6 t/h, feed between 6 t/h and the smaller of 30 t/h and F_{SP}, and block changes cannot exceed 0.10 and 0.25 t/h, respectively. Reduced run capacities replace the command maxima. If requested feed is below 6 t/h, it also lowers the feed minimum. A 10^{-8} diagonal ridge regularizes the quadratic program.

The plan replaces AIC-301's steam demand and can lower the feed request. All inner loops and selectors remain. AIC-301 tracks the applied steam request. A failed, nonoptimal, malformed, nonfinite or out-of-range plan is rejected, including invalid later moves. The first failed scan retains both applied requests, then feed returns at its existing ramp while the tracked PI controls steam. The next valid plan resumes supervision. The optimizer is active only in Auto with steam pairing and solids supervision enabled.

4.6 Cases and Their Measured Outcomes

The original eight cases and the shift appear in Table 4.5. Those original cases start from one reference configuration, steam pairing with averaging levels, the override and FY-700 on, and the ratio and recirculation off, and each departs from it only where its row says. Fouling runs at the design rate in every case, so every steam value read after an hour carries it. An hour of fouling at the design point needs 0.019 t/h more motive steam, or at fixed steam and feed would take the solids from 48.0 to 47.38 %. Every number in this section was read from the engine itself, run without its interface at the default seed, with each setting that departs from a case changed at $t = 0$ before the run starts, as a reader does before pressing Run (Section 1.6.4), and Table 4.6 gathers the scorecard's tiles case by case.

4.6.1 The Chain and the Two Handles

Case 1 prints the chain of Section 4.4 on the stopwatch beside the staircase, which scans every 5 s with the noise on and reads 20, 75, 70 and 400 s. Effects 2 and 3 arrive together or a scan apart, in either order depending on the noise, and over seeds 1 to 5 AT-301 arrives after 375 to 390 s. Ten

Table 4.5 Cases, Their Events and What Each Teaches

#	Case and event	Lesson
1	Solids loop off, steam +0.10 t/h at 1 min	vapor against liquid
2	On steam, w_{SP} 48 to 49 % at 5 min	delay sets the pace
3	On feed, the same step	throughput as the MV
4	w_F 9.0 to 8.4 % at 5 min	feed forward what moved
5	Tight levels, LT-103 +10 % at 10 min	sumps as buffers
6	Tower 25 to 32 °C at 5 min	the staircase lifts
7	20 h from clean	ceiling, cut, CIP due
8	Feed 16 to 13 t/h at 5 min, ratio on	the wetting floor
9	Feed 20 to 22 t/h, steam/feed MPC	delayed-output plan
10	480 min shift	graded challenge

Table 4.6 Scorecard Tiles at the End of Each Run, Which Lasts 60 min Except 12 min for Case 1, 30 min for Case 5, 120 min for Cases 4 and 6 and 20 h for Case 7

Case and configuration	Feed (t)	Off (t)	Worst (%)	Hottest (°C)	ε	Dry (s)
1 as shipped	4.0	0.26	2.28	67.03	4.50	0
2 as shipped	20.0	0.26	0.84	67.04	4.53	0
2 K_c 0.08, τ_I 5	20.0	2.00	2.52	67.04	4.53	0
2 K_c 0.02, τ_I 20	20.0	0.30	0.97	67.04	4.54	0
3 as shipped	19.9	0.30	0.96	67.04	4.54	0
3 K_c 0.16, τ_I 10	19.8	0.28	0.92	67.04	4.54	0
4 as shipped	40.0	0.35	1.18	67.05	4.51	0
4 ratio	40.0	0.35	1.19	67.05	4.51	0
4 ratio and compensation	40.0	0	0.38	67.05	4.51	0
5 tight	10.0	0	0.58	67.03	4.55	0.5
5 averaging	10.0	0	0.09	67.03	4.55	0
6 as shipped	40.0	0	0.74	67.67	4.48	0
6 12 h into the cycle[a]	39.2	1.49	1.24	69.64	4.24	0
7 as shipped	397.2	0	0.80	69.50	4.31	0
7 FY-700 off	399.9	17.39	5.86	69.46	4.31	0
7 override off	399.9	0	0.36	71.18	4.29	0
8 as shipped	13.6	0.61	1.94	67.06	5.25	2670
8 recirculation	13.6	0.75	2.49	67.06	5.28	0

Off: product outside ±1 % of the setpoint. Worst: largest error of the true solids at AT-301 after 10 min. ε: economy since reset, kg/kg.
[a] Hours since CIP at reset set to 12 before the run.

minutes after the move the chest and the three effects stand 0.81, 0.69, 0.54 and 0.40 K warmer and the true solids at AT-301 have climbed from 48.1 to 50.1 %, and to 51.3 % by 30 min, with nothing correcting them. Stepping the motive steam back down by 0.10 t/h once the case has settled, at 60 min, gives arrivals of 20, 75, 80 and 390 s on the same stopwatch, with effects 2 and 3 again within a scan of each other and AT-301 again last.

In case 2 the solids setpoint moves from 48 to 49 % at 5 min. On steam, with 0.04 t/h per % and 10 min, the true solids reach 48.9 % after 9.3 min and peak at 49.13 %, with a scorecard IAE of 11.9 % min over the hour, while the motive steam rises from 3.39 to 3.43 t/h, of which 0.030 t/h is the step's steady-state requirement and the rest mostly the hour's fouling. Doubling the gain and halving the integral time reaches 48.9 % after 6.0 min but swings between 46.48 and 51.31 % and has not settled by 60 min (IAE 61.3), and 0.02 t/h per % with 20 min reaches only 48.67 % in the hour (IAE 29.5). The delay and the holdups set the pace, and the gain chooses how the loop behaves inside it. The crossing time moves with the instrument noise, from 9.6 to 11.8 min over seeds 1 to 5 and 10.1 min with the noise off.

Case 3 pairs AIC-301 with the feed at 0.12 t/h per % and 12 min. The same step reaches 48.9 %

after 12.9 min, peaking at 49.01 % (IAE 16.1), and the feed settles near 19.85 t/h, so the hour processes 19.86 t of milk against 19.97 t in case 2. Most of the difference in speed is the tuning. At the steam loop's loop gain and integral time, 0.16 t/h per % and 10 min, the feed-paired loop reaches 48.9 % after 9.7 min, 0.3 min behind steam pairing, and peaks at 49.27 % against 49.13, while a gain of 0.25 reaches it after 7.8 min and overshoots to 49.63 %. The 0.3 min lies inside the spread the noise gives either loop, and with the noise off the two cross at 10.3 and 10.1 min. The cost does not go away. At steady state one more percent of solids costs 0.12 t/h of feed at fixed steam or 0.030 t/h of motive steam at fixed feed, and the feed pairing also cuts the feed for the fouling, which is why case 3 ends the hour 0.15 t/h below the design feed.

> **TRY IT IN THE APP** **3 · The Same Loop on Feed Rate**
>
> Load the scenario and press Run. At 5 min the solids setpoint moves from 48 to 49 %, and AIC-301 moves the feed instead of the steam. The true solids reach 48.9 % after 12.9 min and peak at 49.01 %, while the feed settles near 19.85 t/h, below the design 20 t/h.

4.6.2 Feed Solids and the Ratio

In case 4 the feed solids fall from 9.0 to 8.4 % at 5 min. More water must now leave each tonne of feed, 16.50 t/h in all against 16.25, which at steady state needs 3.46 t/h of motive steam against 3.39. Without a feedforward the loop waits for AT-301 and the true solids fall to 46.82 % at 17.3 min. The ratio alone follows the feed flow, which did not change, and gives the same run to within the flow transmitter's noise, a fall to 46.81 %. With the feed-solids compensation the ratio raises the steam as soon as AT-200 sees the thinner milk, the product dips only to 47.62 % at 21.1 min and nothing goes off specification, the IAE falling from 28.2 to 15.3 % min. Before the dip the product barely rises, to 48.01 % at 12.3 min, because the steam acts before the thinner milk has reached effect 3. The trim still supplies about a third of the steam change, because the ratio's average economy asks for 0.052 t/h of the 0.076 t/h the plant needs.

Pause. Try Exercise 4.2 before reading on.

4.6.3 Tight Against Averaging

Case 5 starts with tight level loops and a steady feed. In the first 10 min each loop turns its transmitter's boiling noise into pump speed, the transfer flows vary with standard deviations of 0.78 and 0.74 t/h, the product pump runs between 25 and 62 %, and the product flow to the dryer swings between 2.06 and 4.89 t/h around its 3.75, a standard deviation of 0.60 t/h (0.44 to 0.50 t/h over seeds 1 to 5). The swings reach the concentrate, which wanders between 47.45 and 48.60 %, and at 19.8 min they take effect 3's film below its floor for one 0.5 s step. When LT-103 starts to read 10 % high at 10 min, the tight product pump drains effect 3's sump at once, reaching its 7.97 t/h limit at 10.2 min. With the averaging tuning the product flow's standard deviation is 0.023 t/h and that of the true solids 0.056 % (0.012 to 0.034 % over seeds 1 to 5), and the same bias drains the sump slowly, peaking at 4.23 t/h. Either way the true level settles near 39 % while LT-103 reads about 49, because no level loop can see its own transmitter's bias. The variation a tight loop passes on grows along a chain of inventories, which Chapter 11 measures stage by stage for ordering rules.

> **DID YOU KNOW? Tight Level Loops Spoil the Concentrate**
>
> From the simulator: in case 5, tight level loops leave a worst solids error of 0.58 % and 0.5 s of broken film, while averaging loops leave 0.09 % and none. Holding the sumps exactly passes the boiling noise on as flow (Table 4.6).

4.6.4 A Warm Condenser

In case 6 the cooling tower warms from 25 to 32 °C at 5 min, and the water reaching the condenser follows through its 15 min lag. PIC-500 opens the cooling water until the valve saturates at 28.0 min, and from then on the pressure of effect 3 floats up to 120.9 mbar. After two hours the chest and the three effects stand 2.54, 2.35, 2.11 and 1.85 K warmer, the entrainment has fallen from 1.57 to 1.49 kg/kg, the economy from 4.56 to 4.42, the motive steam has risen from 3.39 to 3.49 t/h, and the product dipped to 47.26 % at 40.2 min while the effects took up the heat for their new temperatures. The growth of the rise toward the top of the staircase is the fouling of those two hours. With the fouling off the same warming lifts the chest and the three effects by 1.76, 1.72, 1.72 and 1.73 K, the entrainment falls only to 1.54 kg/kg, the economy to 4.48 and the motive steam rises to 3.45 t/h, so the condenser moves every step together, as Equation (4.11) requires. On a clean evaporator the first effect ends at 67.7 °C. The hours since CIP are a start-up parameter, and setting them to 12 h before the run starts the case twelve hours into a fouling cycle, with T_1 at 68.21 °C. The same lift then puts the first effect on the override at 32.0 min, the feed is cut to 19.15 t/h, the product dips to 46.76 % and 1.49 t go off specification.

4.6.5 Twenty Hours of Fouling

Case 7 runs 20 h from clean at 20 t/h. At 6, 12 and 15 h the calandrias keep 91 to 92, 82 to 84 and 78 to 81 % of their clean film coefficients, the chest climbs from 72.1 to 74.0, 76.1 and 77.3 °C, the entrainment falls from 1.57 to 1.45, 1.31 and 1.24 kg/kg, the economy from 4.56 to 4.42, 4.25 and 4.17, and the motive steam rises from 3.39 to 3.49, 3.64 and 3.71 t/h. The first effect climbs from 65.3 to 66.8 and 68.4 °C and meets the 69.5 °C ceiling at 14.9 h, when the motive steam has peaked at 3.74 t/h of its 6, so the steam supply is not what runs out. From there TIC-601 holds the steam, lowering it to 3.51 t/h by 20 h, the product never passes 69.50 °C, and FY-700 cuts the feed to keep the solids in specification, never below 47.20 %. Once the cut has stayed beyond 5 % for five minutes the supervisor declares CIP due, at 19.6 h, and at 20 h the feed is 18.88 t/h, while the economy of 4.14 and the calandrias' 72, 73 and 75 % of their film coefficients appear in Figure 4.4.

Holding 48 % at 20 t/h this late in the cycle is infeasible within the temperature limit, and the two ways of pretending otherwise are both measured. Without FY-700 the product falls to 42.1 % by 20 h and 17.4 t go off specification, with CIP declared on the solids at 15.8 h. Without the override the first effect passes 70 °C at 16.7 h and reaches 71.18 °C, close to the temperature Table 4.3 gives for the 20 h deposits at the full rate (71.10 °C). The same trade between temperature and throughput appears in the crude preheat train of Volume 11, Chapter 3, where a saturated fuel valve leaves the charge rate as the last handle on the column inlet temperature.

Lowering the bottom of the staircase works the other way, as Equation (4.11) predicts. With the effect 3 pressure setpoint at 100 mbar, set before the run, the case starts with θ_3 at 45.88 °C, the first effect peaks at 68.27 °C, the override never acts and the full 20 t/h is processed for the 20 h, with an economy since reset of 4.41 against 4.31.

4.6.6 Turndown Into the Wetting Floor

Case 8 starts at 16 t/h with the ratio and its compensation on, where effect 3's film runs at 1.14 times its floor of 0.08 kg/(m s), and lowers the feed setpoint to 13 t/h at 5 min, ramping at 0.2 t/h per minute. Effect 3's film breaks at 15.3 min, when the feed has fallen to 13.93 t/h, and stays broken for 2670 s of the hour, at as little as 0.85 of the floor, so the dry share fouls fast and effect 3 ends the hour at 0.961 of its clean film coefficient. With recirculation FY-800 returns 1.12 t/h to the top of effect 3 at the end

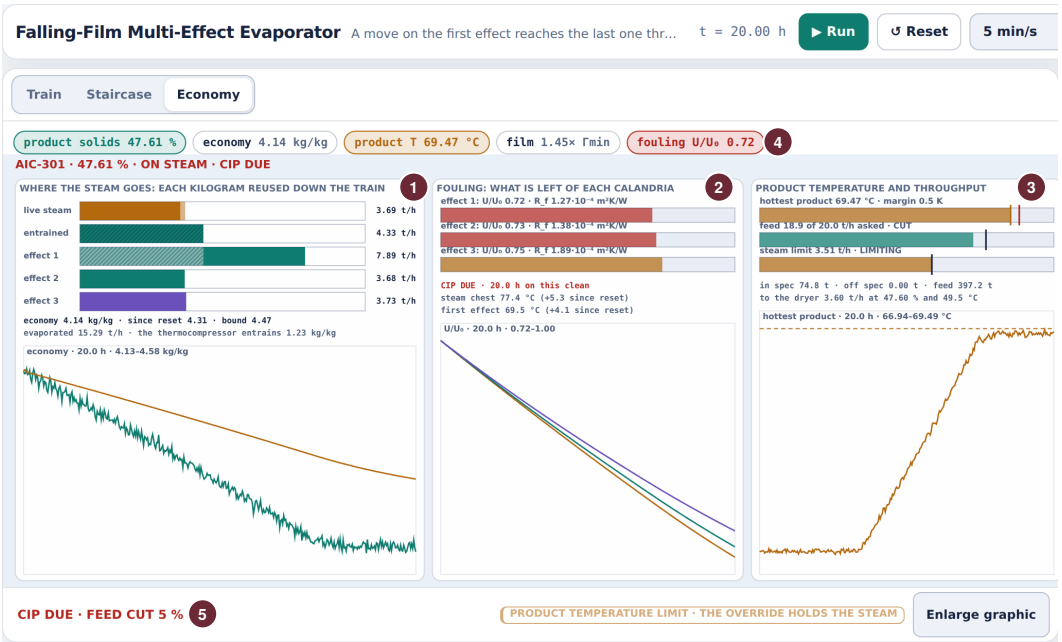

Figure 4.4 The evaporator in case 7, "Twenty hours of fouling," at $t = 20.00$ h with no control changed, in the Economy view. The steam ladder shows 3.69 t/h of live steam, 4.33 t/h of effect 1's vapor entrained at 1.23 kg/kg, and 7.89, 3.68 and 3.73 t/h made by the three effects, 15.29 t/h in all, for an economy of 4.14 kg/kg against 4.31 since reset and a bound of 4.47, and its sparkline falls through the run. The fouling panel gives 0.72, 0.73 and 0.75 of the clean film coefficients, with the chest at 77.4 °C and the first effect at 69.5 °C, 4.1 K above its reset. The product panel gives the hottest product at 69.47 °C, 0.5 K from the limit, the feed cut to 18.9 of the 20.0 t/h asked and the steam limited to 3.51 t/h, with 74.8 t in specification and none off of 397.2 t fed, and its sparkline sits on the preheater's 67 °C until the first effect overtakes it, then rises and flattens on the 69.5 °C ceiling. The badges add the worst film at 1.45 times its floor, and the status line reads CIP DUE · FEED CUT 5 %. The screen is cropped to its header and plant view. Callouts: (1) the steam ladder, (2) the fouling panel, (3) the product panel, (4) the badges and (5) the status line.

of the hour, the film never falls below 1.12 times the floor, and effect 3 keeps 0.988. Recirculation protects the heat transfer, not the concentration. After the ramp the solids wander between 47.06 and 49.94 % as shipped and between 46.79 and 50.49 % with recirculation, because the economy rises from 4.99 to 5.37 kg/kg at the lower load as the thermocompressor entrains more. The ratio, built on 4.8, asks at 60 min for 2.20 t/h of motive steam where 1.84 t/h suffice, and the trim has to take back 0.37 t/h.

4.6.7 The Shift Challenge

The challenge runs 480 min at 20 t/h from clean. The feed solids drop to 8.4 % at 1 h and the tower warms from 25 to 31 °C from 3 h, and the shift starts on feed pairing with tight level loops and no feedforward. The target asks for the true solids within 48 ± 1 % from 10 min on, at least 158 t of feed, a product never above 70 °C, an economy since reset of at least 4.35 kg/kg and no film below the floor. The configurations measured on the shift are compared in Table 4.7.

As shipped the product swings between 46.51 and 48.79 %, a worst error of 1.49 %, only 155.5 t of feed are processed because the loop cuts the feed whenever the product runs thin, and the tight loops break effect 3's film for 13 s. The suggested configuration holds 47.38 to 48.06 %, a worst error

Table 4.7 Configurations Measured on the Eight-Hour Shift

Configuration	Solids (%)	Off (t)	Feed (t)	Hottest (°C)	ε	Dry (s)
as shipped	46.51–48.79	0.64	155.5	67.05	4.49	13
suggested[a]	47.38–48.06	0	160.0	69.35	4.40	0
steam, averaging, no feedforward	46.77–48.06	0.45	160.0	69.36	4.40	0
steam, averaging, ratio alone	46.77–48.06	0.45	160.0	69.35	4.40	0
steam, feedforward, tight	46.94–48.87	0.02	160.0	69.34	4.40	0.5
feed, averaging	46.52–48.05	0.62	155.5	67.05	4.49	0

[a] steam pairing, ratio with feed-solids compensation, averaging levels.

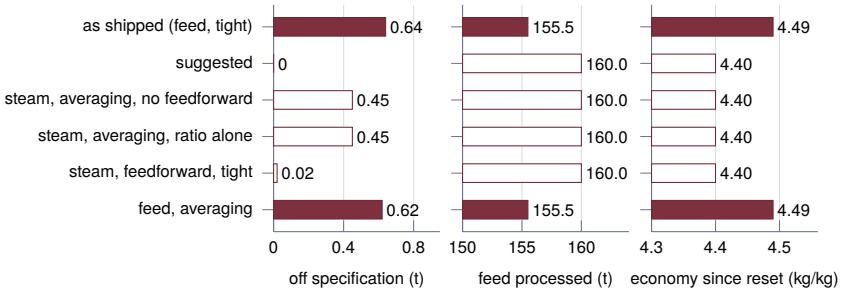

Figure 4.5 The eight-hour shift under the six configurations of Table 4.7. Dark bars hold the solids with the feed and white bars with the steam. The feed and economy axes start at 150 t and 4.3 kg/kg. Notice that both feed-paired runs show the higher economy and process less milk, while the suggested configuration alone sends nothing off specification.

of 0.62 %, processes all 160.0 t, peaks at 69.35 °C and ends at 4.40 kg/kg. Each of its three changes is needed on its own. Without the compensation, or without any feedforward, the product falls to 46.77 %, with tight levels to 46.94 % while effect 3's film breaks for 0.5 s (0 to 6.5 s over seeds 1 to 5), and on feed pairing even averaging levels leave only 155.5 t. The extra 4.4 t of milk cost 1.42 t of steam for 3.7 t more water, 2.6 kg/kg at the margin, so the shipped configuration's higher economy is throughput it gave up. Over seeds 1 to 5 the suggested configuration's worst error stays between 0.52 and 0.79 % and the shipped one's between 1.41 and 1.82 %. Figure 4.5 sets the six configurations side by side.

4.6.8 Steam and Feed Planned Through the Delayed Measurement

Case 9 requests 22 t/h instead of 20 t/h after 5 min, with averaging levels, compensated feedforward and recirculation. Identical seeds, measurements, limits and events are used for both structures. Table 4.8 reports 120 min runs, except 1200 min for fouling and 480 min for the suggested shift. Solids IAE uses the same reported measurement as the scorecard. Steam travel sums changes of the applied steam request after its selector, which differs from demand travel when the override binds.

On the throughput move, MPC reduces IAE from 54.92 to 20.28 % min and off-specification product from 1.123 to 0.199 t, with 43.689 versus 43.730 t processed. It plans around the delayed concentration response, as shown in Figure 4.6. The 22.40 versus 4.10 t/h steam-travel comparison also contains the cost of noisy feedforward. Disabling feedforward gives classical travel of 3.05 t/h, but IAE rises to 138.90 % min. MPC ignores that channel and its response is unchanged. Among nine hindsight classical settings ($K_c = 0.02, 0.04, 0.08$, $\tau_I = 5, 10, 20$ min), the smallest IAE is 44.57 % min at 0.08 and 10 min. This finite search establishes no global tuning optimum. Halving the physics step retains IAE 54.92070 versus 20.276 14 % min.

The predictive model omits AT-200's feed-composition compensation. On the suggested shift it therefore produces 0.466 t off specification, with worst error 1.225 %, and does not meet the

Table 4.8 Matched Classical and Predictive Supervision

Case, law	IAE (% min)	Travel (t/h)	Feed (t)	Off (t)	Plans
Throughput, Loops	54.92	22.40	43.730	1.123	0
Throughput, MPC	20.28	4.10	43.689	0.199	240
Setpoint, Loops	18.24	2.79	39.967	0.255	0
Setpoint, MPC	17.89	3.44	39.959	0.026	240
Feed solids, Loops	28.18	2.75	39.967	0.347	0
Feed solids, MPC	29.40	3.10	39.959	0.412	240
Warm water, Loops	25.90	2.75	39.967	0.000	0
Warm water, MPC	25.88	3.12	39.959	0.000	240
Fouling, Loops	236.71	19.85	397.249	0.000	0
Fouling, MPC	198.74	35.43	397.054	0.000	2400
Shift, suggested, Loops	66.79	82.92	159.962	0.000	0
Shift, suggested, MPC	89.64	12.43	159.924	0.466	960

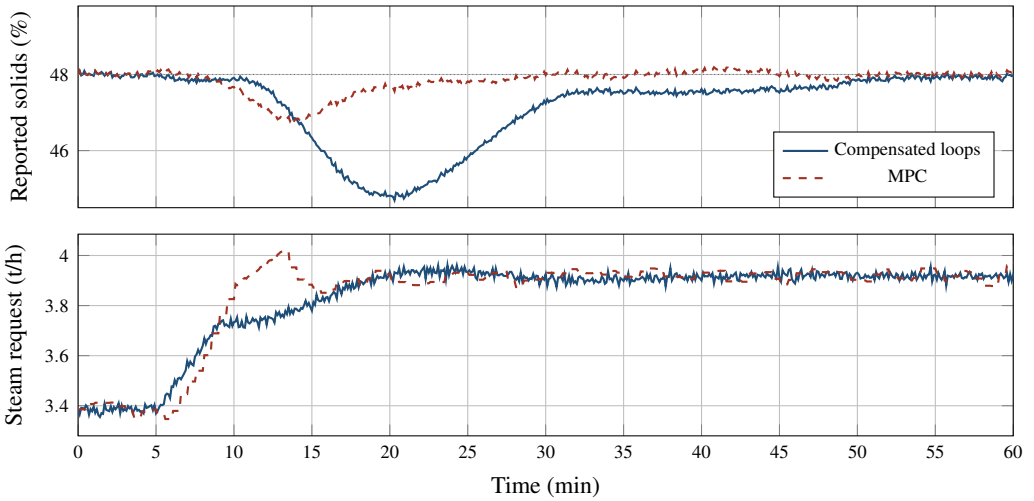

Figure 4.6 Concentration and steam requests after the throughput increase, compensated loops solid and MPC dashed, with the setpoint dotted. The first hour of the matched runs is shown.

unchanged challenge criterion. The suggested classical structure passes with no off-specification product and worst error 0.617 %. Selecting MPC on the as-loaded shift does not activate it because that case uses feed pairing. That run remains the original failure, rather than evidence of predictive control.

On fouling, the smaller IAE costs 35.43 t/h rather than 19.85 t/h of steam-request travel and 0.195 t of feed. Peak temperature is 69.546 °C versus 69.495 °C, with 65.6 versus 0 min above the 69.5 °C supervisory ceiling. Both stay below 70 °C. The fixed clean model and shared temperature slack do not guarantee the actual ceiling. The warm-water result is essentially tied, while the feed-solids case has greater IAE and off-specification mass under MPC.

No plan is rejected in these comparisons. Chrome averages 3.88 ms per solve over 240 plans, or 0.776 % of one core at unit playback. At the nominal 300 setting, solves alone require 2.33 cores. Actual playback on the test host reaches 117.6 simulated min/sec at 8.16 frames/sec, compared with 1.02 min/sec at 60.4 frames/sec at the unit setting. The highest setting is consequently a requested acceleration rather than a guaranteed rate.

4.7 Technology in Practice

What Measures It

The product is judged by its density. The vibrating-tube meter measures density so precisely that such meters often serve mainly as density meters and only secondarily as flowmeters, and the same instrument reports mass flow and the temperature of its tubes [21]. The other readings of the train are flows, temperatures, levels and one pressure, with the feed solids read by the refractometer AT-200. In the chapter AT-301 is the framework's shared sampled analyzer behind a transport delay keyed to the product volume, so at the design flow its reading is about 90 s old.

> **TECHNOLOGY** **Sensor: The Vibrating-Tube Density Meter**
>
> A vibrating-tube meter, the Coriolis flowmeter, keeps its tubes oscillating at their resonant frequency with a force coil. A denser fluid gives the tubes more mass, so they vibrate at a lower frequency, while the phase shift between two pickups gives the mass flow. A resistance thermometer on the tubes compensates the effect of temperature on the metal. Gas bubbles trapped in bent tubes spread the mass unevenly and cause errors in density and flow, and bent tubes may hold product in a way that is unacceptable in hygienic service, where straighter tubes fare better [21]. AT-301 in the chapter turns density into solids exactly, with 0.04 % of noise and a delay keyed to the product line.

What Moves It

Four kinds of final element act on the train. The spindle of the thermocompressor sets the motive steam, the feed pump sets the throughput, the transfer and product pumps move the liquid from sump to sump, and one valve sets the cooling water to the condenser. In the model every one of them follows its command through a first-order lag of 2 s, without stiction.

> **TECHNOLOGY** **Final Element: The Steam Jet and Its Spindle**
>
> An ejector drives a high-velocity gas or vapor such as steam through a tapered tube, which produces a vacuum at its throat [21]. In thermal vapor recompression, high-pressure steam compresses the vapor leaving an evaporator in such an ejector, and the mixture returns to the same unit as its heating medium. Part of the vapor must leave to keep the mass balance of the steam entering the unit [10]. In the chapter the spindle sets the motive flow under FIC-100 through a 2 s lag, and at the design point the jet compresses 25.05 kPa to 34.08 kPa. Real ejectors break down above a critical compression ratio, which the model does not reproduce.

Where the Controller Runs

Industrial loops are often nested or cascaded, so that the setpoint of one loop comes from another, and supervisory and lower-level loops run with cycle times from milliseconds to minutes. A distributed control system adds a supervisory level that oversees the controllers of local processes [22]. In the conventional structure of a processing unit, the dynamic constraint control above the basic loops combines PID algorithms, lead-lag blocks and high and low select logic, executed about every minute over basic controls executed every second, and model predictive control replaces that combination with a single controller [19]. The chapter's structure is of that conventional kind. Ten PI loops, a ratio station, a wetting-floor function and a supervisor execute every 5 s over a plant integrated every 0.5 s, and the low selects on the steam and the feed are its select logic. The optional predictive supervisor replaces only the steam demand of AIC-301 and leaves every inner loop and selector in place. Figure 4.7 draws the solids loop on its devices.

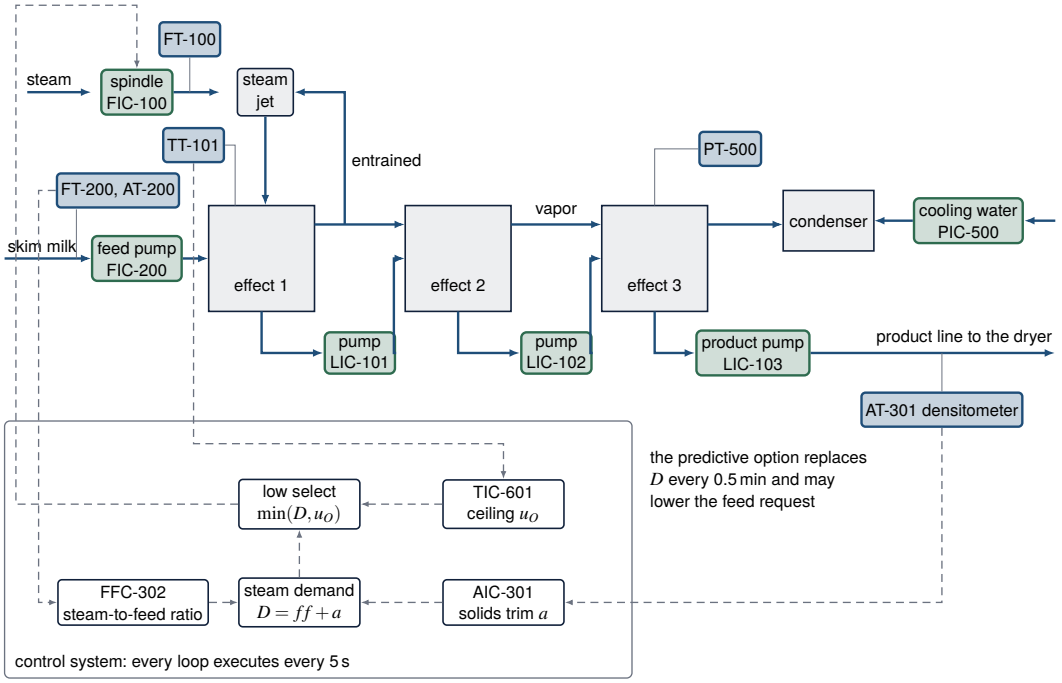

Figure 4.7 The solids loop on plant hardware. The ratio station turns the measured feed into the motive steam the mass balance asks for, AIC-301 trims it from the densitometer at the end of the product line, and the low select lets TIC-601 hold the steam down when the first effect reaches its ceiling. FIC-100 then moves the spindle of the steam jet. The feed, the three sumps and the pressure of effect 3 are held by their own loops on the feed pump, the transfer and product pumps and the cooling-water valve. Notice that the only reading of the product is the last device on the train, behind three effects and the product line.

> ⚠ **IN PRACTICE** **What Keeps an Evaporator Efficient**
>
> A dairy energy guide lists the common sources of inefficiency in evaporators: excessive venting, losses by radiation and convection, poor vacuum system performance, air and water leaks, fouling and poor separator efficiency. Its maintenance program prevents air leaks, because air does not condense and must be vented, cleans the heat transfer surfaces, stops water leaks that dilute the product and keeps the pressure profile at the optimum, since excess pressure raises the boiling point [10]. The chapter's condenser has no non-condensable gas and no vacuum pump, and case 6 shows how a warmer bottom lifts every step of the staircase.

What Is Modernizing It

Vapor recompression is done either by a mechanical compressor or by a steam jet ejector [23]. In mechanical vapor recompression the vapors are compressed by centrifugal compressors or turbo fans and a small amount of steam makes up the condensate formed during compression. Its steam economy can range from 10 to 30, against a typical 4 to 8 for thermal recompression, and it is most economical where electricity is cheap, while the steam jet is most economical where high-pressure steam is cheap. Membranes are a second direction. Reverse osmosis and ultrafiltration can concentrate a product before it reaches the evaporator, which lowers the water the evaporator must remove and so its energy [10]. The same guide puts the savings of process control at around 5 % or more for plants without updated control systems [10]. The chapter models neither a fan nor a membrane, and a fan across effect 1 would only replace the thermocompressor's recycle.

4.8 Limitations

The effects are well mixed, so the model's residence-time distribution is broader than a falling film's plug flow, and the vapor inventories, the swell of the boiling sumps and the flash of condensate are left out. The boiling-point elevation, the fouling law, the wetting floor and the thermocompressor characteristic are illustrative, the thermocompressor has no critical compression ratio, and the film coefficient does not fall with the viscosity of the concentrate.

The densitometer is exact apart from its delay and noise, with no laboratory check, and the pumps and valves are linear first-order actuators without stiction. Cleaning in place is an instantaneous reset, and mechanical vapor recompression is not modeled. The feedforward's economy estimate is fixed, so it over-steams at turndown, and the limits, sizes and capacities are teaching values for one train.

> **OPEN CHALLENGE** **Evaporation Without Live Steam**
>
> A 2022 industrial decarbonization roadmap of the United States Department of Energy names the electrification of evaporation and pasteurization as an important opportunity for food and beverage manufacturing, where these processes mostly run below 200 °C. It asks for studies of how a change of technology affects the final product, so that safety and quality are kept. Mechanically driven heat pumps such as vapor recompression already suit low-grade heat, but research is needed to extend them to higher temperatures [24]. The train of this chapter runs on live steam through a thermocompressor. Progress would be evaporation driven by electricity with a product shown to be unchanged.

The predictive supervisor uses one clean operating point and fixed delays, although product-line residence changes with flow. Output-bias correction does not estimate fouling-dependent gains or supply the omitted composition signal. The shift loss and the 65.6 min above the fouling ceiling quantify these restrictions.

> **OPEN CHALLENGE** **A Predictive Model That Ages With the Deposits**
>
> A survey of industrial model predictive control found a lack of methods to verify whether an identified model is adequate for control, or whether it has deteriorated significantly since the controller was commissioned. Only two adaptive algorithms had reached the market despite a strong incentive for a self-tuning controller, which the survey reads as the difficulty of adaptive control in the real world [19]. The predictive supervisor of this chapter keeps one clean model through a fouling cycle and spends 65.6 min above the supervisory ceiling. Progress would be a supervisor that detects when deposits have aged its model and updates it without losing the limit.

4.9 Conclusion

A three-effect milk evaporator is a train in which heat and concentration travel at different speeds, and the simulator records the true solids next to the densitometer's reading, which makes the gap between them measurable. With the design settings, 3.39 t/h of motive steam and 0.18 t/h of preheat steam evaporate 16.25 t/h, an economy of 4.56 kg/kg against a bound of 4.85. A steam step reaches half of its change in the vapor of the three effects after 16, 75 and 82 s and in the densitometer after 380 s, and the next kilogram of motive steam evaporates 2.4 kg against an average of 4.80. The first effect's temperature rests on the condenser and on the drops through the effects below it, so the deposits of a twenty-hour run lift it by 4.20 K from effects 2 and 3 and by 1.19 K from effect 1, which instead costs 0.38 kg/kg of economy. Over twenty hours the ceiling binds at 14.9 h and the supervisor declares CIP due at 19.6 h rather than run the product above 69.50 °C or off its solids. At equal loop gain the two pairings reach 48.9 % within 0.3 min of each other, inside the spread the instrument noise gives, and the feed pairing pays in milk. The shift improves from 0.64 t off specification and 155.5 t processed to none off and 160.0 t. The effects are well mixed and the fouling, wetting and thermocompressor laws are illustrative. Predictive supervision lowers throughput-step IAE from 54.92 to 20.28 % min, but produces 0.466 t off specification on a shift the compensated loops pass. Further work would

add plug flow in the tubes, a film coefficient that falls with viscosity, an ejector map with its critical compression ratio and a cleaning cycle with its hours offline. A concentration measured at the end of a chain has to be controlled through the whole chain.

KEY POINTS

- A steam move reaches the three vapors within two minutes, and the concentrate reaches the densitometer at the end of the product line only after about six.
- The next kilogram of motive steam evaporates 2.4 kg against an average of 4.80.
- The deposits of a twenty-hour run in effects 2 and 3 lift the first effect by 4.20 K, those in effect 1 by 1.19 K.
- At equal loop gain the steam and feed pairings reach 48.9 % within 0.3 min of each other, and the feed pairing pays in milk.

Where else this applies. Volume 4, Chapter 4 and Volume 4, Chapter 5 cut a feed rate when an equipment limit binds, and Volume 11, Chapter 3 trades temperature against throughput in the same way.

References

[1] C. Ramirez, M. Patel, and K. Blok, "From fluid milk to milk powder: Energy use and energy efficiency in the European dairy industry," *Energy*, vol. 31, no. 12, pp. 1984–2004, 2006.

[2] P. Quaak, M. P. C. M. van Wijck, and J. J. van Haren, "Comparison of process identification and physical modelling for falling-film evaporators," *Food Control*, vol. 5, no. 2, pp. 73–82, 1994.

[3] N. T. Russell, H. H. C. Bakker, and R. I. Chaplin, "A comparison of dynamic models for an evaporation process," *Chemical Engineering Research and Design*, vol. 78, no. 8, pp. 1120–1128, 2000.

[4] J. A. Winchester and C. Marsh, "Dynamics and control of falling film evaporators with mechanical vapour recompression," *Chemical Engineering Research and Design*, vol. 77, no. 5, pp. 357–371, 1999.

[5] "U.S. patent to Norbert Rillieux for an "improvement in sugar-works" (1843)," in *A Hammer in Their Hands: A Documentary History of Technology and the African-American Experience*, C. Pursell, Ed. The MIT Press, 2005, pp. 75–86.

[6] Z. Stefanov and K. A. Hoo, "A distributed-parameter model of black liquor falling film evaporators. Part I. Modeling of a single plate," *Industrial & Engineering Chemistry Research*, vol. 42, no. 9, pp. 1925–1937, 2003.

[7] H. T. El-Dessouky and H. M. Ettouney, "Multiple-effect evaporation desalination systems. thermal analysis," *Desalination*, vol. 125, no. 1–3, pp. 259–276, 1999.

[8] R. M. Price, P. R. Lyman, and C. Georgakis, "Throughput manipulation in plantwide control structures," *Industrial & Engineering Chemistry Research*, vol. 33, no. 5, pp. 1197–1207, 1994.

[9] K. A. McDonald, T. J. McAvoy, and A. Tits, "Optimal averaging level control," *AIChE Journal*, vol. 32, no. 1, pp. 75–86, 1986.

[10] A. Brush, E. Masanet, and E. Worrell, "Energy efficiency improvement and cost saving opportunities for the dairy processing industry: An ENERGY STAR guide for energy and plant managers," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep., Oct. 2011, sponsored by the U.S. Environmental Protection Agency.

[11] M. P. C. M. van Wijck, P. Quaak, and J. J. van Haren, "Multivariable supervisory control of a four-effect falling-film evaporator," *Food Control*, vol. 5, no. 2, pp. 83–89, 1994.

[12] 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.

[13] "U.S. patent to Norbert Rillieux for an "improvement in sugar-making" (1846)," in *A Hammer in Their Hands: A Documentary History of Technology and the African-American Experience*, C. Pursell, Ed. The MIT Press, 2005, pp. 87–88.

[14] J. H. Keenan, E. P. Neumann, and F. Lustwerk, "An investigation of ejector design by analysis and experiment," *Journal of Applied Mechanics*, vol. 17, no. 3, pp. 299–309, 1950.

[15] D. E. Hartley and W. Murgatroyd, "Criteria for the break-up of thin liquid layers flowing isothermally over solid surfaces," *International Journal of Heat and Mass Transfer*, vol. 7, no. 9, pp. 1003–1015, 1964.

[16] S. Paramalingam, J. Winchester, and C. Marsh, "On the fouling of falling film evaporators due to film break-up," *Food and Bioproducts Processing*, vol. 78, no. 2, pp. 79–84, 2000.

[17] J. Visser and T. J. M. Jeurnink, "Fouling of heat exchangers in the dairy industry," *Experimental Thermal and Fluid Science*, vol. 14, no. 4, pp. 407–424, 1997.

[18] N. Rillieux, "Improvement in sugar-works," United States Letters Patent No. 3,237, Aug. 1843, patented August 26, 1843; image of the specification and drawings served by the United States Patent and Trademark Office.

[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] 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] 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] U.S. Department of Energy, Advanced Manufacturing Office, "Use vapor recompression to recover low-pressure waste steam," U.S. Department of Energy, Washington, DC, Energy Tips: Steam, Steam Tip Sheet 11, DOE/GO-102012-3415, Jan. 2012, oSTI 1219562.

[24] U.S. Department of Energy, "U.S. Department of Energy's industrial decarbonization roadmap," U.S. Department of Energy, Washington, DC, Tech. Rep. DOE/EE-2635, Sep. 2022.

4.10 Exercises

4.1 ⌕ Concept check · 1 min

In case 1 the motive steam steps up by 0.10 t/h with the solids loop off. The vapor made in effect 1 shows half of its change after 16 s, and the vapors of effects 2 and 3 after 75 and 82 s. Why does effect 1 answer first?

(a) Effect 1 holds the least liquid, so its boiling temperature rises fastest

(b) AIC-301 reads the densitometer and moves effect 1 before the others

(c) The thermocompressor draws more of effect 1's vapor as soon as its motive flow rises

(d) PIC-500 holds the pressure of effect 1, so its vapor follows the steam at once

4.2 ⌕ Concept check · 2 min

In case 4 the feed solids fall from 9.0 to 8.4 %. With the steam-to-feed ratio on and its feed-solids compensation off, the true solids fall to 46.81 %, the same run as with no feedforward. Why does the ratio alone not help?

(a) Its economy estimate of 4.8 kg/kg is too high, so its correction is too small

(b) It moves the motive steam, whose effect reaches AT-301 only after six minutes

(c) AIC-301 limits its trim to cancel the ratio's output after the disturbance

(d) It follows the measured feed flow, and the feed flow did not change

4.3 ⌕ Concept check · 2 min

Table 4.3 places the deposits of a twenty-hour run in one effect at a time. Those of effects 2 and 3 lift the first effect by 4.20 K, while the deposit of effect 1 alone lifts it by 1.19 K. Which statement explains the difference?

(a) The first effect sits on the condenser and the drops of effects 2 and 3, and its own drop lies above it

(b) Effects 2 and 3 have more heating area, so their deposits hold more resistance in total

(c) Deposits in effects 2 and 3 raise the boiling-point elevations of the concentrate

(d) The deposit of effect 1 lowers the entrainment, which cools the first effect's liquid

4.4 Theory and calculation · 6 min

The evaporator takes 20 t/h of skim milk to 48.0 % solids, and no solids leave with the vapor.

(a) Find the product flow and the water evaporated with feed at 9.0 %, and again with feed at 8.4 %. *Checkpoint:* 3.75 and 16.25 t/h, then 3.50 and 16.50 t/h.

(b) Find the motive steam that the compensated ratio of Equation (4.12) asks for at each feed solids, with $\hat{\varepsilon} = 4.8\,\mathrm{kg/kg}$, and its change. *Checkpoint:* 3.385 and 3.438 t/h, a change of 0.052 t/h.

(c) At steady state the thinner feed needs 0.076 t/h more motive steam. Find the water evaporated per added kilogram of motive steam for this disturbance, and the share of the steam change left to the trim of AIC-301. *Checkpoint:* About 3.3 kg/kg, and about one third.

4.5 Theory and calculation · 7 min

Use Equations (4.3) and (4.11) and the design point of Section 4.3.

(a) Find the boiling-point elevation of the product at 48 % solids. *Checkpoint:* 1.76 K.

(b) Add up the temperature budget of the first effect at the design point, with $\theta_3 = 47.76\,^{\circ}\mathrm{C}$, $\Delta T_3 = 7.75\,\mathrm{K}$, $\Delta T_2 = 7.43\,\mathrm{K}$ and elevations of 0.39 and 0.23 K in effects 2 and 1, and find the margin to the 69.5 °C ceiling. *Checkpoint:* 65.32 °C and 4.18 K.

(c) The effect 3 pressure setpoint is lowered from 110 to 100 mbar. Find the new θ_3 from the Antoine equation of Section 4.3. Without the override, case 7 at 110 mbar takes the first effect to 71.18 °C at 20 h. Estimate its peak at 100 mbar if every step of the staircase moves with the bottom, and say whether the ceiling would bind. *Checkpoint:* 45.88 °C, and about 69.30 °C, just under the ceiling.

4.6 Theory and calculation · 6 min

Effect 3's sump holds 260 kg between empty and full, and the product pump passes 8 t/h at full speed. Its design duty is 3.75 t/h, a command of 46.9 %. LIC-103 runs at 0.6 %/% with 20 min (averaging) or 6 %/% with 1 min (tight).

(a) Find the time constant with which proportional action alone would return the level, $M_{sump}/(K_c L_{max})$, for each tuning. *Checkpoint:* 3.25 min and 19.5 s.

(b) At 10 min in case 5 LT-103 starts to read 10 % high. Find the immediate change of the pump command from the proportional term, and the product flow it asks for, for each tuning. *Checkpoint:* Averaging: +6 %, 4.23 t/h. Tight: +60 %, the full 8 t/h.

(c) Once the loop has brought LT-103 back to its setpoint, find the true level and the concentrate drained from the sump. *Checkpoint:* About 40 %, and 26 kg.

4.7 App challenge · 8 min

Load scenario 1, Step the steam, time the chain, at the default seed and a speed of 1. The motive steam steps from 3.39 to 3.49 t/h at 1 min with the solids loop off.

(a) Run to 12 min and read the four half-change times that the vapor chain panel prints ten minutes after the move. *Checkpoint:* 20, 75, 70 and 400 s.

(b) Run on to 30 min and read the true solids at AT-301, the `product_solids_pct` column of the exported file. *Checkpoint:* 51.3 %.

(c) Run on to 60 min, set the Motive steam setpoint on the FIC-100 card back to 3.39 t/h, run to 72 min and read the panel again. *Checkpoint:* 20, 75, 80 and 390 s.

Accept times within one 5 s scan and solids within 0.05 %. Say why AT-301 always comes last.

4.8 App challenge · 8 min

Use the default seed and a speed of 5, and read the true solids from the "true at AT-301" trace or the `product_solids_pct` column. The solids setpoint steps from 48 to 49 % at 5 min.

(a) Load scenario 2, Solids loop on live steam, and run to 60 min as shipped. Read how long after the step the true solids first reach 48.9 %, their peak, and the Feed processed tile. *Checkpoint:* About 9.3 min, 49.13 % and 20.0 t.

(b) Load scenario 3, The same loop on feed rate, set the Solids gain on feed to 0.16 t/h per % and the Solids integral time on feed to 10 min before Run, and run to 60 min. Read the same three values. *Checkpoint:* About 9.7 min, 49.27 % and 19.8 t.

Accept crossing times within 0.2 min and peaks within 0.02 %. Compare the loop gains of the two runs and say what the feed pairing costs.

4.9 App challenge · 8 min

Load the last scenario, Challenge, an eight-hour shift, and run the full 480 min at the default seed and a speed of 60. Read the scorecard tiles Worst solids error, Feed processed, Hottest product, Economy since reset and Dry-out at the end.

(a) Run the shift as shipped, with feed pairing, tight level loops and no feedforward. *Checkpoint:* 1.49 %, 155.5 t, 67.05 °C, 4.49 kg/kg, 13 s.

(b) Reload the scenario. Before Run, set Solids held by to Live steam, turn the Steam-to-feed ratio and the Feed-solids compensation On, and set the three Level gains to 0.6 % per % and the three Level integral times to 20 min. Run the shift again. *Checkpoint:* 0.62 %, 160.0 t, 69.35 °C, 4.40 kg/kg, 0 s.

Accept the tiles as displayed. Say which targets the first run misses and what the second run's higher steam buys.

4.10 Application area · 4 min

At the design point 3.39 t/h of motive steam gives up 2.51 MJ/kg in the chest, 0.18 t/h of preheat steam gives up 2.20 MJ/kg, effect 1's vapor condenses at 65.09 °C and the jet entrains 1.566 kg/kg. The train evaporates 16.25 t/h.

(a) Find the latent heat $\lambda(\theta_1)$ and the bound of Equation (4.10), and compare it with the economy. *Checkpoint:* 2.347 MJ/kg, 4.85 kg/kg against 4.56 kg/kg.

(b) Find the bound with no thermocompressor ($R_e = 0$) and the same heat per kilogram of steam, and the live steam that a train at that bound would need for the same water. *Checkpoint:* 3.19 kg/kg and about 5.1 t/h.

4.11 Application area · 4 min

A falling film needs at least $\Gamma_{min} = 0.08\,\mathrm{kg/(m\,s)}$ of liquid per meter of wetted perimeter, and the tube bundle of effect 3 has 17.5 m of perimeter. FY-800 keeps each film 15 % above the floor by recirculation.

(a) Find the flow effect 3 needs to stay wet, and its margin at the design inflow of 7.46 t/h. *Checkpoint:* 5.04 t/h and 1.48 times the floor.

(b) At the end of case 8 with recirculation on, FY-800 returns 1.12 t/h to the top of effect 3. Find the inflow FT-302 measures and the share of the floor the film would receive without recirculation. *Checkpoint:* About 4.68 t/h, about 0.93 of the floor.

4.12 Review the agent · 5 min

An engineer asks an AI assistant whether the solids loop can be slowed down for thinner milk, and it replies as follows. "The live-steam economy is 4.56 kg/kg. The thinner feed of case 4 needs 0.25 t/h

more water evaporated, which takes $0.25/4.56 = 0.055\,\mathrm{t/h}$ of steam. The compensated ratio already adds 0.052 t/h, within 6 %. Detune AIC-301 to 0.02 t/h per % and 20 min, and the feedforward will carry the disturbance.” Find the flaw and say how to check it.

CHAPTER 5

Roll-to-Roll Coating Simulator for Web Tension and Coat Weight During Speed Changes

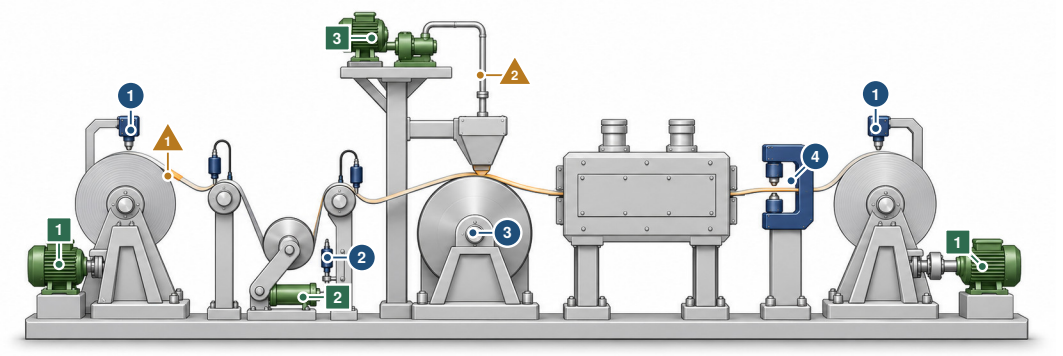

The controller moves 1 unwind and rewind motors 2 dancer air load 3 coating pump

The controller reads 1 roll diameter sensors 2 dancer position 3 line speed at the coating roll 4 coat-weight gauge

What disturbs it 1 splice to a web of another stiffness 2 coating batch of another viscosity

Notice that the die sits on the coating roll far upstream of the gauge, so a thinner coat is made long before it is weighed. Illustration, not to scale.

IN PLAIN TERMS **One Speed Change Spoils Two Things at Once**

The main idea. A coating line pulls a thin sheet of film or foil between two rolls and spreads a wet layer on it. The sheet must be pulled just tight enough, and the layer must weigh the right amount.

The challenge. A faster line spreads the same coating over more sheet, so the layer thins at once, while the scale that weighs it sits far downstream. The rolls at each end must spin faster, and they change size every turn.

Why it matters. Making the pump and the roll motors follow the speed, not the late readings, saves most of the sheet a speed change would waste.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the two histories in Section 5.2, the try-its and the facts of Section 5.6 and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the two open challenges

The control question A coating line sells the tension a web was wound at and its coat weight, and one change of line speed upsets both: the slot die thins the coating at once, and the winders must spin up rolls whose radius and inertia change every turn. Which part of the remedy belongs to feedforward and compensation, and which part can a trim on a late gauge still do?

Learning objectives

- Derive the coat weight, the pump capacity and the gauge delay of a pre-metered die from the flow, the width and the line speed.
- Compute the inertia of a winding roll and the torque, or dancer offset, that a speed ramp demands of it.
- Explain with a mass-on-a-spring model why a softer spliced web, not a stiffer one, destabilizes a hot tension trim.
- Show on the app what pump feedforward and inertia compensation each remove from the scrap of an order change.
- Judge which part of a predictive controller's advantage comes from its model of the pump lags and of a delay counted in meters.

Prerequisites Sections 1.4.3 (line speed as the throughput handle), 1.4.8 (disturbances passed along a line), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/web-coating/

See also Volume 3, Chapter 6 for a feedforward from the cause with the feedback demoted to a slow trim, and Volume 3, Chapter 12 for the feedforward of a known, fading load. Volume 5, Chapter 5 for a product quality fixed on a moving ribbon long before its gauge reads it, and Volume 5, Chapter 6 for material tracked past a blind instrument.

5.1 Introduction

Paper, plastic film, metal foil and textiles are made and processed as webs, sheets so much longer than they are wide that they travel from one roll to another through the machinery that works on them. The tension in each span between driven rolls is held by the drives at its two ends, and tension control on paper-making and converting machinery, models of tension variation in pliable material and the decentralized control of whole processing lines each have a literature of their own [1–3]. Coating adds a liquid layer to the moving web. A slot die is a pre-metered coater, which lays down everything its pump delivers, so the wet film is the flow per unit width divided by the web speed [4, 5]. The same machine makes adhesive films and labels, and it makes the cathodes and anodes of lithium-ion cells [6, 7], where the dry coat weight sets the capacity per unit area.

Both properties of the coated roll are ratios with the line speed in them, which is exactly what a block diagram with constant gains hides. At steady speed a winder's tension is its torque divided by its radius, the radius of the simulated unwind falls from 300 mm toward a 76 mm core, and any torque that accelerates the roll must also spin up an inertia that grows with the fourth power of that radius. The coat weight made at the die is the pump's flow divided by the width and the speed. On the simulated film line, raising the speed from 120 to 180 m/min with the pump flow held makes the coating 33 % thin at once, and the gauge 19 m downstream first sees it 7.07 s later. The same ramp swings the tension in the first span by $+12.1$ and -16.6 %, because the web itself has to accelerate a full unwind roll that behaves like 201 kg at its surface. The simulator corrects the belief that tension is a torque setting and coat weight a pump setting, each to be held by its own feedback loop. Both follow the speed through a known law, so the lasting remedies are feedforward and compensation, and feedback is left to trim what they miss. The coagulant dosing of Volume 3, Chapter 6 divides the work between a feedforward and a slow trim in the same way.

The class of problem reaches past coating. A winder is a plant whose gain and inertia change continuously and predictably while it runs. Models that carry the varying radius and inertia underlie the robust, parameter-varying and decentralized designs proposed for web lines [3, 8], and adaptive

proportional-integral (PI) tension control has been implemented on roll-to-roll lines [9]. A quality variable that is set upstream as a flow divided by a throughput and measured downstream inherits a transport delay equal to a distance divided by a speed, so its feedback loop changes character whenever the line changes speed. The control and optimization of coating lines are treated as a subject of their own in the coating literature [10], and even the choice of instrument that closes the tension loop, a dancer roll or a load cell, has been argued out on coating machinery [11].

However, the two halves of the problem are usually taught apart. Tension appears, if at all, as a mass on a spring, and coat weight as a first-order loop with a dead time of a given number of seconds. Neither view shows that one speed ramp disturbs both, that the same gauge is 9.49 s late at 120 m/min and 38.0 s late at 30 m/min, or that a torque-driven tension loop is easier to hold on a stiff web than on a soft one. This chapter traces one speed change along the two physical paths by which it reaches a coated roll, on a simulated line where feedforward from the speed, compensation of the changing radius and inertia, and feedback on a late gauge can each be switched off and priced in meters of scrap. The line speed is the throughput handle of the line in the sense of Section 1.4.3, and the pump and both winders follow it. The tension transport model, radius and inertia compensation, dancers, load cells and speed feedforward are all established practice.

Section 5.2 traces how tension and coat weight came to be controlled on converting lines. Section 5.3 states the model, a line with 24 states and a web map that carries every 5 cm of product to the rewind, in enough detail to rebuild it. Section 5.4 gives the speed master, the winder torques and the pump with their feedback trims, and an optional model predictive controller of the rewind torque and the pump [12]. Section 5.5 derives three reduced models and checks them against the full nonlinear engine: a third-order dancer loop that reproduces the dancer travel and the span tension through a ramp to within 0.7 mm and 0.1 N, a mass-on-a-spring model of the rewind that places the stability limit of the hot load-cell trim of case 4 at 0.80 times the stiffness of the film once the 5 ms control interval is counted, and a delay loop whose critical integral time of 6.4 s at 30 m/min the engine reproduces. Section 5.6 measures nine scenarios and a graded challenge, and sets the predictive controller against the loops case by case. Section 5.7 describes the real sensors, drives and controllers, Section 5.8 states the limitations and Section 5.9 concludes.

5.2 Tension and Coat Weight on Converting Lines

A web line is built from driven sections, and the draw between two of them, the difference of their surface speeds, is what stretches the web in between. An unwind held back by a brake needs a braking torque that falls as the roll runs down, because the torque that gives a fixed tension is proportional to the radius. Tension control on paper-making and converting machinery was being treated as an engineering subject in its own right by 1964 [1], and models of the tension variations in pliable material passing through production machinery followed [2].

Two instruments came to close the loop. A dancer roll rides in a loop of web on a weighted or air-loaded arm, so that its load sets the tension and its position reports the mismatch between the speeds on either side, while a load cell under an idler reads the tension directly. The first history segment follows the choice between them and the estimators that replace both.

> ⧖ HISTORY **1993: Dancer Roll or Load Cell**
>
> Whether paper machinery should close its tension loop on a dancer roll or on load cells was still a question in 1993. A study of that year answered it with a computer model of an existing off-machine coater that carried both systems, analyzed each for stability and dynamic response, and set the results beside measurements made on site [11]. A later answer drops the sensor. For a two-drum paper winder, two tension estimators were built from the unwinding and the rewinding variables, with the dynamics, the friction and the changing inertia taken into account, because load cells are expensive and need frequent maintenance and calibration. Through typical winder cycles the estimators gave stable and precise tension control without a load-cell loop [13].

The winders make the problem hard, because the radius and inertia of the rolls vary widely during winding and unwinding, which changes the plant within every roll [8]. The second history segment traces the designs that decouple tension from speed on such a plant, beyond the decentralized proportional-integral-derivative (PID) loops common in industry, and adaptive PI schemes were implemented for roll-to-roll manufacturing [9]. The structure simulated here, a speed master with torque-driven winders under decentralized PI trims, is the kind of industrial baseline those studies compare against.

> ⧖ HISTORY **2002 and 2007: Decoupling Tension From Speed**
>
> On a winding line, a web break or a fold is prevented by decoupling the web tension from the web velocity, and the wide variation of the rolls' radius and inertia changes the dynamics considerably while they wind. A 2002 study compared a robust controller that reduced the coupling between tension and velocity with the decentralized PID loops common in industry, gave it gains that vary with the radius, and combined it with smoothly scheduled controllers on an experimental winding system [8]. A 2007 design divided a whole processing line into tension zones, each set from its reference tension and from a master speed roller that fixes the transport speed, and tested its decentralized controller against an often-used industrial PI controller on a large experimental platform [3]. In a study of 2023, decentralized PID controllers were still the mainstream in the roll-to-roll industry [14].

Coating science followed a parallel path. A pre-metered coater applies everything it is fed, so its wet thickness is fixed by the flow and the speed, and what the flow in the bead must do is stay stable. The limiting flow of a pre-metered coating device was analyzed in 1976 [4], and the operating limits of slot die coating, from the low-flow limit to dynamic wetting failure and the highest coating speed for a given flow, have since been reviewed as a body of work [5]. Control and optimization of coating processes have a chapter of their own in an edited reference work on liquid film coating [10].

Battery electrodes gave the same machinery a new product. The web is an aluminum or copper current collector, the coating is a thick slurry of active material, and the coating window of such slurries in a slot die has been characterized [6]. Metrology for electrode slurry coating, including its integration into the line, has been reviewed [15], and electrode processing has been costed in enough detail to show that three processing changes, among them replacing the N-methylpyrrolidone solvent, could together lower the pack cost by 26.4 % [16]. Automotive cell production has turned these lines into a manufacturing concern at scale [7]. On an electrode line a coat weight a few percent light is capacity lost, and a torn foil stops the line for a rethread.

5.3 Model of the Coating Line

The simulated line carries a web 1.0 m wide from an unwind to a rewind through two spans. Span 1 runs 4 m from the unwind over a dancer to the coating roll, which is the speed master of the line and carries the slot die. Span 2 runs 24 m from the coating roll through a 15 m dryer that starts 2 m after the die, past the coat-weight gauge 19 m after the die, to the rewind. Two products run on it, a packaging film of polyethylene terephthalate (PET) 50 µm thick carrying a water-based adhesive, and a battery electrode of aluminum foil 15 µm thick carrying a cathode slurry of nickel manganese cobalt oxide (NMC) in N-methylpyrrolidone. The arrangement and every symbol used below are shown in Figure 5.1. The controlled (CV), manipulated (MV) and disturbance (DV) variables and the states are classified in Table 5.1.

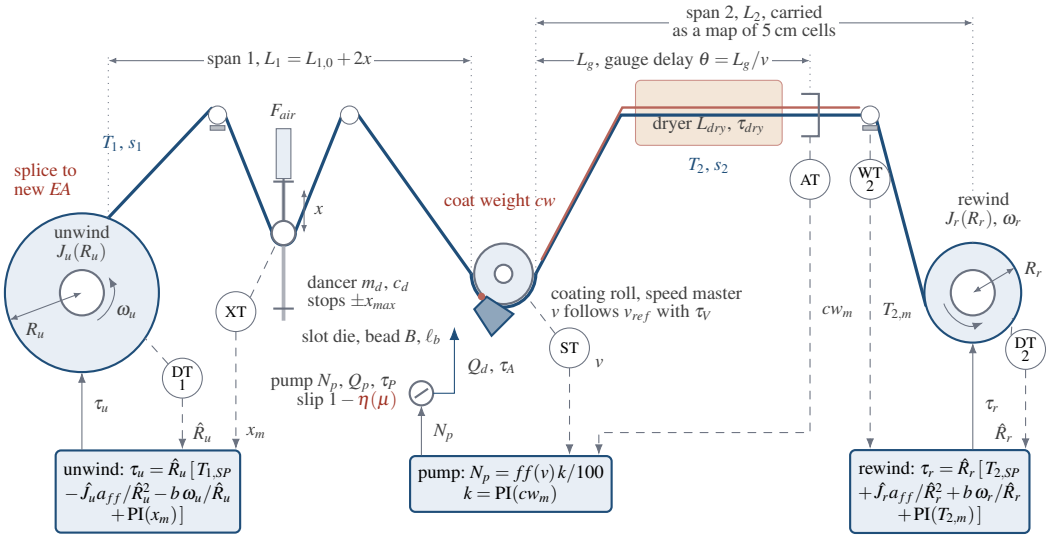

Figure 5.1 Roll-to-roll coating line as modeled. The unwind (R_u, J_u, ω_u) is held back by τ_u and can be spliced to a web of another EA. Span 1 (s_1, T_1) wraps a dancer of mass m_d and damping c_d under the air load F_{air}, whose travel x between stops at $\pm x_{max}$ makes the span $L_{1,0} + 2x$ long. The coating roll is the speed master at v, and its slot die, fed by a gear pump at speed N_p whose slip grows as the viscosity μ falls, holds a bead B of length ℓ_b on the web. Span 2 (s_2, T_2) carries the coat weight cw through the dryer to the gauge AT at L_g and the load cell WT2 and winds onto the rewind (R_r, J_r, ω_r) under τ_r. DT1 and DT2 are ultrasonic diameter sensors, XT the dancer sensor and ST the speed encoder, and the three controllers set τ_u, N_p and τ_r.

Table 5.1 Variables of the Film Line at Its Reset State

Symbol	Role	Nominal	Range
T_1	CV, span 1 tension, set by the dancer	200 N	SP 20–600
T_2	CV, span 2 tension, load cell	150 N	SP 20–600
cw_g	CV, dry coat weight at the gauge	12.0 g/m^2	SP 2–400
x	CV of the dancer loop, position	0 mm	±150
τ_u	MV, unwind holdback torque	59.87 N m	±120
τ_r	MV, rewind torque	18.33 N m	±120
N_p	MV, coating pump speed	52.91 %	0–100
v_{SP}	MV, line speed order	120 m/min	0–300
μ	DV, coating viscosity, % of recipe	100	25–300
f_{EA}	DV, next roll stiffness, $\times EA$	1.00	0.2–3
σ_g	DV, gauge noise (1σ), % of target	0.5	0–3
σ_{lc}	DV, load-cell noise (1σ)	1 N	0–10
—	DV, order, stop, restart, splice, batch	—	events
ω_u, R_u	states, unwind	6.661 rad/s, 300 mm	—
ω_r, R_r	states, rewind	16.663 rad/s, 120 mm	—
x, $\dot{x}$, F_{air}	states, dancer	0, 0, 400 N	—
s_1, s_2	states, span stretch	3.552, 15.989 mm	—
v, ξ	states, line speed, web coordinate	2 m/s, 0	—
Q_p, Q_d	states, pump and die flow	3.175, 3.048 L/min	—
B	state, bead per unit width	0.127 mm^2	—
—	states, three flags, six balances, drive energy	—	—

5.3.1 Rolls That Change Size

Each winder loses or gains one layer of web per turn, and its inertia grows with the fourth power of its radius, as shown in Equations (5.1) and (5.2).

$$\frac{\mathrm{d}R_u}{\mathrm{d}t} = -\frac{h\,\omega_u}{2\pi}, \qquad \frac{\mathrm{d}R_r}{\mathrm{d}t} = \frac{h_c\,\omega_r}{2\pi} \tag{5.1}$$

$$J(R) = J_{core} + \frac{\pi}{2}\,\rho\,W\left(R^4 - R_{core}^4\right) \tag{5.2}$$

where h is the web thickness, W the width, ρ the web density, J_{core} the inertia of the core, shaft and motor, and R_{core} the core radius. The rewind builds with the coated web, of thickness $h_c = h + cw_{SP}/\rho_{dry}$ and density $\rho_c = (\rho h + cw_{SP})/h_c$, which is 60.0 µm and 1358 kg/m^3 for the film. A full 300 mm film roll carries 18.11 kg m^2 against 0.50 kg m^2 for the empty core, and it behaves like $J/R^2 = 201$ kg at its surface, although the film on it weighs 368 kg. A full coil of foil carries 34.71 kg m^2, or 386 kg at its surface, and weighs 714 kg. The web leaves or arrives at the surface speed, so the angular momentum it carries is exactly $\omega\,\mathrm{d}J/\mathrm{d}t$ and the torque balances are written for $J\dot{\omega}$ rather than for $\mathrm{d}(J\omega)/\mathrm{d}t$, as shown in Equation (5.3).

$$J_u\dot{\omega}_u = T_1 R_u - \tau_u - b\,\omega_u, \qquad J_r\dot{\omega}_r = \tau_r - T_2 R_r - b\,\omega_r \tag{5.3}$$

where b is the bearing friction. At the unwind the web pulls the roll round and the motor holds it back, while at the rewind the motor drives and the web resists. At steady speed the holdback torque is $T_1 R_u - b\,\omega_u$, which is 59.87 N m for 200 N on the full film roll.

5.3.2 Spans and the Dancer

The state of each span is its stretch s_i, the path length minus the unstretched length of web in it. Conservation of the unstretched web gives Equation (5.4).

$$\begin{aligned} \dot{s}_1 &= 2\dot{x} + \frac{v}{1+\varepsilon_1} - v_0, \qquad \dot{s}_2 = \frac{v_2}{1+\varepsilon_2} - \frac{v}{1+\varepsilon_1} \\ T_i &= \begin{cases} \max[0,\ (s_i + \tau_w \dot{s}_i)/C_i] & s_i > 0 \\ 0 & s_i \le 0 \end{cases} \end{aligned} \tag{5.4}$$

where $v_0 = \omega_u R_u$ and $v_2 = \omega_r R_r$ are the surface speeds of the winders, v is the speed of the coating roll, C_i is the compliance of the span, the integral of $1/EA$ over the unstretched web in it, τ_w is a retardation time that lumps web, air and idler damping, and $\varepsilon_i = s_i/(C_i EA)$ is the strain of the web leaving span i, evaluated with the EA of the web at that roller. For one web, $C_i = (L_i - s_i)/EA$, and linearizing gives the standard transport equation $\dot{T}_i = (EA/L_i)(v_i - v_{i-1}) + (T_{i-1}v_{i-1} - T_i v_i)/L_i$. A negative stretch is slack web, which carries no tension and must be taken up before tension can build again. After a splice a span holds two webs, and the controller's record of the web supplies C_i as the sum of the two parts. The film has $EA = 225$ kN, so span 2 is a spring of $EA/L_2 = 9.4$ kN/m, and the foil has 1050 kN.

The web wraps the dancer twice, so the dancer lengthens span 1 by $2x$ and its air load balances twice the tension, as shown in Equation (5.5).

$$\begin{aligned} m_d\ddot{x} &= F_{air} - 2T_1 - c_d\dot{x}, \qquad |x| \le x_{max} \\ \tau_{air}\dot{F}_{air} &= 2T_{1,SP} - F_{air} \end{aligned} \tag{5.5}$$

At a stop the dancer lands inelastically. Its velocity is set to zero and held there while it is pushed into the stop, and any travel past the stop within one step is taken from the path, so the stretch loses twice the overshoot. On a stop, span 1 is a fixed length of web between two rolls turning at different speeds.

5.3.3 Coating Roll, Pump, Die and Bead

The coating roll is a stiff speed-regulated drive, and the web coordinate at the die, ξ, counts the unstretched web that has passed it, as shown in Equation (5.6).

$$\tau_V \dot{v} = v_{ref} - v, \qquad \dot{\xi} = \frac{v}{1+\varepsilon_1} \tag{5.6}$$

A gear pump follows its speed command through a lag, loses part of what it displaces to leakage past its gears, and feeds the die through a hose and manifold that act as a second lag, as shown in Equation (5.7).

$$\begin{gathered} \tau_P \dot{Q}_p = \frac{N_p}{100} Q_{max} - Q_p, \qquad \tau_A \dot{Q}_d = \eta\, Q_p - Q_d \\ \eta = 1 - \beta \frac{100}{\max(10,\ \mu)} \end{gathered} \tag{5.7}$$

where Q_p is the displaced flow, Q_d the flow leaving the die, β the slip at the recipe viscosity and μ the viscosity in percent of the recipe. A batch at half the viscosity therefore doubles the slip from 4 to 8 %. The die feeds a small bead between its lips and the web, and the moving web drags the bead away, as shown in Equation (5.8).

$$\frac{\mathrm{d}B}{\mathrm{d}t} = \frac{Q_d}{W} - \frac{v B}{\ell_b}, \qquad h_w = \frac{B}{\ell_b}, \qquad cw_{die} = \varphi\, \rho_{wet}\, h_w \tag{5.8}$$

where B is the bead volume per unit width, ℓ_b its length, h_w the wet film laid down, φ the solids fraction and ρ_{wet} the density of the wet coating. At steady state $h_w = Q_d/(Wv)$, the pre-metered law, which is 25.4 µm wet for 12 g/m^2 of adhesive and 124.2 µm for 200 g/m^2 of cathode. At a standstill the bead grows instead of dividing by zero, and above $h_{bead}\,\ell_b$ it spills. When the interlock lifts the die, its flow and the bead go to the drip tray. The bead time constant ℓ_b/v is 2.5 ms at 120 m/min, so the bead matters only near a standstill.

5.3.4 The Web Map and the Gauge

The controller keeps a record of the web in 5 cm cells indexed by ξ. Each cell carries the dry coat weight it received at the die, averaged over its length, the target it was coated to, the time it has spent in the dryer and a set of defect flags, and it keeps all of them until it is wound. The gauge reads the cell now under it, as shown in Equation (5.9).

$$\begin{gathered} cw_g(t) = cw_{die}\left(\xi(t) - L_g \frac{L_{u2}}{L_2}\right) \\ \theta = \frac{L_g\, L_{u2}}{L_2\, v} \approx \frac{L_g}{v} \end{gathered} \tag{5.9}$$

where $L_{u2} = L_2 - s_2$ is the unstretched web in span 2 and the argument is the coordinate of the cell at the gauge. The delay is 9.49 s at 120 m/min, 38.0 s at 30 m/min and 47.5 s for the foil at 24 m/min. It is fixed in meters, so a change of speed changes it at once, and during a ramp the gauge reads web that was coated at whatever speed it had when it passed the die. The coating-weight gauge of Volume 16, Chapter 10 stands far downstream of its air knives, and its delay moves with the line speed in the same way.

A cell is scrapped at the rewind if it carries a flag, if its coat weight is more than 3 % from the target it was coated to, or if its residual solvent $[(1-\varphi)/\varphi]\exp(-t_{dry}/\tau_{dry})$ exceeds 1 % of the dry coat. A tension excursion beyond ±10 % of its setpoint flags every cell in that span at that instant, because the tension in a span is uniform and all of that web lived through it. Wrinkles below 40 % of the setpoint, a spill, a break, bare web and the half meter on either side of a splice are flagged the same way. The dryer never limits the scenarios, because a film cell needs 2.88 s and a foil cell 11.27 s in the 15 m dryer, which caps the film at 312 m/min and the foil at 80 m/min.

Table 5.2 Line Parameters. Hidden Constants Are Never Displayed, and All Others Can Be Changed in the Model Panel.

Symbol	Meaning	Value	Class
W	web width	1.0 m	adjustable
E/E_0	web stiffness, times the product's	1.00	adjustable
$L_{1,0}$	span 1 path, dancer centered	4 m	adjustable
L_2	span 2 path, coating roll to rewind	24 m	adjustable
L_g	coating roll to gauge	19 m	adjustable
—	coating roll to dryer entry	2 m	hidden
L_{dry}	dryer length	15 m	adjustable
R_{core}	core radius	76 mm	hidden
J_{core}	core, shaft and motor inertia	0.50 kg m^2	adjustable
—	rewind radius at transfer	350 mm	adjustable
b	bearing friction	0.02 N m s	adjustable
τ_{max}	winder torque limit	120 N m	adjustable
τ_V	coating roll speed response	50 ms	adjustable
—	emergency stop deceleration	60 m/min per s	adjustable
—	coating roll radius, inertia (drive power only)	100 mm, 0.5 kg m^2	hidden
m_d, c_d	dancer mass, damping	8 kg, 1000 N s/m	adjustable
x_{max}	dancer travel to each stop	150 mm	adjustable
τ_{air}	dancer loading regulator lag	0.5 s	adjustable
τ_w	span retardation time	5 ms	adjustable
Q_{max}	pump displacement at full speed	6.0 L/min	adjustable
β	pump slip at recipe viscosity	4 %	adjustable
τ_P, τ_A	pump lag, hose and manifold lag	0.3, 0.5 s	adjustable
ℓ_b	coating bead length	5 mm	adjustable
h_{bead}	bead height before it spills	300 µm	adjustable
v_{min}	coating interlock speed	5 m/min	adjustable
—	tension, coat weight bands	±10, ±3 %	adjustable
—	wrinkle threshold, solvent limit	40, 1 %	adjustable

5.3.5 Breaks, Splices and Instruments

The web breaks when either tension exceeds T_{break}, and the line then stops at 60 m/min per s with the pump off. A splice replaces the unwind with a full roll at the same surface speed, either on command or when the unwind is within 5 mm of its core, and it can bring a web whose EA is the product's times the next-roll factor. The rewind transfers to an empty core at 350 mm. The load cells add first-order-filtered Gaussian noise of σ_{lc} with a 20 ms correlation time, the dancer sensor 0.2 mm with 20 ms, and the ultrasonic diameter sensors 0.3 mm with 0.1 s before a 1 s average. The gauge adds noise of σ_g times the target with a 0.5 s correlation time and averages through a 1 s lag. The line speed and the roll speeds come from encoders without noise.

5.3.6 Initialization and Numerics

Table 5.2 gathers the line's constants and Table 5.3 the two recipes. Every scenario opens on the recipe's steady state. The stretch of each span is $s_i = T_{i,SP}L_i/(EA + T_{i,SP})$, the winders turn at $v_0 = v/(1+\varepsilon_1)$ and $v_2 = v(1+\varepsilon_2)/(1+\varepsilon_1)$, the die flow is Wvh_w, the pump flow Wvh_w/η and the bead $h_w\ell_b$, and the unwind is full at 300 mm while the rewind holds 120 mm unless the scenario sets either. The web map is filled with coated web at the target from the rewind to the die and with bare web from the die to the unwind. The application checks that the derivatives of the eleven winder, dancer, span, drive and coating states are below 10^{-9} there, leaving only the two radii, the web coordinate and the running balances in motion. The equations advance in fixed 1 ms steps of fourth-order Runge-Kutta integration, the loops act on a 5 ms cycle, and a historian samples the line at 20 Hz. Halving the physics step while retaining the control interval gives the same splice break

Table 5.3 The Two Products. A Product Is Loaded With Its Recipe on a Freshly Threaded Line.

Quantity	Packaging film	Battery electrode
web, coating	50 µm PET, adhesive	15 µm Al, NMC cathode
E, ρ	4.5 GPa, 1390 kg/m^3	70 GPa, 2700 kg/m^3
full roll radius	300 mm	300 mm
$T_{1,SP}$, $T_{2,SP}$	200, 150 N	120, 100 N
T_{break}	800 N	360 N
cw_{SP}	12 g/m^2	200 g/m^2
φ, ρ_{wet}, ρ_{dry}	0.45, 1050, 1200 kg/m^3	0.70, 2300, 2400 kg/m^3
τ_{dry}	0.6 s	3.0 s
line speed	120 m/min	24 m/min
trim integral time	6 s	30 s

time of 37.848 s to 1 ms precision. Seeding, replay and the playback speed, here 0.5× to 20×, work as Section 1.6.1 describes.

5.3.7 Assumptions

Each span is one spring with uniform tension. A tension wave crosses the 24 m span in 13 ms in the PET and 5 ms in the foil, at wave speeds $\sqrt{E/\rho}$ of 1.8 and 5.1 km/s, which is far faster than anything the controllers do. The idealization leaves out the drag and inertia of the idlers in a long dryer.

Span damping is a retardation time. Writing it as $\tau_w \dot{s}_i / C_i$ makes the damping proportional to the stiffness, and Section 5.5 shows that this proportionality sets the stability limit of the load-cell trim. Damping from air and idler bearings need not scale with the web.

The web never slips, the drives are ideal and the rolls are rigid. The web leaves each roller with the strain it had in the span before it, although a driven roller slips once the tension ratio across it exceeds its capstan limit. The winders are torque drives with ultrasonic diameter sensors and the coating roll a stiff speed master. Wound-in tension, telescoping, crushed cores and taper tension are absent, and the rewind builds one coated thickness per turn.

The coating is pre-metered through one bead. The coat weight follows $Q/(Wv)$ at every speed, so ribbing, air entrainment, bead break-up and the low-flow limit that bound a real slot die never occur [4–6].

Viscosity acts only through pump slip. In a real die circuit the discharge pressure also scales with viscosity, so the viscosity slider stands for any unmeasured change in the delivered flow.

The dryer is one time constant and the gauge sees dry solids only. Residual solvent is scored at the rewind but never reaches the gauge reading.

The break and wrinkle limits are illustrative. They are four and three times the span 1 setpoint for the film and the foil, and 40 % of setpoint, because real webs let go at edge nicks, pinholes and weak splices far below their intact strength.

5.4 Implemented Controllers

Every loop reads a measurement rather than the true state, runs every 5 ms, and is configured as listed in Table 5.4.

5.4.1 Discrete Algorithm

All three loops run the shared velocity-form PID routine of Section 1.6.2, Equations (1.11) to (1.14), with the tracking time $T_t = \tau_I$. Two details are particular to this line. In the winder loops the output bias u_b is not a constant. It carries the feedforward of Equation (5.11) below and is updated at every execution, so the trims add to a feedforward that moves with the radius and the acceleration. The coat-weight trim instead multiplies the pump feedforward of Equation (5.12). The derivative time is zero in all three loops and none has a rate limit, because the drives, the pump and the hose already

Table 5.4 The Three Feedback Loops. Each PI Works in the Units of Its Feedforward, and Each Executes Every 5 ms.

Loop	CV	MV	K_c	τ_I
dancer	x_m, SP 0	$\tau_u/\hat{R}_u$	0.4 N/mm	5 s
load cell	$T_{2,m}$	$\tau_r/\hat{R}_r$	0.5 N/N	2 s
coat weight	$100\,cw_m/cw_{SP}$	trim k	0.3 %/%	6 or 30 s

A long span raises the holdback, a low tension raises the winding torque and a light coat raises the trim. The winder outputs are limited to $\pm\tau_{max}/\hat{R}$, which is ±120 N m of torque. The trim is limited to 50 to 150 %, with its upper limit lowered to $10^4/ff$ whenever that keeps N_p at or below 100 %.

supply the lags. In manual all three PIs track the operator's output, so the return to automatic is bumpless.

5.4.2 Speed Master

The line speed follows a ramped reference, and the acceleration handed to the winders is that ramp passed through the drive's own response, as shown in Equation (5.10).

$$\begin{aligned} v_{ref,k} &= v_{ref,k-1} + \mathrm{clamp}\left(v_{tgt} - v_{ref,k-1},\ \pm a\Delta t\right) \\ a_{ref,k} &= \left(v_{ref,k} - v_{ref,k-1}\right)/\Delta t \\ a_{ff,k} &= a_{ref,k}\left(1 - \mathrm{e}^{-\Delta t/\tau_V}\right) + a_{ff,k-1}\,\mathrm{e}^{-\Delta t/\tau_V} \end{aligned} \tag{5.10}$$

where the ramp rate a defaults to 6 m/min per s and the target v_{tgt} is the order v_{SP}. With the capacity limit on, $v_{tgt} = \min(v_{SP},\ 0.95\,\eta_{nom}Q_{max}/(Wh_{w,SP}))$, which is 215.5 m/min for the film and 44.1 m/min for the foil at the nominal slip. The filter matters. Fed the raw ramp acceleration instead, in a test made for this chapter with a wrapper around the unchanged controller, the suggested configuration of the challenge swings span 2 from −27.5 to +22.1 % and scraps 26.50 m instead of 2.55 m, because the feedforward torque then arrives about 50 ms before the line accelerates and excites the lightly damped rewind on its span.

5.4.3 Winder Torques

Both winders apply a feedforward torque from the tension setpoint, the estimated radius and the planned acceleration, trimmed by a PI in tension units, as shown in Equation (5.11).

$$\begin{aligned} \tau_u &= \hat{R}_u\left[T_{1,SP} - \frac{\hat{J}_u a_{ff}}{\hat{R}_u^2} - \frac{b\,\omega_u}{\hat{R}_u} + \mathrm{PI}_u(x_m)\right] \\ \tau_r &= \hat{R}_r\left[T_{2,SP} + \frac{\hat{J}_r a_{ff}}{\hat{R}_r^2} + \frac{b\,\omega_r}{\hat{R}_r} + \mathrm{PI}_r(e_r)\right] \end{aligned} \tag{5.11}$$

where $\hat{J} = J(\hat{R})$ from Equation (5.2), $e_r = T_{2,SP} - T_{2,m}$, and the bracketed feedforward enters each PI as its bias. Because the trims work in newtons and are multiplied by the radius, their gains need no rescaling as a roll grows or empties. The unwind's PI keeps the dancer mid-stroke, since the air load already sets the tension, and the rewind's PI corrects the tension read by its load cell. Multiplying by $\hat{R}$ is radius compensation. The *fixed torque* option replaces $\hat{R}$ by the radius at the moment it is selected, and the load-cell trim can be switched off, leaving the rewind on its feedforward alone. Inertia compensation is the $\hat{J}a_{ff}$ term and can also be switched off.

The radius comes from an ultrasonic diameter sensor on each winder, preset from the new roll at a splice or a transfer. A radius estimated instead from the ratio of line speed to unwind speed would read $R_u(1 - 2\dot{x}/v_0)$ on this line, because the web leaving the unwind moves at $v + 2\dot{x}$ while the dancer travels. Section 5.5 shows that this error feeds the dancer's motion back into the holdback torque with a destabilizing sign.

5.4.4 Coat Weight

The pump speed is a feedforward from the measured line speed times a trim, a ratio of pump to line speed, as shown in Equation (5.12).

$$\begin{aligned} N_p &= \frac{ff\,k}{100}, \qquad k = \mathrm{PI}_{cw}\left(\frac{100\,cw_m}{cw_{SP}}\right) \\ ff &= \frac{100\,W\,v\,h_{w,SP}}{\eta_{nom}\,Q_{max}} \end{aligned} \tag{5.12}$$

where $h_{w,SP} = cw_{SP}/(\varphi\rho_{wet})$ and $\eta_{nom} = 0.96$ is the efficiency at the recipe viscosity. The feedforward knows the recipe and nothing about the coating actually in the pump. With the speed feedforward switched off, ff is held at its last value. The trim's upper limit moves with the feedforward so that N_p never asks for more than 100 %, which keeps the integral from winding up while the pump is saturated.

Below $v_{min} = 5\,\mathrm{m/min}$ the coating interlock lifts the die, stops the pump and makes the trim track its held value, and on the restart it holds the trim until web coated 2 m after the die re-engaged reaches the gauge. With the interlock off, the die stays on the web whatever the speed. After a break the winders brake their rolls to rest in about 0.5 s and the pump stops. In manual the operator sets τ_u, τ_r and N_p and all three PIs track.

5.4.5 Simulation, Control and Controller Models

Equations (5.1) to (5.9), evaluated with the true radii, strains, splices and viscosity, form the simulation model. The loops' internal model is the feedforward of Equations (5.10) to (5.12), which uses the sensed radii, the recipe density and slip, and the drive's 50 ms response, and never the true viscosity or the true EA of a spliced web. The control models of Section 5.5 are linearizations made for analysis, and the application does not identify or use them. The predictive controller of Section 5.4.6 carries a model of its own.

5.4.6 Model Predictive Control of the Rewind and the Pump

A structure switch hands the rewind torque and the coating pump either to the loops above or to a model predictive controller (MPC), set up in the general form of Section 1.6.3. The loops stay the default everywhere except in case 9. In automatic the MPC decides the force the rewind torque puts on the web, $F = \tau_r/\hat{R}_r$, and the pump speed N_p together every $T_s = 0.05\,\mathrm{s}$, which is ten control intervals. Its horizon of $N = 60$ steps (3 s) is cut into eight move blocks of 1, 1, 2, 3, 5, 8, 13 and 27 steps, so each quadratic program carries 16 moves and one slack, and each plan's opening move stands until the next plan [12]. The unwind, the dancer and the speed master keep their loops. Horizons of 1.5 and 6 s leave the scrap of case 9, case 2 and the challenge unchanged, and so does 5 s ($N = 100$), at a mean of 228 µs a plan in Node against 148 µs for 3 s in the same run on a shared host.

The controller's model is written from Equations (5.3) to (5.8) with the declared parameter defaults and the recipe, frozen when the application loads, and never with the run's parameters. Its states are the coating roll speed v_1, the rewind surface speed relative to that roll z, the span 2 stretch s_2, and the pump and die flows q_p and q_d in percent of Q_{max}, as shown in Equation (5.13).

$$\begin{aligned} \tau_V\dot{v}_1 &= v_{ref} - v_1, \qquad \hat{T}_2 = ks_2 + cz \\ \hat{m}\dot{z} &= F + d - \hat{T}_2 - \frac{b}{\hat{R}_r^2}(v_1 + z) - \delta_J\hat{m}\dot{v}_1 \\ \dot{s}_2 &= \left(1 - \frac{T_{2,SP}}{EA}\right)z - \frac{\bar{v}}{L_2}s_2 \\ &\quad + \frac{T_{1,SP} - T_{2,SP}}{EA}v_1 + \frac{\bar{v}\,T_{2,SP}}{EA} \\ \tau_P\dot{q}_p &= N_p - q_p, \qquad \tau_A\dot{q}_d = q_p - q_d \end{aligned} \tag{5.13}$$

where $\hat{m} = J(\hat{R}_r)/\hat{R}_r^2$ from Equation (5.2) with the coated density, $k = EA/L_2$, $c = \tau_w k$, $\bar{v}$ is the measured speed at which the transport term is linearized, d is an estimated force on the rewind, and δ_J is 0 with inertia compensation on and 1 with it off. On the film $\hat{m}$ is 60.5 kg at 120 mm, k is 9375 N/m and c is 46.9 N s/m, so the rewind rings at 1.98 Hz with a damping ratio of 0.035 at 120 m/min, while on the foil $\hat{m}$ is 81.1 kg and k is 43 750 N/m. The model is discretized exactly with a zero-order hold at T_s and rebuilt when the measured radius has moved 1 % or the speed 10 %. Against steps of ± 1 N m and ± 5 % of pump on the engine, with the noise off, on the film at 120 and 30 m/min and the foil at 24 m/min, it leaves 1.5 to 8.7 % of the rms span 2 response and 0.12 to 0.58 % of the die's unexplained. The preview $v_{ref,i}$ at step i is the speed master's own ramp from its present reference to its target at the ordered rate, taken at the middle of each step. With the preview off the present reference is held.

Two estimators run every 5 ms under either structure, so a switch to the MPC starts warm. A steady-state Kalman filter on $[z, s_2, d]$ reads the load cell with a design variance of 1 N^2 and lets d drift as a random walk of 1×10^{-5} N^2 per scan, the integrating disturbance that removes offset under a mismatched model [17]. It estimates 63 % of an unmodeled 1 N force in 1.635 s on the film and 1.620 s on the foil. That variance is the fastest of the four tried, 4×10^{-4} to 4×10^{-6}, that held a splice to every next-roll stiffness from 0.2 to 3 times the product's, while 4×10^{-5} let span 2 swing from −67 to +71 % on the softest web and 4×10^{-4} broke it. Kalman filtering of the tension sensor, combined with a disturbance observer, has improved tension control on an industrial roll-to-roll system [18]. With the load-cell trim switched off the load cell counts as out of service and the filter runs on its model alone.

The coat weight is estimated across the delay. The controller integrates the encoder speed into a web odometer $\hat{\xi}$, records the coat weight its model laid at the die on every 2 cm of web, and reads that record $L_g = 19$ m back through the gauge's 1 s average. The ratio of the gauge reading to this model estimates the gain $g = \eta/\eta_{nom}$ of Equation (5.20), the pump efficiency relative to the recipe's, as shown in Equation (5.14).

$$\tau_g \dot{\widehat{cw}}_g = \widehat{cw}_{die}(\hat{\xi} - L_g) - \widehat{cw}_g, \qquad \tau_{est}\dot{\hat{g}} = \frac{cw_m}{\widehat{cw}_g} - \hat{g} \tag{5.14}$$

where $\widehat{cw}_{die} = \varphi \rho_{wet} \eta_{nom} Q_{max} q_d/(100 W v_1)$ is the model's coat weight at the die, cw_m the gauge reading and $\tau_{est} = 10$ s. The estimate $\hat{g}$ is clamped to 0.5 to 2 and held while the interlock holds the trim, while the line runs below v_{min} or the web under the gauge was coated with the die lifted or below it, and after a break. Because the delay is counted in meters it follows the speed at every speed and stops with the line. Through pump steps the model knows, $\hat{g}$ stays within 0.07 % of 1 at 10, 30 and 120 m/min, where the gauge read against what the die lays at the same instant would be up to 15.8 % off, and within 0.0101 % through the stop and restart of case 8, during which the odometer does not move. Delays defined implicitly by an integral of a transport rate, as this one is, have been treated with prediction-based feedback where the rate is itself the control input [19]. The plan predicts the die, because the web between the die and the gauge is already coated, and needs no chain of delay states, which at T_s would number 760 at 30 m/min.

The plan minimizes the objective of Equation (5.15).

$$\begin{aligned} J = &\sum_{i=1}^{N} \left[w_T\, e_{T,i}^2 + w_C\, e_{C,i}^2 \right] + \gamma_2\, \nu^2 + \gamma_1\, \nu \\ &+ \sum_j \left[\kappa_F \lambda_T\, \Delta F_j^2 + \kappa_P \lambda_P\, \Delta N_{p,j}^2 \right] \end{aligned} \tag{5.15}$$

where $e_{T,i} = 100(\hat{T}_{2,i} - T_{2,SP})/T_{2,SP}$ and $e_{C,i} = 100(q_{d,i} - q_i^*)/q_i^*$ are the tension and die-flow errors in percent at step i, $w_T = 1$ and $w_C = 4$ their weights, $\kappa_F = 0.2\,\mathrm{N}^{-2}$ and $\kappa_P = 0.5$ per percent squared the prices of the moves at each block boundary j, $\lambda_T = \lambda_P = 1$ two price multipliers set on the app, and ν the slack of the soft tension limit. The die-flow target q_i^* is the pump feedforward of

Equation (5.12) at the predicted speed $v_{1,i}$ divided by $\hat{g}$ and capped at 100 %, with the weight's denominator floored at its value for v_{min}. The bounds $|F| \le \tau_{max}/\hat{R}_r$, the shipped trim's clamp, and $0 \le N_p \le 100\,\%$ are hard. The tension band of $\pm 10\,\%$ is soft, judged every third step with one slack over the whole horizon priced at $\gamma_2 = 10^3$ per percent squared and $\gamma_1 = 10^2$ per percent, and with a 0.5 s funnel where the free response already lies outside it. The weights were chosen once, on case 9, and are never retuned per case. With inertia compensation on, the shipped compensation stays beneath the plan at every 5 ms scan, so the applied torque is $\tau_r = \hat{R}_r F + \hat{J}_r a_{ff}/\hat{R}_r$ within $\pm\tau_{max}$. The coating interlock, the braking after a break and the dancer's stops act above the plan unchanged.

A plan is used only when its solver status reads optimal, every planned command lies inside its bounds with the first equal to the applied move, and the cost, the slack and every prediction converted to newtons and grams per square meter are finite. Otherwise the shipped trims take over, which under the MPC track the applied torque and the applied ratio of pump speed to feedforward at every scan. With the noise off a forced rejection moves the torque by at most 0.0007 N m and the pump by at most 0.06 % on the hand-over scan, except in the challenge as loaded, where the trim's 150 % limit on a feedforward frozen at the old speed moves the pump by 21.0 %. At ten captured decisions an independent condensation reproduces the solver's objective to a relative 2.4×10^{-12} with the optimality conditions met. In Chrome a plan costs 104 µs, from the difference in wall time between 300 s of case 9 under the MPC and under the loops, which is 0.21 % of one core in real time and 4.1 % at the 20× playback speed.

The simulation model, the instruments and the controller's model stay apart. The plant integrates Equations (5.1) to (5.9) with the true radii, strains, splices and viscosity. The MPC reads only the encoder, the rewind's diameter sensor, the noisy load cell, the noisy averaged gauge, its own applied commands, the interlock and break signals, the setpoints and the declared constants. A replay of the same measurements over a different hidden plant reproduces its commands bit for bit, and the application's tests catch each of 25 copies of the controller doctored to read one true quantity or run parameter or to keep its state outside the engine.

5.5 Three Reduced Models

5.5.1 The Dancer Through a Ramp

With $F_{air} = 2T_{1,SP}$, Equation (5.5) says exactly what span 1 tension does while the dancer moves, as shown in Equation (5.16).

$$T_1 - T_{1,SP} = -\tfrac{1}{2}\left(c_d \dot{x} + m_d \ddot{x}\right) \tag{5.16}$$

The dancer holds the tension only while it is still. In the ramp of case 2 the identity holds in the engine to 0.004 N. Span 1 is so stiff, 56 kN/m for the film, that its strain barely changes, so the unwind surface moves at $v_0 = v + 2\dot{x}$. Substituting into the unwind balance of Equation (5.3) with the torque of Equation (5.11), the sensed radius equal to the true one and the friction canceled by its feedforward, gives a third-order loop driven by the line acceleration, as shown in Equation (5.17).

$$\begin{aligned} &\left(2m_u + \tfrac{1}{2}m_d\right)\ddot{x} + \tfrac{1}{2}c_d\dot{x} + K_c x \\ &\qquad + \frac{K_c}{\tau_I}\int x\,dt = -m_u\left(\dot{v} - \hat{a}\right) \end{aligned} \tag{5.17}$$

where $m_u = J_u/R_u^2$ is the roll's equivalent mass at its surface and $\hat{a}$ the compensated acceleration, a_{ff} with inertia compensation and zero without it. For the full film roll, $m_u = 201$ kg and the roots are $-0.272\ \mathrm{s}^{-1}$ and $-0.479 \pm 0.703\,j\ \mathrm{s}^{-1}$, a damped oscillation of 8.9 s period. Without compensation the proportional term alone would need a dancer offset of $m_u a/K_c$ to supply the accelerating torque, 101 mm for the film at 12 m/min per s against stops at 150 mm, and the integral removes it only over several seconds. Driven by the case 2 ramp, Equation (5.17) gives a dancer swing from −93.1 to +76.2 mm and a span 1 tension from 224.1 down to 166.7 N at 11.16 s and 16.06 s, while the engine with its noise switched off gives −92.8 to +75.5 mm and 224.1 to 166.8 N at the same instants. For

the full coil of foil, $m_u = 386\,\text{kg}$, and the challenge ramp gives -115.0 to $+85.5\,\text{mm}$ and 86.0 to 177.6 N against the engine's -115.1 to $+85.7\,\text{mm}$ and 85.9 to 177.7 N.

The same model prices the speed-ratio radius estimate. With $\hat{R}_u = R_u(1 - 2\dot{x}/v_0)$ the holdback, in tension units, changes by $-2T_{1,SP}\dot{x}/v_0$, and the damping of Equation (5.17) becomes $c_d/2 - 2T_{1,SP}/v_0$. That is +300 N s/m for the film at 120 m/min, but −100 N s/m for the foil at 24 m/min and −300 N s/m for the film at 30 m/min, where the model's roots move to $0.155 \pm 0.740\,j$ and $0.455 \pm 0.966\,j\ \text{s}^{-1}$. In a test made for this chapter with a wrapper that replaces the unwind radius estimate by the ratio and leaves the controller otherwise unchanged, the foil line running steadily at 24 m/min loses its dancer, whose swing roughly doubles every 5 s until the foil breaks at 56.9 s, and the film line at 30 m/min breaks at 14.45 s. The error grows as the speed falls, which is why the application measures the diameter.

5.5.2 The Rewind on Its Span

Span 2 is a spring between the coating roll, which holds its speed, and the rewind, which has an equivalent mass $m_r = J_r/R_r^2$ of 61 kg at 128 mm. The load-cell trim acts on the tension that the spring itself applies to that mass. Linearized with the feedback delayed by T_d, Equations (5.3), (5.4) and (5.11) give the characteristic equation of Equation (5.18).

$$m_r s^2(s+\alpha) + (k+cs)\left[s + K_c\left(s + \frac{1}{\tau_I}\right)\mathrm{e}^{-sT_d}\right] = 0 \tag{5.18}$$

where $k = EA/L_2$, $c = \tau_w k$ and $\alpha = v/L_2$ is the transport term. The first term in the bracket is the tension acting on the roll directly, and the proportional trim adds K_c times as much again, so the oscillation has a period near $2\pi\sqrt{m_r/[(1+K_c)k]}$. With $T_d = 0$ and a small α, the Routh criterion gives Equation (5.19).

$$\tau_I > \frac{m_r K_c}{(1+K_c)^2\,\tau_w k} \tag{5.19}$$

The span's damping is itself proportional to its stiffness, so the stabilizing side of the Routh condition grows with the square of the stiffness and the destabilizing side only with the stiffness, and the shortest usable integral time is inversely proportional to the stiffness of the web. For the hot trim of case 4, $K_c = 1$ and $\tau_I = 0.5\,\text{s}$, Equation (5.19) with the transport term restored gives a critical integral time of 0.30 s on the PET and 0.91 s on the web of 0.3 times its stiffness. The PET is held and the softer web is not, and a splice to a stiffer web moves the loop away from instability rather than toward it.

🕮 **Pause.** Try Exercise 5.2 before reading on.

The 5 ms control interval is not negligible here. The oscillation is so lightly damped, at a damping ratio of 0.018 on the PET without delay, that the half-sample delay of the zero-order hold, $T_d = 2.5\,\text{ms}$, removes a large share of the margin. The growth rates measured on the engine after a 10 N setpoint step, with the load-cell noise off, fall on the delayed model and not on the continuous one, as shown in Figure 5.2. At 0.3 times the PET stiffness the oscillation grows at $0.285\ \text{s}^{-1}$ with a 0.651 s period against $0.282\ \text{s}^{-1}$ and 0.652 s predicted, and on the PET it decays at $0.114\ \text{s}^{-1}$ against $0.117\ \text{s}^{-1}$. The stability limit of the hot trim is 0.80 times the PET stiffness with the delay and 0.59 without it.

5.5.3 The Coat-Weight Trim and Its Delay

From the trim k to the gauge reading in percent of target the plant is a gain $g = \eta/\eta_{nom}$ behind the transport delay and four small lags, as shown in Equation (5.20).

$$G(s) = \frac{g\,\mathrm{e}^{-\theta s}}{(\tau_P s+1)(\tau_A s+1)(\tau_g s+1)(\ell_b s/v+1)} \tag{5.20}$$

where $\tau_g = 1\,\text{s}$ is the gauge average and $g = 0.958$ after a thin batch. With that gain and the film recipe trim, $K_c = 0.3$ and $\tau_I = 6\,\text{s}$, the loop has a gain margin of 2.87 at 120 m/min, where $\theta = 9.49\,\text{s}$,

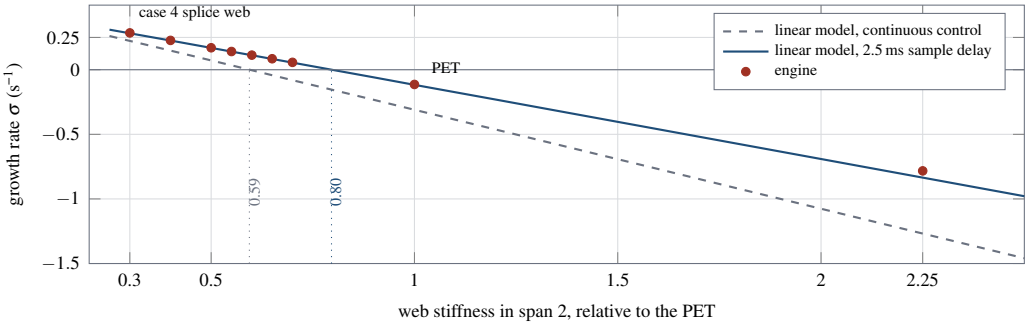

Figure 5.2 Growth rate of the rewind tension oscillation under the hot load-cell trim ($K_c = 1$, $\tau_I = 0.5\,\mathrm{s}$) against the stiffness of the web in span 2. The lines are Equation (5.18) at $R_r = 128\,\mathrm{mm}$ without the half-sample delay (dashed) and with it (solid), and the dots are the engine's envelope after a 10 N setpoint step at 27 s of case 4. The dotted verticals mark the two stability limits, 0.59 and 0.80.

and of 0.93 at 30 m/min, where $\theta = 37.97\,\mathrm{s}$, with a crossover period of 135.7 s. The critical integral time at 30 m/min is 6.40 s. The engine agrees. Swept with the gauge noise off, its oscillation after the thin batch grows at $1.56 \times 10^{-3}\,\mathrm{s}^{-1}$ for $\tau_I = 6\,\mathrm{s}$ against $1.52 \times 10^{-3}\,\mathrm{s}^{-1}$ predicted, is neutral at 6.4 s, and decays at $2.12 \times 10^{-3}\,\mathrm{s}^{-1}$ for 7 s against $2.17 \times 10^{-3}\,\mathrm{s}^{-1}$. The rule of thumb for integral action on a pure delay, $K_c\theta/\tau_I < \pi/2$, gives 1.90 at 30 m/min and 0.47 at 120 m/min, which is on the conservative side of the same limit.

The recipe trims follow a standard rule. For a pure delay the simple internal model control (SIMC) rule asks for integral action with $K_i = 1/[g(\tau_c + \theta)]$, and with $\tau_c = \theta$ and $K_c = 0.3$ that is an integral time of $0.6\,g\,\theta$ [20]. With $g = 1$ this gives 5.7 s for the film at 120 m/min, against the recipe's 6 s, and 28.5 s for the foil at 24 m/min, against 30 s. At 30 m/min the same rule asks for 22.8 s, close to the 25 s that case 6 offers as its fix. Each recipe trim is right at one speed only, because the delay it faces is a distance. The SIMC integral time is proportional to $\theta \approx L_g/v$, so an integral time scheduled on the line speed would keep about the same margin at every speed. The basis weight loops of Chapter 3 face the same accounting, with a scanning gauge whose averaging adds to the transport delay.

5.6 Scenarios and Measured Results

The ten cases of Table 5.5 were run on the engine at the default seed 20261023. Settings for each comparison were made before Run, as a reader makes them on the app (Section 1.6.4), with scheduled events applied at their nominal boundary. The case 1 order was fired at 60 s. Deviations are those of the true tensions and of the true coat weight under the gauge, as the scorecard reports them, and scrap is counted in 5 cm cells and quoted to 0.05 m. Table 5.6 collects the results of cases 1 to 8 and Table 5.7 the challenge. Section 5.6.7 compares the predictive controller with the loops on case 9 and the others.

5.6.1 A Steady Line and a Ramp With Feedforward

Over 120 s of case 1 the span 1 tension holds 200 N to 0.047 N rms, because the dancer's air load sets it, while span 2 holds 150 N to 0.78 N rms as its trim answers load-cell noise of 1 N. The true coat weight at the gauge stays within 0.30 % of target and the line makes 7195 m^2 of good product an hour. An order of +50 % placed after a minute moves span 1 by less than 0.2 N and the dancer by less than 0.5 mm with both feedforwards on. The coating still dips 3.60 % at the die during the 10 s ramp at 6 m/min per s, because the pump and hose lag the speed by their $\tau_P + \tau_A = 0.8\,\mathrm{s}$, and that scraps 10.65 m. At 60 m/min per s the dip is 21.5 % deep but lasts about a second, and it scraps 10.30 m. The shallow dip sits only a few tenths of a percent past the band, so its scrap depends on

Table 5.5 The Ten Cases, in Teaching Order

#	Scenario	Teaching point
1	Steady coating, film at 120 m/min	the reference line
2	Speed ramp, both feedforwards off, to 180 m/min at 10 s	one speed, two paths
3	Rewind from 80 mm at 210 m/min on fixed torque, trim off	$T = \tau/R$
4	Splice at 10 s to a web of 0.3 times the stiffness, hot trim	the span is a spring
5	Electrode at 24 m/min, half viscosity at 20 s	a change behind a delay
6	Film at 30 m/min, thinner batch at 10 s	a delay fixed in meters
7	Order of 240 m/min at 10 s	a capacity limit
8	Stop at 15 s, restart at 45 s	standstill logic
9	Film at 30 m/min, thinner batch at 10 s, order of 120 m/min at 150 s, MPC	a plan against the loops
10	Challenge, foil order 24 to 42 m/min at 20 s	the whole design

Table 5.6 Measured Results, Cases 1 to 8. Ramp Rates in m/min per s

Case and condition	Result
1, over 120 s	T_1 0.047, T_2 0.78 N rms, *cw* within 0.30 %
1, order +50 %, ramp 6 / 60	die dip 3.60 / 21.5 %, 10.65 / 10.30 m
2, as shipped	die 8.02 g/m^2, gauge 7.07 s later, 208.90 m
2, span 1 / span 2	+12.1, −16.6 / −9.6, +6.9 %
2, dancer	−93.0 to +75.7 mm
2, pump feedforward / both	22.30 / 12.80 m, span 1 within 0.2 N
3, fixed torque	122.0, 73.9, 60.3 N at 99, 163, 200 mm
3, fixed torque, 700 s	wrinkles 508.5 s, 2362.9 of 2447.8 m
3, $T \times R$ / fixed with trim	149.7–152.7 / 145.4–151.7 N
4, hot trim, soft web	breaks at 37.85 s
4, τ_I 2 s / default, soft	within 2.44 / 1.44 %
4, hot trim, 2.25 times as stiff	within 2.94 %
5, half viscosity	die 191.39 g/m^2, gauge 47.63 s later, 20.25 m
5, τ_I 15 / 60 s	20.15 / 20.35 m, peak 202.07 g/m^2
6, recipe trim	peaks at 13.11, reaches 10.81 g/m^2 at 600 s, 188.15 m
6, τ_I 25 s / K_c 0.1 / at 120	20.55 / 25.40 / 23.85 m
7, order 240 m/min	pump full at 223.1 m/min, 11.34 g/m^2, 644.5 m
7, capacity limit	held at 215.5 m/min, 22.05 m
8, interlock	die off 24.64 to 45.47 s, 4.52 mL, 20.60 m
8, no interlock	1.891 mL spilled, 141.75 g/m^2, 27.65 m

where the noise-driven trim happens to be when the order arrives. Placed at 0 s, 10 s or 30 s instead, the same order scraps 10.00, 9.20 and 17.45 m.

5.6.2 One Speed Change, Two Paths

Case 2 repeats the order at 12 m/min per s with the pump feedforward and inertia compensation off. The coating at the die falls 33.2 %, from 12.0 to 8.02 g/m^2, as soon as the ramp begins, because the fixed flow is spread over half as much web again. The gauge first sees it 7.07 s later, less than the 9.49 s of the steady line because the web accelerates on its way, and the slow trim is not back inside the band until 87.1 s. The order costs 208.90 m of scrap. At the same moment the full unwind has to be spun up through the web. Span 1 rises 12.1 % as the ramp starts and falls 16.6 % as it ends, the dancer travels from −93.0 to +75.7 mm, and span 2 swings from −9.6 to +6.9 % as the rewind lags and catches up. The screen 4 s into the ramp is shown in Figure 5.3. The pump feedforward alone cuts the scrap to 22.30 m, of which 16.35 m is web that lived through the span 1 excursion, and adding inertia compensation holds span 1 within 0.2 N and the dancer within 0.5 mm and leaves 12.80 m,

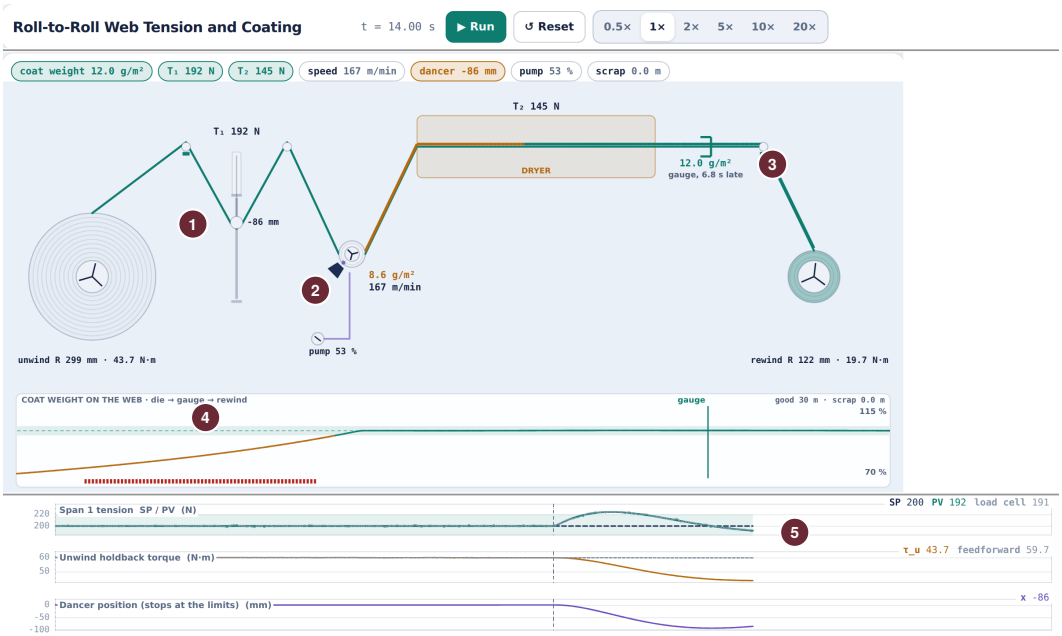

Figure 5.3 Case 2, "Speed ramp, feedforward off", paused at $t = 14.00$ s, 4 s into the ramp from 120 to 180 m/min at 12 m/min per s with the pump feedforward and inertia compensation off. The picture is a composite of the app's header, its plant view with the web map, and the first three of its trend strips. The side view shows the unwind at 299 mm and 43.7 N m, the dancer (1) 86 mm toward its short-span stop with a warning badge, span 1 at 192 N, span 2 at 145 N, the die (2) laying 8.6 g/m^2 in the warning color with the pump at 53 %, and the gauge (3) still reading 12.0 g/m^2, 6.8 s late. On the web map below (4), the thin section rises from the die toward the ±3 % band short of the gauge, and marks flag web that was in span 1 while its tension was outside ±10 %. The trends (5) show span 1 rising above 220 N and falling back through its setpoint, the holdback torque 16 N m below its 59.7 N m feedforward and the dancer at −86 mm. No scrap is yet wound, and the worst tension deviation of the run, 12.1 %, is the span 1 peak of Table 5.6.

all of it the pump's lag during the ramp. Inertia compensation alone removes the tension scrap and leaves all 208.90 m of coat-weight scrap.

> **TRY IT IN THE APP** **2 · Speed Ramp, Feedforward Off**
>
> Load the scenario, press Run and watch the coat weight printed beside the die. When the order of 180 m/min starts the ramp at 10 s, the coating falls from 12.0 to 8.02 g/m^2 at once, while the gauge downstream keeps reading the target for another 7.07 s.

5.6.3 The Radius and the Splice

Case 3 builds the rewind from 80 mm at 210 m/min on the torque that gave 150 N at the start, with its load cell out of service. The tension follows $T = \tau/R$, reading 122.0 N at 99 mm, 73.9 N at 163 mm and 60.3 N at 200 mm, slightly above $150 \times 80/R$ because the rewind decelerates as it grows and returns a little torque. The tension leaves its band at 24.3 s and the web wrinkles at 508.5 s, so 2362.9 m of the 2447.8 m wound in 700 s is scrap. Multiplying the setpoint by the measured radius holds 149.7 to 152.7 N over the whole build, and the load-cell trim on fixed torque holds 145.4 to 151.7 N, removing by feedback an error that radius compensation prevents outright.

DID YOU KNOW? A Fixed Torque Loses Tension as the Roll Grows

From the simulator: a rewind built from 80 mm on the torque that gave 150 N at the start reads 60.3 N at 200 mm, and the web wrinkles at 508.5 s, so 2362.9 of the 2447.8 m wound is scrap. Multiplying the setpoint by the measured radius holds 149.7 to 152.7 N (Table 5.6).

Case 4 splices an almost empty roll to a web with 0.3 times the stiffness of the PET, under a trim tuned hot on the PET that holds it within 2.94 % before the splice. The soft web reaches the coating roll at 12.00 s and fills span 2 at 24.01 s. The oscillation grows as the soft web fills the span and afterward at about the rate of Figure 5.2 until the web breaks at 37.85 s. The same hot trim holds a web 2.25 times as stiff within 2.94 %, the hot gain with a 2 s integral time holds the soft web within 2.44 %, and the default trim of 0.5 and 2 s holds it within 1.44 % through 40 s.

TRY IT IN THE APP 4 · Splice to a Softer Web

Load the scenario, press Run and watch span 2 after the splice at 10 s. The oscillation grows as the softer web fills the span until the web breaks at 37.85 s. Press Reset, set Load-cell trim integral time to 2 s, keep the gain at 1 and run again: span 2 now stays within 2.44 %.

5.6.4 An Unmeasured Change and a Late Gauge

Case 5 runs the electrode recipe, and at 20 s a batch at half the viscosity doubles the pump slip. The die makes 191.39 g/m^2, 4.30 % light, and the gauge first sees it 47.63 s after the batch arrives. The trim restores the target and the gauge is back in band at 118.8 s, after 20.25 m of scrap. With integral times of 15 and 60 s the scrap is 20.15 and 20.35 m, and the faster trim overshoots at the die, to 202.07 g/m^2. The 19 m already coated between the die and the gauge when the gauge first sees the change is most of that scrap, so the delay rather than the tuning sets what an unmeasured disturbance costs.

Pause. Try Exercise 5.1 before reading on.

Case 6 runs the film at 30 m/min with the trim tuned for 120 m/min and a thinner batch at 10 s. The gauge is 37.97 s late and the loop sits just past its limit, so the gauge rings with a period of about 133 s and growing swings, peaking at 12.59, 12.74, 12.85 and 13.11 g/m^2 and falling to 11.34, 11.24 and 11.05 g/m^2, and it reaches 10.81 g/m^2 at 600 s, the bottom of the next swing. In those ten minutes 188.15 m of the 299.8 m wound is scrap. A 25 s integral time recovers without overshoot and scraps 20.55 m, a gain of 0.1 scraps 25.40 m, and on a line started at 120 m/min the recipe trim meets the same batch with 23.85 m. Reloading case 6 and setting the order to 120 m/min before Run starts at that speed, with the same 23.85 m of scrap, while an order given after Run ramps from the current line speed instead. Because Reset keeps the reader's settings (Section 1.6.1), the recipe integral time comes back only when the scenario is reloaded.

5.6.5 The Pump's Capacity and a Stop

Case 7 orders 240 m/min at 12 g/m^2, more than the 226.8 m/min that the pump can coat at full speed. The pump reaches 100 % at 223.1 m/min during the ramp, and at 240 m/min the coating is 11.34 g/m^2, 5.50 % light, so 644.5 m of the 689.3 m made in 180 s is scrap. The trim sits at its moving limit of 94.5 % without winding up. Holding the speed within the pump's capacity stops the ramp at 215.5 m/min, where the coat weight holds and the scrap is the 22.05 m of the ramp itself. The capacity limit is an override, because the speed the customer ordered gives way to the speed at which the product can still be made. Volume 16, Chapter 9 likewise caps its process speed, there just below the rate at which its entry can supply strip.

Case 8 stops the film line at 15 s and restarts it at 45 s. The pump follows the speed down through its lag, so the last meters are coated heavy, up to 32.9 g/m^2 at the die just before it lifts. The interlock lifts the die at 24.64 s, sends 4.52 mL of coating to the drip tray, re-engages at 45.47 s and holds the

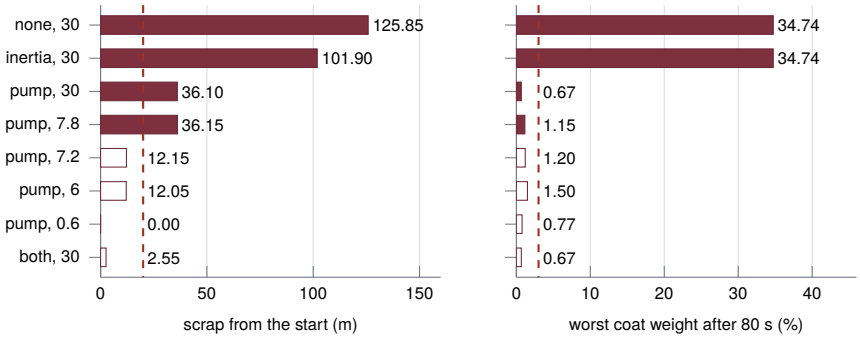

Figure 5.4 The challenge configurations of Table 5.7, labeled by feedforward and ramp rate in m/min per s. Dark bars fail the challenge and white bars meet it. The dashed lines are its 20 m scrap limit and its ±3 % coat-weight band. Notice that the pump feedforward alone brings the coat weight inside the band, and that what remains, the tension scrap of a steep ramp, goes either to inertia compensation or to a gentler ramp.

trim until fresh web reaches the gauge, and the stop costs 20.60 m. Without it the bead grows on the stationary web and 1.891 mL spills from 25.22 s, the restart lays down a stripe of 141.75 g/m^2, which is the full bead, and the trim chases the defect past the gauge, for 27.65 m of scrap.

5.6.6 The Graded Challenge

The challenge changes the electrode order from 24 to 42 m/min at 20 s with a 30 m/min per s ramp and both feedforwards off. It is met with no break, both tensions within ±10 % and the coat weight at the gauge within ±3 % from 80 s to the end of the 200 s run, and no more than 20 m of scrap in total, with the model parameters and disturbance settings left as supplied. The configurations measured are listed in Table 5.7. As shipped the foil survives, peaking at 177.6 N in span 1 against its 360 N limit, but the coat weight is 34.74 % off after 80 s and 125.85 m is scrap. Only the pump feedforward is indispensable, because without it the trim, capped at 150 % of a feedforward frozen at the old speed, cannot supply the 75 % more flow the new speed needs. With inertia compensation on and no pump feedforward it reaches 148.6 % while the gauge still reads 164.8 g/m^2 at 200 s. With the pump feedforward alone the steep ramp scraps 36.10 m, of which 34.95 m carries tension defects, adding inertia compensation leaves 2.55 m, and leaving the compensation off but ramping at 7.2 m/min per s or less also passes, with scrap falling from 12.15 m at that rate to none at 0.6 m/min per s. At 7.8 m/min per s span 2 dips to 89.5 N, 10.5 % low, which flags all 24 m of web in the span at once and lifts the scrap to 36.15 m. The 2.78 % in span 2 after 80 s is the same in every configuration, because it is load-cell noise at 106.7 s that the fixed seed repeats, and it falls to 0.15 % with that noise off. Figure 5.4 sets the scrap and the coat weight of each configuration against the two limits of the challenge.

5.6.7 Model Predictive Control Against the Loops With Feedforward

The MPC is set against the loops in Table 5.8, every case at its own seed, duration, instruments, noise and events. The loops run with both feedforwards on, case 3 with radius compensation and its load cell out of service, case 4 with its hot trim and the challenge with its suggested settings, while the MPC carries case 9's settings into every case unchanged. The measures accumulate at the 1 ms step: the scorecard's scrap, the time and depth of the true coat weight under the gauge outside ±3 %, the rms span 2 tension about its setpoint, and the pump travel summed over every 5 ms scan. Three loop variants, none of them a setting in the application, hand the loops what the plan carries. Lead evaluates the pump feedforward at the speed master's reference 0.75 s ahead, the preview the MPC receives, accel evaluates it at $v + (\tau_P + \tau_A)\, a_{ff}$, the measured speed advanced by the two lags with no

Table 5.7 The Challenge. Worst Deviations Count From 80 s and Scrap From the Start

Feedforward, ramp	T_1 max (N)	Dancer (mm)	*cw* worst (%)	Scrap (m)	Met
none, 30	177.6	−115/+86	34.74	125.85	no
inertia, 30	120.1	−0.2/+0.3	34.74	101.90	no
pump, 30	177.6	−115/+86	0.67	36.10	no
pump, 7.8	145.2	−101/+80	1.15	36.15	no
pump, 7.2	143.2	−98/+79	1.20	12.15	yes
pump, 6	139.4	−91/+76	1.50	12.05	yes
pump, 0.6	121.9	−10/+10	0.77	0.00	yes
both, 30	120.1	−0.2/+0.3	0.67	2.55	yes

Ramp rates in m/min per s. "Pump" is the speed feedforward to the pump, "inertia" the inertia compensation and "both" the suggested configuration. The worst span deviations after 80 s are 0.09 and 2.78 % in every row.

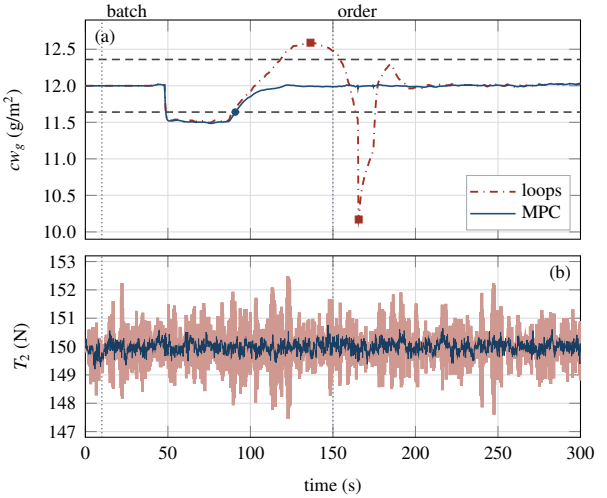

Figure 5.5 Case 9 under the MPC and under the loops with both feedforwards at the case's own seed. (a) True dry coat weight of the web under the gauge against the ±3 % band (densely dashed). (b) True span 2 tension. The MPC is the solid line in (a) and the thin dark trace in (b), the loops the dash-dotted line in (a) and the wide light trace in (b). The thinner batch arrives at 10 s and the order of 120 m/min at 150 s (dotted verticals). The square marks show the loops' overshoot to 12.59 g/m^2 and their dip to 10.17 g/m^2 after the ramp, and the round mark the MPC's return into the band for good at 90.9 s.

preview at all, and aligned replaces the coat-weight trim by the MPC's own $\hat{g}$. Hindsight is the best of 20 coat-weight trims, K_c from 0.1 to 0.5 and τ_I from 3 to 50 s, found on case 9 after the fact.

Case 9 runs the film at 30 m/min with both feedforwards and the recipe trim. A batch at half the viscosity arrives at 10 s, the die lays 11.50 g/m^2, 4.2 % light, and the gauge 38.0 s downstream first reads the thin web 1.5 % light at 48.2 s. The MPC compares the gauge with what its model laid on the same web, so $\hat{g}$ has moved 0.5 % by 50.4 s, settles at 95.83 %, and the coat weight returns inside the band for good at 90.9 s, as shown in Figure 5.5. Its 21.05 m of scrap is the web coated light before the gauge could see it. The recipe trim, tuned for 120 m/min, overshoots to 12.59 g/m^2 at 136.4 s and still reads 12.48 g/m^2 when the order of 120 m/min arrives at 150 s, for 30.90 m of scrap. Through the 15 s ramp the loops' pump follows the measured speed through the 0.8 s of pump and hose lag and the die dips to 10.17 g/m^2, for 32.10 m more, while the plan drives the pump ahead of the flow and holds the die between 11.94 and 12.01 g/m^2. Span 2 holds 150 N to 0.212 N rms against the loops' 0.760 N.

Removing one part of the plan at a time locates the gain. Without the pump and hose lags in its

Table 5.8 The MPC Against the Loops With Both Feedforwards, and the Loops Handed What the Plan Carries

Case	Structure	Scrap (m)	*cw* off band s (% s)	T_2 rms (N)	Pump (%)
9	loops	63.00	88.4 (158.3)	0.760	54.4
	MPC	21.05	42.1 (44.1)	0.212	51.7
	MPC, preview off	21.05	42.1 (44.1)	0.212	50.7
	MPC, no pump lags	40.40	52.3 (72.6)	0.212	43.9
	MPC, delay in time	71.00	94.0 (122.6)	0.212	52.8
	loops, τ_I 25 s	39.95	50.8 (79.1)	0.760	52.8
	loops, lead, τ_I 25 s	20.55	41.1 (43.5)	0.760	50.6
	loops, accel, τ_I 25 s	20.80	41.3 (43.5)	0.760	54.6
	loops, lead, aligned	21.40	42.8 (44.8)	0.760	43.9
	hindsight	39.35	— (72.4)	—	—
	hindsight, lead	19.70	— (41.8)	—	—
	hindsight, accel	19.80	— (41.8)	—	—
	loops, load cell 0.1, 5 s	63.00	88.4 (158.3)	0.218	54.4
2	loops	12.80	4.3 (10.9)	0.742	43.1
	MPC	0.00	0	0.215	38.3
	MPC, no pump lags	12.75	4.3 (10.3)	0.215	29.4
	loops, accel	0.00	0	0.742	48.6
3	loops	0.00	0	0.233	56.3
	MPC	0.00	0	0.274	16.7
4	loops, hot trim	29.45[a]	0	—	54.9
	loops, default trim	1.05	0	0.814	5.0
	MPC	1.05	0	0.263	1.6
5	loops	20.25	50.6 (55.0)	0.670	31.7
	MPC	20.70	51.8 (55.7)	0.297	11.8
	loops, K_c 0.3, τ_I 12 s	20.10	50.3 (52.7)	0.670	32.7
6	loops	188.15	386.6 (950.8)	0.767	18.3
	MPC	21.05	42.1 (44.1)	0.214	3.4
	MPC, delay in time	207.35	418.7 (1723.9)	0.214	21.7
	loops, τ_I 25 s	20.55	41.1 (43.5)	0.767	8.1
7	loops	644.50	162.7 (402.3)	0.734	47.5
	MPC	617.55	155.8 (388.8)	0.213	49.6
	loops, capacity limit	22.05	6.2 (13.3)	0.737	62.1
	MPC, capacity limit	0.00	0	0.214	54.3
8	loops	20.60	10.3 (120.5)	0.773	115.7
	MPC	0.20	0.1 (2.7)	0.216	312.4
	MPC, preview off	4.25	2.1 (4.2)	0.216	226.4
	loops, accel	0.50	0.3 (5.8)	0.773	121.4
10	loops, suggested	2.55	3.6 (28.0)	0.662	65.5
	MPC, suggested	0.60	0.9 (6.6)	0.294	69.5
	MPC, preview off	0.70	1.0 (11.0)	0.294	69.5
	loops, accel	0.80	1.1 (8.7)	0.662	79.2
	loops, as loaded	125.85	152.7 (3897.0)	1.695	28.6
	MPC, as loaded	10.95	0.9 (6.6)	0.322	69.5

Runs last 300, 180, 600, 90, 600, 600, 180, 180 and 200 s for cases 9, 2 to 8 and the challenge 10. The band is ±3 % of the target, the depth in parentheses integrates the excursion beyond it, and pump travel sums the changes of the pump command. The MPC made 20 plans a second, 6000 on case 9, and rejected none, at a mean of 160 µs a plan in Node on a shared host.
[a]The web breaks at 37.85 s.

model the plan scraps 40.40 m on case 9 and 12.75 m on case 2. Reading the record a fixed 9.5 s late, a delay fixed in time rather than in meters, it rings on the slow line and scraps 71.00 m on case 9 and 207.35 m on case 6. The preview leaves every scrap count of cases 2 to 6 and 9 unchanged and matters on the steep stop of case 8 (0.20 against 4.25 m) and a little in the challenge (0.60 against 0.70 m). The soft band never binds, and removing it changes no number. The estimate time moves the result in the expected direction. At 2 and 30 s instead of the shipped 10 s, the plan scraps 19.65 and

24.30 m on case 9, because a faster $\hat{g}$ takes up the thinner batch sooner once the gauge sees it. The loops reach the same result when handed the same knowledge. A trim retuned to 25 s for the slow line scraps 39.95 m, and 20.55 m with a lead on the pump or 20.80 m with accel, against the plan's 21.05 m. With the plan's own $\hat{g}$ in place of the trim and either lead they scrap 21.40 m. The tension gain comes from acting gently on a noisy load cell, which the loops can also do, since the load-cell trim calmed to $K_c = 0.1$ and $\tau_I = 5$ s holds 0.218 N rms on case 9. On case 9, therefore, neither the preview, the soft band nor the optimization adds a measurable share.

🕮 **Pause.** Try Exercise 5.3 before reading on.

Elsewhere the MPC removes the 12.80 m of case 2, holds the softer web of case 4 with only the 1.05 m of the splice zone scrapped, where the hot trim breaks it at 37.85 s, and with the capacity limit on scraps nothing on case 7 against the loops' 22.05 m. The losses are reported as measured. On case 5 the plan scraps 20.70 m against the loops' 20.25 and a retuned trim's 20.10 m, and it loses on each of four noise seeds, because its 10 s estimate is slower than the recipe trim at 24 m/min. A 2 s estimate, not shipped, would scrap 19.70 m. On cases 6 and 9 a trim retuned to 25 s scraps 20.55 m against 21.05 m, a trim of 0.2 and 6 s matches the plan's 21.05 m on case 6 with less depth (41.5 against 44.1 % s), and the hindsight trims with a lead or accel reach 19.70 and 19.80 m on case 9. With the load cell out of service in case 3 the plan runs on its model alone and holds 0.274 N rms against the radius-compensated loops' 0.233 N. On case 8 it moves the pump 312.4 % to the loops' 115.7 %, driving it past its steady value to push the flow through the lags at the stop and the restart. Both structures lose most of case 7 to the pump's capacity. At seeds 11, 12 and 13 case 9 scraps 21.30, 21.00 and 21.45 m under the MPC against 63.00, 62.65 and 64.25 m under the loops, and halving the physics step changes no scrap count of cases 2, 5, 6, 8, 9 and the challenge under either structure.

> ✦ **DID YOU KNOW? The Optimizer Is Not Where the Gain Lies**
>
> From the simulator: on case 9 the predictive controller scraps 21.05 m against the loops' 63.00 m. Loops handed the same knowledge, a 25 s trim and a lead on the pump, scrap 20.55 m, and the plan without its model of the pump lags scraps 40.40 m (Table 5.8).

With its suggested settings the challenge is met by both structures, with 2.55 m under the loops and 0.60 m under the MPC. As loaded, with both feedforward switches off, the loops fail with 125.85 m while the MPC meets the target with 10.95 m, all of it tension scrap. The plan supplies the rewind's inertia torque itself, but the unwind keeps the shipped law without inertia compensation, so span 1 still peaks at 177.6 N, 48 % over its setpoint against a 360 N break, and the dancer swings from −115 to +86 mm before scoring starts at 80 s. The challenge does not score those limits, and they are the same under the loops as loaded.

5.7 Technology in Practice

What Measures It

The line has two ways to read tension, and the chapter uses both: the dancer of span 1, whose position XT reports the mismatch of speeds while its air load sets the tension, and the load cell WT2 under an idler of span 2, which reads the tension directly. The radius that both winder torques are scaled by can be read by reflective laser sensors beside the rollers or, where no sensor is wanted, estimated by integrating the web thickness over the turns of the roll [14]. The chapter's DT1 and DT2 are ultrasonic diameter sensors with 0.3 mm of noise and a 1 s average. For the coat itself, laser caliper sensors mounted and calibrated at the oven inlet of a slot-die coater have monitored the wet thickness of cathodes and anodes with a precision better than ±2 % [21]. The chapter's gauge AT reads the dry coat weight 19 m after the die, through a 1 s average and noise of 0.5 % of the target.

TECHNOLOGY **Sensor: The Tension Load Cell Under an Idler**

A load cell is one or more strain gauges bonded to a metal structure of precisely known elastic properties, which stretches with the applied force like an extremely stiff spring, so the gauges turn force into a change of resistance [22]. Web tension measuring devices usually mount load cells under the bearings of a roller that the web wraps at a fixed angle, and the force on the cells is then proportional to the tension [23]. As tension sensors, load cells are expensive and need frequent maintenance and calibration [13]. An electric current through a load cell, from a ground fault or welding, is read as a change of resistance and so as a false force [22], which a load-cell trim would hold on its setpoint. The chapter's WT2 is ideal apart from noise of 1 N with a 20 ms correlation time.

What Moves It

Three drives move the line. The coating roll is the speed master, a stiff speed-regulated drive with a 50 ms response, and the transport speed of every tension zone is referred to such a roller [3]. The two winders are torque drives, and the gear pump that feeds the die is a third speed-commanded drive whose delivered flow falls short of its displacement by a slip that grows as the viscosity falls. The chapter gives the pump a 0.3 s lag and its hose and manifold a 0.5 s lag, and those 0.8 s are what the pump feedforward cannot remove during a ramp.

TECHNOLOGY **Final Element: The Winder Drive in Torque Mode**

A variable frequency drive achieves precise speed control of an alternating-current motor by varying the frequency of the power it feeds, and precise torque control by varying the voltage and current. An unwind that holds the web back works its motor as a generator. Dynamic braking dissipates that energy in a resistor, while a regenerative drive pumps it back into the supply grid. Every electronic braking method depends on the drive working, and a drive whose braking needs line power has no braking capacity at all when that power is lost [22]. An unwind left without braking can no longer hold its roll back. The chapter's winders are ideal torque drives, commanded as $\hat{R}$ times a force and limited to ± 120 N m.

Where the Controller Runs

In industrial control, loops are often nested or cascaded, with cycle times from milliseconds to minutes, and programmable logic controllers serve as local controllers within a supervisory control scheme [24]. The chapter's structure is of that kind. The speed master hands its filtered acceleration to both winders, each winder loop carries its feedforward torque as the bias of a PI trim in newtons, and the pump runs as a ratio to the measured line speed, all on a 5 ms cycle. The optional predictive controller plans the rewind force and the pump speed every 0.05 s above those loops. Figure 5.6 draws the loops on their devices.

IN PRACTICE **Where the Load-Cell Roller Goes**

A load-cell roller reads tension cleanly where the web wraps it at a fixed angle. Next to a coiler the wrap angle changes as the coil builds, so the usual solution is a second roller, a little farther from the coil, over which the web passes at a fixed wrap. A single roller beside the coil can serve instead if its signal is corrected with the current coil diameter, which most coiler drives already measure or compute [23]. Load cells read every force on their structure, and vibration from nearby machinery puts oscillating forces on them [22]. The chapter's WT2 sits on an idler in span 2 and is read as the tension itself.

What Is Modernizing It

Signal processing and disturbance estimation are one direction. On an industrial roll-to-roll system, a disturbance observer gave robustness to the changing inertia of the web roll and to disturbances from web slip, a feedforward cut the tracking error where the tension setpoint changes, and Kalman filtering of the tension sensor greatly improved the control performance [18]. Constrained predictive control is the second: a design of 2023 wrote the rated motor torque into its constraints and updated

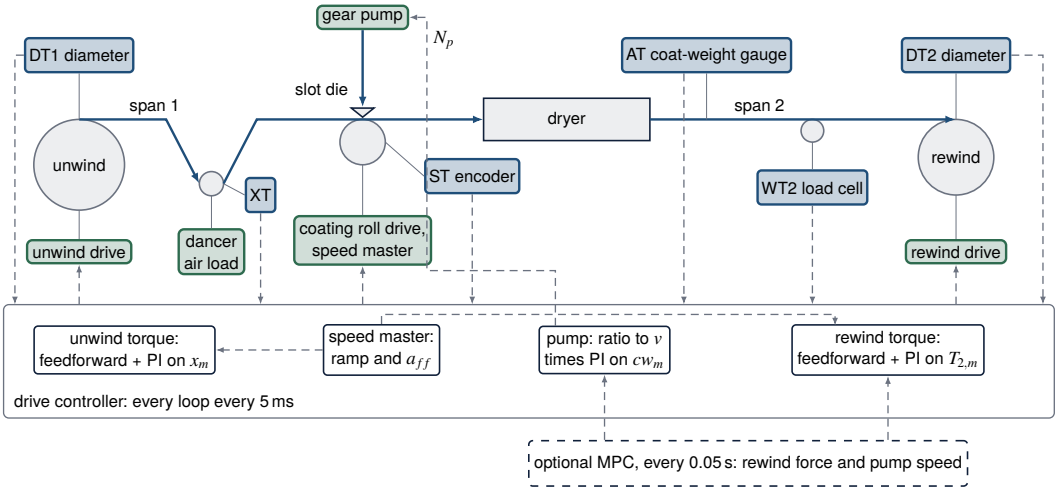

Figure 5.6 The loops on line hardware. Every sensor and drive is one of the chapter's: DT1 and DT2 read the roll diameters, XT the dancer, ST the line speed, AT the dry coat 19 m after the die and WT2 the span 2 tension. One drive controller runs all loops every 5 ms: the speed master hands its ramp and filtered acceleration to the winder feedforwards, and the pump runs as a ratio to the line speed under a slow trim on the late gauge. Notice that only the dancer's air load sets a tension directly, and that the optional MPC replaces the pump and rewind loops but not the unwind.

its model from the measured or estimated roll radius [14]. The chapter's own predictive controller carries both ideas, a Kalman filter with an integrating disturbance on the load cell and the torque limit as a hard bound, and Section 5.6.7 sets it against the loops case by case.

5.8 Limitations

The spans are lumped springs whose damping scales with their stiffness, and the stability limit of the load-cell trim in Section 5.5 rests on that scaling, so a line damped mostly by idler bearings and air would place it elsewhere. Tension waves, roller slip and the capstan limit, backlash, and the mechanics of the wound roll, among them wound-in tension, telescoping and taper tension, are outside the model. The coating obeys the pre-metered law at every speed, with no coating window, edge build-up or cross-web profile, and viscosity reaches the coat weight only through pump slip [5, 6]. The dryer is one time constant and the gauge a point sensor on the dry solids.

> ⚠ **OPEN CHALLENGE** **Seeing What the Electrode Coat Becomes**
>
> The link between the physical processes of electrode manufacture and its outputs is not yet fully understood, partly because so many parameters are involved and complete data sets under different processing conditions are lacking. A 2021 review of metrology for slurry coating found a range of measurements that need little improvement, with in-line particle sizing the exception, but few studies that bring them together, which multiphysics and data-driven models need [15]. The chapter's gauge is one point sensor on the dry solids. Progress would be in-line measurements joined into one picture of the coat as it is made.

Switching the product during a run re-threads the line at the new recipe's steady state instead of splicing the new web in, so a product change cannot be run as a transient. The break limits, the wrinkle threshold and the scrap rules are illustrative design choices of the application, and its scrap counts compare controllers with one another rather than predicting a real line's yield. The three reduced models of Section 5.5 are checked against the simulator and not against a physical line, so they establish what the teaching model does rather than what a given line would do.

The predictive controller's model is the line at its declared constants and recipe, linearized at the measured radius and speed. It does not know the stiffness of a spliced web, which reaches it only through the load cell's force estimate, nor the true viscosity, which reaches it only through $\hat{g}$ one gauge delay later. It leaves out span 1 and the dancer, the bead's 2.5 ms lag at 120 m/min, the growth of the radius within a plan, and the stretch of span 2, which moves the gauge's place on the film web by about 13 mm. It also takes the width, the pump displacement and the other constants at their declared values while the shipped feedforward reads the run's, so a reader's change to those favors the loops. The tension band carries one slack over the whole horizon rather than one per step, and the weights, chosen once on case 9, lose case 5.

OPEN CHALLENGE **Tension Tracking While the Web Changes**

Roll-to-roll lines still make products less consistent than batch processes, and tension-induced distortion is among the root causes. Coating and drying change the elasticity of the substrate, the radius and inertia of the rolls change all the time, and the rated torque of the motors must not be exceeded. A study of 2023 met these with a robust predictive controller whose parameters are scheduled, shown in simulation of an actual line [14]. Progress would be tension tracking that stays precise as the web's elasticity changes, shown on a running line.

5.9 Conclusion

This chapter put a torque-driven winder pair, a dancer, a pre-metered slot die and a late gauge on one speed, so that each remedy for a speed change could be priced separately. A ramp from 120 to 180 m/min with the pump flow held makes the coating 33.2 % thin at the die and costs 208.90 m of film, and the pump feedforward and inertia compensation together cut that to 12.80 m. A reduced third-order model of the dancer loop reproduces the dancer travel and span tension through that ramp and shows why a speed-ratio radius estimate destabilizes the foil line at 24 m/min. A mass-on-a-spring model of the rewind explains the break of a softer spliced web at 37.85 s and places the hot trim's stability limit at 0.80 times the PET stiffness once the 5 ms control interval is counted. The coat-weight trim tuned at 120 m/min has a gain margin of 0.93 at 30 m/min after a thin batch, and a batch change behind the 19 m gauge costs the electrode line 20.15 to 20.35 m whatever the integral time.

The optional model predictive controller scraps 21.05 m against the loops' 63.00 m when a thinner batch and a faster order reach the film at 30 m/min, a gain owed to its model of the pump's lag and its delay counted in meters, which loops with a 25 s trim and a lead on the pump match with 20.55 to 20.80 m.

The model omits the coating window, roller slip, the mechanics of the wound roll and any damping not proportional to the web's stiffness. A fuller model would add a taper-tension schedule, a second coating station with register control, an estimator that reads the flow into the die ahead of the gauge, and a speed-scheduled trim. On the app, feedforward, compensation and delayed feedback are learned by running them against one another.

KEY POINTS

- One speed change thins the coating at once and upsets the tension, because the winders must spin up rolls that change every turn.
- The pump feedforward and inertia compensation cut the scrap of a ramp from 120 to 180 m/min from 208.90 to 12.80 m.
- A softer spliced web, not a stiffer one, breaks the hot load-cell trim.
- The predictive controller's gain on case 9 comes from its model of the pump's lag and its delay in meters.

Where else this applies. Volume 16, Chapter 10 weighs its coating through a delay set by the line speed.

References

[1] K. P. Grenfell, "Tension control paper-making and converting machinery," *IEEE Transactions on Applications and Industry*, vol. 83, no. 73, pp. 234–240, 1964.

[2] D. P. D. Whitworth and M. C. Harrison, "Tension variations in pliable material in production machinery," *Applied Mathematical Modelling*, vol. 7, no. 3, pp. 189–196, 1983.

[3] P. R. Pagilla, N. B. Siraskar, and R. V. Dwivedula, "Decentralized control of web processing lines," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 1, pp. 106–117, 2007.

[4] K. J. Ruschak, "Limiting flow in a pre-metered coating device," *Chemical Engineering Science*, vol. 31, no. 11, pp. 1057–1060, 1976.

[5] X. Ding, J. Liu, and T. A. L. Harris, "A review of the operating limits in slot die coating processes," *AIChE Journal*, vol. 62, no. 7, pp. 2508–2524, 2016.

[6] M. Schmitt, M. Baunach, L. Wengeler, K. Peters, P. Junges, P. Scharfer, and W. Schabel, "Slot-die processing of lithium-ion battery electrodes—coating window characterization," *Chemical Engineering and Processing: Process Intensification*, vol. 68, pp. 32–37, 2013.

[7] A. Kwade, W. Haselrieder, R. Leithoff, A. Modlinger, F. Dietrich, and K. Droeder, "Current status and challenges for automotive battery production technologies," *Nature Energy*, vol. 3, no. 4, pp. 290–300, 2018.

[8] H. Koc, D. Knittel, M. de Mathelin, and G. Abba, "Modeling and robust control of winding systems for elastic webs," *IEEE Transactions on Control Systems Technology*, vol. 10, no. 2, pp. 197–208, 2002.

[9] P. R. Raul and P. R. Pagilla, "Design and implementation of adaptive PI control schemes for web tension control in roll-to-roll (R2R) manufacturing," *ISA Transactions*, vol. 56, pp. 276–287, 2015.

[10] P. M. Schweizer, "Control and optimization of coating processes," in *Liquid Film Coating*, S. F. Kistler and P. M. Schweizer, Eds. Dordrecht: Springer Netherlands, 1997, pp. 735–768.

[11] N. A. Ebler, R. Arnason, G. Michaelis, and N. D'Sa, "Tension control: dancer rolls or load cells," *IEEE Transactions on Industry Applications*, vol. 29, no. 4, pp. 727–739, 1993.

[12] 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.

[13] R. Carrasco and M. A. Valenzuela, "Tension control of a two-drum winder using paper tension estimation," *IEEE Transactions on Industry Applications*, vol. 42, no. 2, pp. 618–628, 2006.

[14] Z. Chen, B. Qu, B. Jiang, S. R. Forrest, and J. Ni, "Robust constrained tension control for high-precision roll-to-roll processes," *ISA Transactions*, vol. 136, pp. 651–662, 2023.

[15] C. D. Reynolds, P. R. Slater, S. D. Hare, M. J. H. Simmons, and E. Kendrick, "A review of metrology in lithium-ion electrode coating processes," *Materials & Design*, vol. 209, p. 109971, 2021.

[16] D. L. Wood, J. Li, and C. Daniel, "Prospects for reducing the processing cost of lithium ion batteries," *Journal of Power Sources*, vol. 275, pp. 234–242, 2015.

[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] H. Hwang, J. Lee, S. Eum, and K. Nam, "Kalman-filter-based tension control design for industrial roll-to-roll system," *Algorithms*, vol. 12, no. 4, p. 86, 2019.

[19] D. Bresch-Pietri, J. Chauvin, and N. Petit, "Prediction-based stabilization of linear systems subject to input-dependent input delay of integral-type," *IEEE Transactions on Automatic Control*, vol. 59, no. 9, pp. 2385–2399, 2014.

[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] D. Mohanty, J. Li, R. Born, L. C. Maxey, R. B. Dinwiddie, C. Daniel, and D. L. Wood, III, "Non-destructive evaluation of slot-die-coated lithium secondary battery electrodes by in-line laser caliper and IR thermography methods," *Analytical Methods*, vol. 6, no. 3, pp. 674–683, 2014.

[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] J. W. Turley, "Web tension measuring device for use with web coiling equipment," U.S. Patent 5,275,062, assigned to T. Sendzimir, Inc., 1994, filed August 9, 1991, issued January 4, 1994.

[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.

5.10 Exercises

5.1 Concept check · 2 min

In case 5 a batch of cathode paste at half the viscosity reaches the electrode line at 20 s, with the line at 24 m/min. With coat-weight trim integral times of 15, 30 and 60 s the run scraps 20.15, 20.25 and 20.35 m of foil. What sets most of that scrap?

(a) The integral time of the trim, because the fastest trim overshoots at the die

(b) The foil already coated between the die and the gauge before the gauge sees the change

(c) The dryer's time constant, which leaves solvent in the foil while the batch changes

(d) The load-cell noise, which flags tension defects on the foil during the batch change

5.2 Concept check · 1 min

A rewind tension trim tuned hot on the PET film ($K_c = 1$, $\tau_I = 0.5\,\mathrm{s}$) holds span 2 within 2.94 %. The next roll will be spliced in with a different web. According to the mass-on-a-spring model of Section 5.5, which splice is the dangerous one?

(a) A softer web, but only because it stretches further before it breaks

(b) Neither, because the trim works in newtons and is multiplied by the radius

(c) A softer web, because the shortest stable integral time grows as the stiffness falls

(d) A stiffer web, because the tension rises faster for the same speed mismatch

5.3 Concept check · 2 min

On case 9 the predictive controller scraps 21.05 m and the loops with both feedforwards 63.00 m. Which result from Table 5.8 locates the source of the difference?

(a) With the speed preview off the plan scraps far more, so the preview carries the gain

(b) Without the soft tension band the plan scraps more, so the constraint carries the gain

(c) No retuning brings the loops near the plan, so the optimization carries the gain

(d) Without the pump lags, or with the delay counted in time, the plan scraps far more

5.4 Theory and calculation · 6 min

The film recipe coats $12\,\mathrm{g/m^2}$ of dry adhesive with a solids fraction $\varphi = 0.45$ and a wet density $\rho_{wet} = 1050\,\mathrm{kg/m^3}$ on a web $W = 1.0\,\mathrm{m}$ wide. The gear pump displaces $Q_{max} = 6.0\,\mathrm{L/min}$ at full speed and delivers $\eta_{nom} = 0.96$ of it at the recipe viscosity.

(a) Compute the wet film $h_{w,SP}$ the die must lay down. *Checkpoint:* 25.4 µm.

(b) Compute the fastest line speed at which the pump can still coat the target at full speed, and the speed at which the capacity limit of Equation (5.10) holds the line. *Checkpoint:* 226.8 m/min and 215.5 m/min.

(c) Case 7 orders 240 m/min. Compute the coat weight the pump makes there at full speed and how light it is, and say whether the coat-weight trim can restore the target. *Checkpoint:* $11.34\,\mathrm{g/m^2}$, 5.5 % light. It cannot.

5.5 Theory and calculation · 5 min

A full film roll on the unwind has $R_u = 300\,\mathrm{mm}$ on a core of $R_{core} = 76\,\mathrm{mm}$, a web density $\rho = 1390\,\mathrm{kg/m^3}$, a width $W = 1.0\,\mathrm{m}$ and $J_{core} = 0.50\,\mathrm{kg\,m^2}$. The span 1 setpoint is 200 N, the bearing friction $b = 0.02\,\mathrm{N\,m\,s}$, and at 120 m/min the roll turns at 6.661 rad/s.

(a) Compute the roll's inertia from Equation (5.2) and its equivalent mass $m_u = J_u/R_u^2$ at its surface, and compare it with the mass of the film on the roll. *Checkpoint:* $18.11\,\mathrm{kg\,m^2}$, 201 kg, against 368 kg of film.

(b) Compute the holdback torque the unwind motor applies at steady speed. *Checkpoint:* 59.87 N m.

(c) The line ramps at 12 m/min per s with inertia compensation off. From Equation (5.17), find the steady dancer offset the proportional term alone would need to supply the accelerating force, with the dancer loop gain of 0.4 N/mm. *Checkpoint:* 101 mm, against stops at 150 mm.

5.6 Theory and calculation · 8 min

The rewind of case 4 has an equivalent mass $m_r = 61$ kg at 128 mm. Span 2 is $L_2 = 24$ m of PET with $EA = 225$ kN and a retardation time $\tau_w = 5$ ms. The hot load-cell trim has $K_c = 1$ and $\tau_I = 0.5$ s.

(a) From Equation (5.19), compute the critical integral time on the PET and on a web of 0.3 times its stiffness, and say which of the two the hot trim holds. *Checkpoint:* 0.33 s and 1.08 s. It holds the PET only.

(b) Find the stiffness, as a multiple of the PET's, at which $\tau_I = 0.5$ s is just critical by the same equation, and compare it with the two limits of Figure 5.2. *Checkpoint:* 0.65, between 0.59 and 0.80.

(c) Estimate the period of the tension oscillation on the softer web, and compare it with the 0.651 s measured on the engine. *Checkpoint:* 0.654 s.

5.7 App challenge · 8 min

Load scenario 2, Speed ramp, feedforward off, and change no setting. The order from 120 to 180 m/min arrives at 10 s and ramps at 12 m/min per s. Set the playback speed to 20×, run to 180 s with the default seed and export the CSV file. Read the scrap in the last row of `metric_scrap_m` and the Worst tension deviation tile at 180 s.

(a) Read the lowest coat weight at the die, column `cw_die_gm2`, and compare it with the pre-metered law for a flow held while the speed rises by half. Read the scrap. *Checkpoint:* 8.02 g/m^2 (within 0.02) against 8.0 g/m^2, and 208.90 m (within 0.1 m).

(b) Press Reset, set Line speed feedforward to the pump to On before Run, and run to 180 s. Read the scrap and the Worst tension deviation tile. *Checkpoint:* 22.30 m and 16.6 %.

(c) Press Reset, set Inertia compensation during ramps to On as well, and run to 180 s. Read the same two values and say what the remaining scrap is. *Checkpoint:* 12.80 m and 1.6 %.

5.8 App challenge · 8 min

Load the challenge, Challenge · order change on foil, and change no setting. The electrode order goes from 24 to 42 m/min at 20 s with a ramp of 30 m/min per s, and both feedforward switches are Off. Set the playback speed to 20× and run the full 200 s with the default seed. Read the Worst coat-weight deviation tile and, for two decimals, the last row of `metric_scrap_m` in the exported CSV file.

(a) Read the scrap and the coat-weight deviation as loaded, and explain why the coat-weight trim cannot recover the order. *Checkpoint:* 125.85 m (within 0.1 m) and 34.7 %.

(b) Press Reset, set Line speed feedforward to the pump to On and Speed ramp rate to 6 m/min per s before Run, leave Inertia compensation during ramps Off, and run 200 s. Is the target met? *Checkpoint:* 12.05 m, met.

(c) Press Reset, set the ramp rate back to 30 m/min per s and Inertia compensation during ramps to On, and run 200 s. Compare with the pump feedforward alone at that ramp in Table 5.7. *Checkpoint:* 2.55 m, met, against 36.10 m.

5.9 App challenge · 8 min

Load scenario 9, A slow line, a thinner batch and a faster order: MPC against the loops with feedforward, and change no setting. The predictive controller is selected. Set the playback speed to

20× and run to 300 s with the default seed. Read the scrap in the last row of `metric_scrap_m` of the exported CSV file after each run.

(a) Read the scrap under the MPC. Then press Reset, set Tension and coat weight to Loops before Run, and run to 300 s again. *Checkpoint:* 21.05 m and 63.00 m (within 0.1 m).

(b) Reload the scenario, switch to Advanced mode, set MPC coat-weight estimate time to 2 s before Run and run to 300 s. Reload again and repeat with 30 s. *Checkpoint:* 19.65 m and 24.30 m.

(c) The estimate time moves the scrap by a few meters. Use Table 5.8 to say which parts of the plan carry the remaining 40 m or so, and which two changes to the loops close the gap. *Checkpoint:* The pump-lag model and the delay in meters. A 25 s trim and a lead on the pump.

5.10 Application area · 3 min

An electrode line coats 200 g/m^2 of cathode paste with $\varphi = 0.70$ and $\rho_{wet} = 2300\,\mathrm{kg/m^3}$ on foil 1.0 m wide, through the pump of Table 5.2 and a 15 m dryer. The residual solvent of Section 5.3 must fall below 1 % of the dry coat, with $\tau_{dry} = 3.0\,\mathrm{s}$.

(a) Compute the fastest speed the dryer allows and the fastest speed the pump can coat at full speed. Which one caps the line? *Checkpoint:* 79.8 m/min and 46.4 m/min. The pump.

(b) Production asks for 60 m/min. What pump displacement would coat it at full speed at the recipe slip? *Checkpoint:* 7.76 L/min.

5.11 Application area · 4 min

A drive engineer proposes to remove the ultrasonic diameter sensor from the unwind and estimate the radius from the ratio of the line speed to the unwind speed, a common practice. Section 5.5 shows that the estimate turns the damping of Equation (5.17) into $c_d/2 - 2T_{1,SP}/v_0$, with $c_d = 1000\,\mathrm{N\,s/m}$.

(a) Compute the damping for the film ($T_{1,SP} = 200\,\mathrm{N}$) at 120 and 30 m/min, and for the foil ($T_{1,SP} = 120\,\mathrm{N}$) at 24 m/min. *Checkpoint:* $+300$, -300 and -100 N s/m.

(b) Below what line speed does each product lose its damping, and what do you advise the engineer? *Checkpoint:* 48 m/min for the film, 28.8 m/min for the foil.

5.12 Review the agent · 5 min

An AI assistant is asked how to tune the coat-weight trim of the film line for a heavy order that will run at 30 m/min. It answers: “The recipe trim, gain 0.3 and integral time 6 s, has a gain margin of 2.87 on this line, with the gauge 19 m from the die. The gauge and the loop do not change with the order, so the trim is safe at 30 m/min. To correct a thinner batch faster at the lower speed, raise the gain to 0.5.” Find the flaw, and say how you would check the advice on the app before using it.

CHAPTER 6

Orbit Feedback for a Storage Ring Whose Weak Modes Amplify Noise

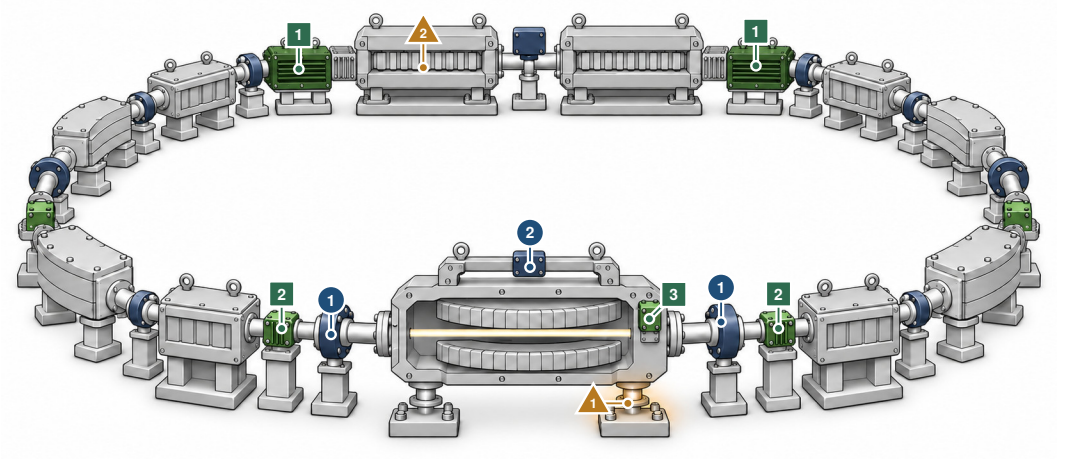

The controller moves 1 slow correctors 2 fast correctors 3 device end coils (gap feedforward)

The controller reads 1 beam position monitors 2 gap readback

What disturbs it 1 girder drift and vibration 2 insertion device kick as its gap closes

Ninety-six monitors and two corrector sets hold the beam round the whole ring, while the light is made inside the insertion devices, where no monitor sits. Illustration, not to scale.

IN PLAIN TERMS Ninety-Six Sensors Steering One Beam

The main idea. A storage ring keeps a beam of electrons on one closed path so that it can give off bright light for experiments. Small steering magnets around the ring push the beam back wherever ninety-six position sensors see it drift.

The challenge. Some steering patterns move the beam a lot, but others barely move it, so correcting those takes large magnet swings for tiny gains. Sensor noise in those weak patterns turns into magnets that never sit still.

Why it matters. Choosing which patterns to correct decides how steady the light is and how quietly the magnets work.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the two histories in Section 6.2, the fact in the Introduction, the try-its and the fact of Section 6.6 and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 6.6), Technology in Practice (Section 6.8), Limitations (Section 6.9) and the two open challenges

The control question A storage ring's 96 by 96 orbit response is known exactly, yet its singular values span a factor of 405, so a full inverse turns monitor noise into magnet motion. How many directions should the feedback correct, and what must a fast corrector set, a feedforward and a monitor check add?

Learning objectives

- Compute the optics of a thin-lens FODO lattice and the closed-orbit response of one corrector from its β and the tune.
- Explain why the full inverse amplifies monitor noise, and compute the filter factors of a truncated and a Tikhonov-regularized inverse.
- Place the truncation on the measured trade between residual orbit and corrector jitter, and say why it falls at the spectrum's cliff.
- Distinguish what feedback, offloading, a gap feedforward and a residual test each contribute to the graded challenge.
- Show from the loop-gain eigenvalues why a corrector wired backward defeats the full inverse at every bandwidth.

Prerequisites Sections 1.4.7 (singular values and directionality), 1.4.6 (a static inverse of the gain matrix), 1.6.1 (engine and seed) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/synchrotron-orbit/

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control, a fast handle kept centered by a slow one. Volume 5, Chapter 11 for a controlled temperature that no instrument reaches, as no monitor sits at a source point, and Volume 5, Chapter 12 for a sensor whose reading cannot be taken at face value.

6.1 Introduction

One corrector magnet in a storage ring, asked for a kick of 5 µrad, moves the electron orbit at every monitor in the ring. In the teaching ring of this chapter, corrector 37 sits at a focusing quadrupole, and its kick moves the closed orbit by up to 74.8 µm at monitor 85, far round the ring, and by 40.0 µm rms over all 96 monitors. The same kick at corrector 38, beside a defocusing quadrupole, moves it at most 24.8 µm. Each column of the 96 by 96 response matrix is a pattern of this kind, a cosine whose sign changes 27 times round the 96 monitors, and orbit feedback corrects the beam by inverting the whole matrix [1].

The simulator is built against a natural belief, that a plant model known to machine precision should be inverted exactly. The response of this ring is exact by construction, and its singular values, the gains along the strong and weak directions of Section 1.4.7, still fall from 379 to 0.94 µm/µrad. A controller that divides by all of them turns monitor noise into magnet motion. On the noise of case 2, 1 µm per sample, the full inverse makes the slow correctors jitter 0.198 µrad about their own mean, twelve times the 0.0164 µrad of an inverse truncated to the 48 strongest modes. Truncation and Tikhonov regularization, the classical answers to an ill-posed inverse [2], trade that noise against the part of the orbit the weak modes would have corrected, and the ring shows the trade one mode at a time.

> **+ DID YOU KNOW? Corrections by the Thousand Each Second**
>
> One light source's orbit feedback, in routine operation for more than 20 years, computes corrections 1.6 thousand times a second from 160 monitors to 38 correctors in each plane. The system planned for the ring's upgrade computes them 22.6 thousand times a second, from 560 monitors to 160 correctors [3].

The same shape appears wherever many actuators act on one field through a matrix whose directions differ greatly in gain. A high-purity distillation column is the classic ill-conditioned plant of process control, where the physics of its poor conditioning decides how much performance a robust controller can keep under actuator uncertainty [4], and the cross-direction profile of a paper

machine is a spatially distributed plant whose controllers are designed in spatial as well as temporal frequency [5]. Two more features of orbit feedback are general. Two actuator sets of different speed and strength share one controlled field, which is the mid-ranging problem of process plants [6], and a failed sensor must be found among many by the residual its readings leave.

The orbit feedback literature is written for accelerator physicists working on real machines, from the decomposition of a measured response [7] to controllers designed per singular mode [8, 9], and a control student rarely meets a hundred-square plant whose modes can be switched on and off one at a time. This chapter traces what an inverse, its truncation or regularization, two corrector sets, a gap feedforward and a residual test each do to the true orbit at the photon source points of a declared storage ring simulated in a browser. The closed-orbit response, the singular value decomposition, Tikhonov regularization, internal model control and mid-ranging are established methods. Every result is a measurement made on this one ring, and two of them overturn the plan the simulator was built from. No lattice error the simulator allows makes the full inverse diverge, so its infeasible case is a corrector wired backward, and the regularization in the graded challenge is not what lets it pass.

The chapter is organized as follows. Section 6.2 follows orbit correction from strong focusing to feedback in mode space. Section 6.3 states the ring: its optics, a closed-form response that agrees with independent tracking at all 11,520 entries, the correctors, the disturbances and what the monitors report. Section 6.4 splits the response into strong and weak directions and shows where the disturbances lie among them. Section 6.5 describes the controllers: a one-sided Jacobi decomposition, a per-mode controller whose closed loop is derived and then measured, the fast set and offloading, the gap feedforward and the residual test, behind a line between measurement and truth that no controller crosses. Section 6.6 measures the ten cases, each credited in the app by a completion rule, and Section 6.8 describes the real monitors, correctors and controllers. Sections 6.7, 6.9 and 6.10 place the ring among the volume's plants, state its limits and summarize. The exercises follow in Section 6.11.

6.2 From Strong Focusing to Feedback in Mode Space

Storage rings focus their beams by alternating gradients. The strong-focusing principle of 1952 showed that a train of lenses that alternately converge and diverge, with equal strengths, focuses on balance, and that the oscillation amplitude and the magnet aperture of a synchrotron shrink with it [10]. The theory that followed wrote the transverse motion with an amplitude function and a betatron phase [11], the β and φ in which every closed-orbit calculation is still written. A dipole error anywhere in such a lattice displaces the closed orbit round the whole ring, by an amount set by β at the error and at the observer and by the tune, which is the response this simulator computes [1].

The first history segment follows the 1993 decomposition of the response matrix into channels. The numerical tools were older. A stable algorithm for the singular values and the pseudo-inverse was published in 1965 [12], and the one-sided Jacobi method, which never forms the product of the matrix with its transpose, was later shown to compute the singular values with high relative accuracy [13]. Because a model's response is only as good as the model, orbit responses measured on the machine were then used to determine a storage ring's optics [14].

> ⧖ HISTORY **1993: The Orbit Split Into Channels**
>
> In 1993 the correction of a storage ring's whole orbit was written as a singular value decomposition of its response matrix, and its use was simulated for a third-generation light source that was to carry 320 corrector magnets and 360 beam position monitors. The decomposition turns the monitors and the correctors into the same number of transformed monitors and transformed correctors. Each transformed monitor is coupled to at most one transformed corrector, with a strength that sets how efficiently that channel corrects, and the channels do not interact. Decoupling the weakly coupled pairs makes the correction more efficient at the expense of accuracy, and it lets the corrector strengths be kept within what their power supplies can deliver [7].

Table 6.1 Variables by Role, With Their Values When a Case Starts

Variable	Role	Start	Range
x_d, x'_d	CV, source point d (truth)	0	±3 µm, ±1 µrad
y (rms)	CV, true orbit at the 96 BPM places	0	≤ 2 µm (challenge)
y_m	measurement, 96 BPM readings at 1 kHz	0	noise σ_n
g_d	measurement, gap readbacks	40 mm	10–40 mm
$\theta_{s,cmd}$	MV, 96 slow commands	0	±50 µrad (5 A)
$\theta_{f,cmd}$	MV, 24 fast commands	0	±5 µrad (5 A)
θ_{ff}	MV, end coils of the devices	0	from the table
θ_s, θ_f	states, kicks after the lags	0	—
δ_g	DV, girder drift (24)	0	0–0.41 µm/$\sqrt{\text{s}}$
v_g	DV, girder vibration	0	0–2 µm, 10–50 Hz
r	DV, dipole ripple at 50 Hz	0	0–0.2 µrad
κ_d	DV, device kicks	0	to 1.6, −2.2, 1.9 µrad
fault	DV, BPM 37	none	+60 µm stuck, or noisy
ν, ε_k	lattice tune, quadrupole error	14.2, 0	14.05–14.35, 0–2 %

The users set the target. The stability an experiment needs depends on how it samples, integrates and scans, over time scales from microseconds to hours, and at third-generation rings it can come to micron and submicroradian stability of the electron orbit [15]. The second history segment follows fast orbit feedback to a controller per singular mode. Internal model control came from the process industries, where its structure was reviewed and unified in 1982 [16]. Two actuator arrays acting on one beam were decomposed by their controllable subspaces, with mid-ranging where the arrays' control directions align [9], an internal-model form of the mid-ranging that process plants use when a quick, small handle and a slow, large one share a single controlled variable [6, 17]. Rate limits on the corrector supplies brought an anti-windup problem of their own [18].

> ⧖ **HISTORY** **2013: A Controller for Every Mode**
>
> Fast orbit feedback splits its work into space and time. The spatial part inverts a response matrix that is normally ill-conditioned, so the less significant singular values are removed or filtered, and the usual dynamics are proportional-integral. Most systems apply the same controller to every mode with a gain per mode, because separate dynamics per mode cost two more matrix products, too heavy for most storage-ring systems with more than a hundred actuators and update rates above a few kilohertz. A study of 2013 gave each mode its own internal-model controller, chosen from the disturbance spectrum in that mode, and ran it on a booster synchrotron of 22 sensors and 22 actuators operated with stored beam, the test bed for a storage ring with 172 of each [8].

6.3 A Declared Ring of 96 Monitors and 120 Correctors

The simulated ring is a declared teaching lattice, not a model of any facility, as shown in Figure 6.1. It has 24 cells of 20 m, 480 m round, each carrying two periods of a thin-lens FODO lattice: a focusing quadrupole (QF), a drift of $L = 5$ m, a defocusing quadrupole (QD) and another drift. A beam-position monitor (BPM) sits 0.4 m downstream of each of the 96 quadrupoles and a slow corrector at each quadrupole. One fast corrector sits at 12.5 m in each cell, two dipoles at 2.5 and 7.5 m, and in cells 3, 11 and 19 an insertion device occupies the middle of the straight from 15 to 20 m, where its photon source point lies. Only the horizontal plane is modeled. The role of each variable is given in Table 6.1, and the Ring view of the simulator, which draws the same ring live, is shown in Figure 6.2.

6.3.1 Optics of the Teaching Lattice

For thin quadrupoles of focal length $\pm f$ separated by drifts L, the phase advance μ of one FODO period and the extreme values of the amplitude function are given by Equation (6.1).

$$\sin\frac{\mu}{2} = \frac{L}{2f}, \qquad \mu = \frac{2\pi\nu}{48}, \qquad \beta_{\pm} = \frac{2L\,(1 \pm \sin(\mu/2))}{\sin\mu} \tag{6.1}$$

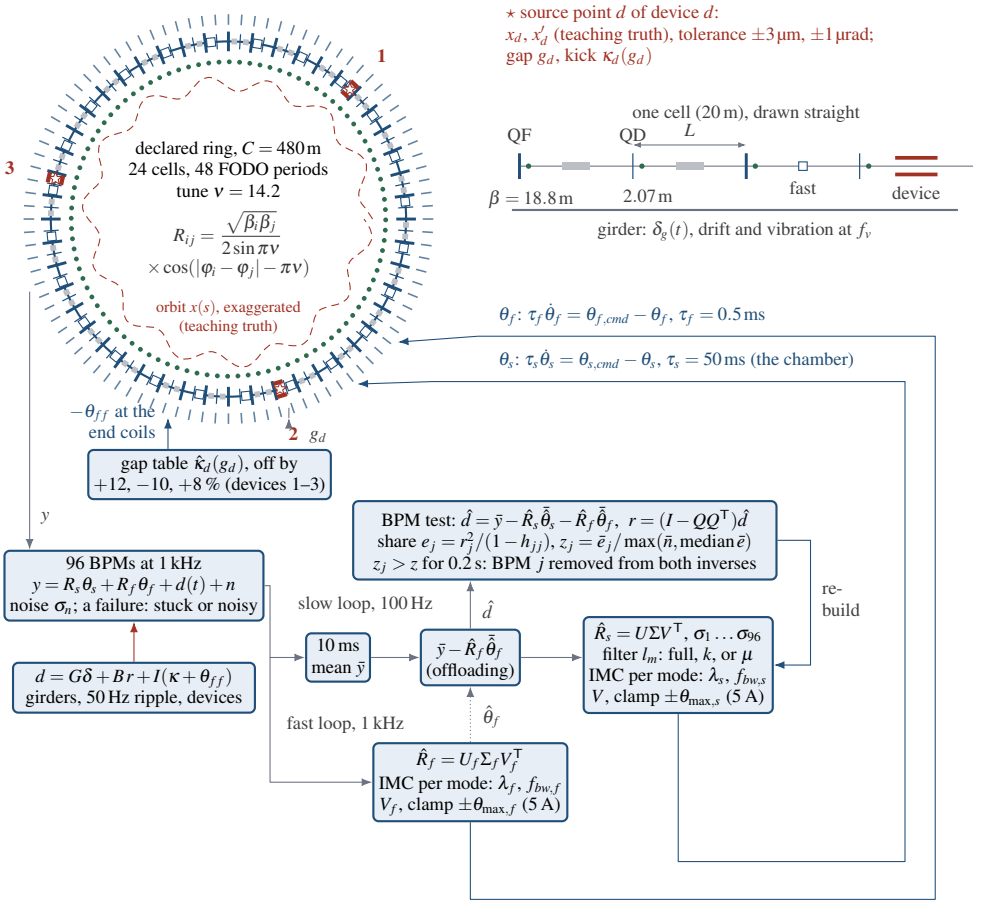

Figure 6.1 The declared ring as modeled, with what its controllers read and drive. The ring carries 24 cells (one is drawn straight at the right) with focusing quadrupoles thick and defocusing ones thin, dipoles shaded, the 96 BPMs as dots inside, the 96 slow correctors as bars outside and the 24 fast correctors as open boxes. The insertion devices 1 to 3 kick the beam by $\kappa_d(g_d)$ at their source points (stars). The dashed orbit and the source points' x_d and x_d' are teaching truth that no controller reads. The controllers see the BPM readings y and the gap readbacks g_d. The slow loop at 100 Hz and the fast loop at 1 kHz each invert their own response in mode space, offloading removes the fast set's share from the slow loop's input, the slow loop's disturbance estimate $\hat{d}$ feeds the BPM test, and the gap table feeds $-\theta_{ff}$ forward. Not to scale.

where ν is the horizontal tune, the number of betatron oscillations in one turn, and 48 is the number of periods. At $\nu = 14.2$ the period advances 106.5°, $f = 3.12$ m, and β is 18.79 m at a QF and 2.07 m at a QD. One period per cell would need 213°, past the 180° at which a FODO period stops being stable, which is why each cell carries two. The simulator does not use Equation (6.1) directly. It

builds the ring from transfer matrices and carries the optics round it, as shown in Equation (6.2).

$$\begin{aligned}
&\cos\mu_t = \tfrac{1}{2}(M_{11}+M_{22}), \quad \beta_0 = \frac{M_{12}}{\sin\mu_t} \\
&\alpha_0 = \frac{M_{11}-M_{22}}{2\sin\mu_t} \\
&\text{drift } \ell: \ \beta \leftarrow \beta - 2\alpha\ell + \gamma\ell^2, \quad \alpha \leftarrow \alpha - \gamma\ell \\
&\qquad\qquad \varphi \leftarrow \varphi + \arctan\frac{\ell}{\beta - \alpha\ell} \\
&\text{quadrupole: } \alpha \leftarrow \alpha \pm \beta/f, \quad \gamma = \frac{1+\alpha^2}{\beta}
\end{aligned} \tag{6.2}$$

where M is the one-turn product of the drift matrices $\left[\begin{smallmatrix} 1 & \ell \\ 0 & 1 \end{smallmatrix}\right]$ and the thin-lens matrices $\left[\begin{smallmatrix} 1 & 0 \\ \mp 1/f & 1 \end{smallmatrix}\right]$, taken from just before the first QF. The integrated phase gives $\nu = 14.200000$ and $\cos\mu_t = 0.309017$. A BPM 0.4 m after a QF sees $\beta = 16.46\,\text{m}$, one after a QD 2.42 m, and a fast corrector and a source point both 7.08 m. With the lattice error on, each quadrupole strength $1/f$ is multiplied by $1+\varepsilon_k g_q$, with g_q a fixed seeded unit normal, and the optics are carried round again without assuming periodicity.

6.3.2 The Closed-Orbit Response

A thin kick θ_j at a place with optics (β_j, φ_j) moves the closed orbit at a place i and changes its angle, as shown in Equation (6.3).

$$\begin{aligned}
R_{ij} &= \frac{\sqrt{\beta_i\beta_j}}{2\sin\pi\nu}\cos\Delta_{ij}, \qquad \Delta_{ij} = |\varphi_i - \varphi_j| - \pi\nu \\
R'_{ij} &= \frac{\sqrt{\beta_j/\beta_i}}{2\sin\pi\nu}\big[-\alpha_i\cos\Delta_{ij} - \operatorname{sgn}(\varphi_i-\varphi_j)\sin\Delta_{ij}\big]
\end{aligned} \tag{6.3}$$

where R is in m/rad, which is µm per µrad, and R' in µrad per µrad. The first line is the closed-orbit response of Equation 15.47 in Wiedemann's text [1]. From it the simulator builds R_s (96 BPMs by 96 slow correctors), R_f (96 by 24), the rows of the three source points' positions and angles, and three disturbance matrices. A girder displaced by δ_g moves its four quadrupoles, which kick the beam by $\pm\delta_g/f$, so the column c of G sums R_{iq}/f over the cell's QFs and $-R_{iq}/f$ over its QDs. The vector B sums the response to a unit kick at every one of the 48 dipoles, and I holds the response to a kick at each device's middle.

6.3.3 Correctors, Disturbances and the True Orbit

The orbit is algebraic in the kicks, because the beam settles on its closed orbit within a few hundred turns, a few microseconds. The kicks follow their commands through first-order lags, the slow ones through the eddy currents of the vacuum chamber they sit over, as shown in Equation (6.4).

$$\begin{aligned}
y &= R_s\theta_s + R_f\theta_f + d \\
d &= G\delta + B\,r + I(\kappa + \theta_{ff}) \\
\theta(t+\Delta t) &= \theta_{app} + \big(\theta(t) - \theta_{app}\big)\,\mathrm{e}^{-\Delta t/\tau} \\
\theta_{app} &= \operatorname{clip}(\theta_{cmd}, \pm\theta_{\max})
\end{aligned} \tag{6.4}$$

where τ is $\tau_s = 50\,\text{ms}$ for the slow set and $\tau_f = 0.5\,\text{ms}$ for the fast one. The commands are held between control executions, so the update is exact for any step. The disturbances run on their own

Table 6.2 Parameters and Defaults. A Is Adjustable in the Parameter Panel, H Hidden From It, S a Control or Disturbance Slider, F Fixed. All Are Teaching Values

Symbol	Meaning	Default (range)	
ν	horizontal tune	14.2 (14.05–14.35)	A
ε_k	rms quadrupole error	0 (0–2 %)	A
L, f	drift, focal length (derived)	5 m, 3.12 m	F
τ_s, τ_f	slow and fast coil lags	50 (10–200), 0.5 (0.1–2) ms	A
$\theta_{\max,s}$	slow limit, 5 A	50 (10–200) µrad	A
$\theta_{\max,f}$	fast limit, 5 A	5 (1–20) µrad	A
σ_n	BPM noise per 1 ms sample	0.5 (0–5) µm	A
κ_s	scale of the device kicks	1 (0–3)	H
q	ground drift rate	0 (0–0.41) µm/$\sqrt{\text{s}}$	S
A_v, f_v	girder vibration	0 µm (0–2), 10 Hz (10–50)	S
r_0	dipole ripple	0 (0–0.2) µrad	S
k, μ	modes kept, Tikhonov weight	48 (1–96), 10 (0.1–1000)	S
$f_{bw,s}$, $f_{bw,f}$	loop bandwidths	1 (0.1–10), 100 (10–300) Hz	S
z	BPM test threshold	100 (10–1000)	S
δ_m, N_m	response measurement	10 µrad, 200	F
	gap table error, devices 1–3	+12, −10, +8 %	F

clocks so that a run does not depend on the physics step, as shown in Equation (6.5).

$$\begin{aligned} w_g &\leftarrow w_g + q\sqrt{\Delta t_d}\,\varepsilon_g, \qquad \delta = Kw + \delta_{step} \\ v_g(t) &= A_v\,\mathrm{Re}\{p_g\,z(t)\,\mathrm{e}^{\mathrm{i}2\pi f_v t}\}, \qquad r(t) = r_0\sin(2\pi 50 t) \\ \kappa_d(g) &= \kappa_{\max,d}\left(\frac{g_{open}-g}{g_{open}-g_{\min}}\right)^2 \end{aligned} \tag{6.5}$$

where each of the 24 girders walks every $\Delta t_d = 10$ ms with rate q in µm/$\sqrt{\text{s}}$, and K smooths the walk along the ring with a row-normalized Gaussian of 3 cells, so that each girder's rms after a time T is $q\sqrt{T}$. The vibration has a fixed complex girder pattern p_g, strongest round cells 7 to 9, and an envelope $z = 1 + \zeta_1 + \mathrm{i}\zeta_2$ whose two parts are first-order colored noise of time constant 1 s and standard deviation 0.35, updated every 1 ms, which makes a narrow line at f_v. A burst multiplies A_v by 4 for 10 s. A ground step moves girders 6 to 9 by 3 µm. Each gap moves toward its target at 5 mm/s between $g_{open} = 40$ mm and $g_{\min} = 10$ mm, and $\kappa_{\max,d}$ is 1.6, −2.2 and 1.9 µrad for devices 1 to 3. The position and angle at each source point are the source-point rows applied to the same kicks and disturbances.

6.3.4 What the Monitors Report

Each BPM is read every 1 ms as the true orbit plus white noise of $\sigma_n = 0.5$ µm per sample (1 µm in case 2), drawn from the case's seed. A failed BPM 37 either jumps 60 µm from its last reading and freezes, or reads the true orbit plus 30 µm rms of noise. The gap readbacks are exact. The correctors' currents are the controllers' own commands, and the source points carry no monitor. The alternative to the model's response is a measured one, obtained as Equation (6.6).

$$\tilde{R}_{ij} = R_{ij} + \frac{\sigma_n}{\delta_m\sqrt{2N_m}}\,\varepsilon_{ij} \tag{6.6}$$

where each corrector would be kicked by $\pm\delta_m = 10$ µrad and the BPMs averaged over $N_m = 200$ samples each way, an element error of 0.0025 µm/µrad. The simulator computes it at once from the ring's true response and that noise, as a measurement on a quiet beam would read, and does not spend the 48 s of beam time a real one would take.

6.3.5 Clocks, Starting State and Numerics

The BPMs are read and the fast loop runs every 1 ms, the physics step, and the slow loop runs every 10 ms on the mean of the ten readings since its last update. A log row is written every 10 ms with

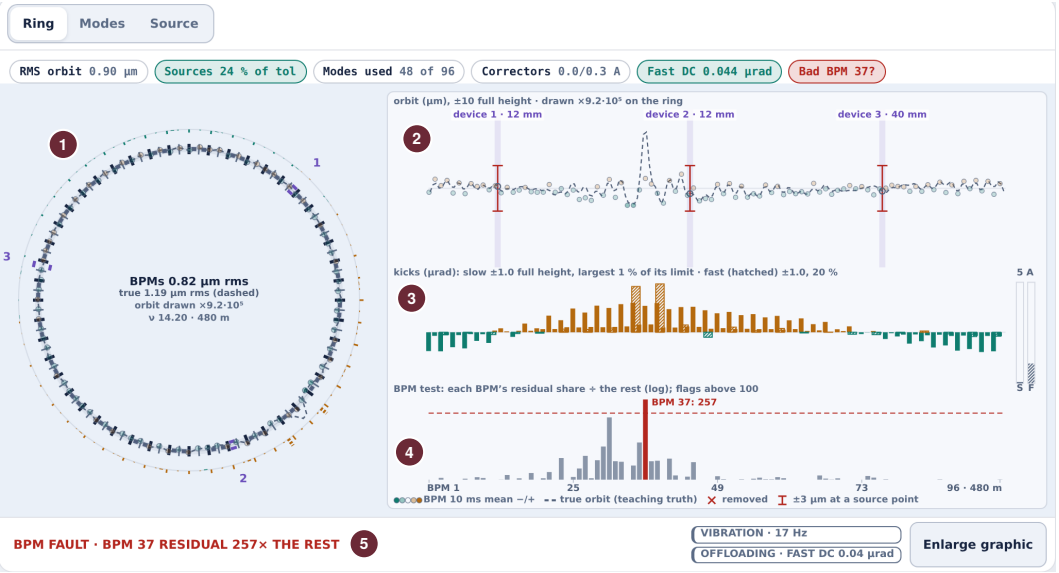

Figure 6.2 The challenge at $t = 300.19$ s on the Ring view with the suggested settings, 0.19 s after BPM 37 turned to noise and 0.02 s before the test removes it, cropped to the plant view. (1) The ring with its BPMs and correctors, the true orbit at 1.19 µm rms in its middle. (2) The unrolled strip with devices 1 and 2 closed to 12 mm and device 3 open at 40 mm, the BPM readings and the true orbit dashed. (3) The kick strip, with the fast set at 20 % of its limit and the slow set at 1 %. (4) The BPM test's strip, which points at BPM 37 with 257 times the median share against a threshold of 100. (5) The status line, which reads the fault. The challenge's results are in Table 6.5.

the minimum and maximum of the source-point positions over its ten steps, so that a 10 to 50 Hz vibration survives the log. Every case starts on the golden orbit, zero at every BPM, with every kick at zero except where the case sets one, and the case's seed drives every noise stream (Section 1.6.1). The parameters are listed in Table 6.2. The scripted events (a gap change, a burst, the BPM failure, the ground step, the lattice error) keep the engine's order of Figure 1.10. Each fires once, at the start of the first step at or after its time, and the controllers then run before the exact update.

6.3.6 Simplifications

The quadrupoles are thin lenses, the dipoles' weak focusing is left out of the optics, and only the horizontal plane is modeled, without coupling, dispersion from the correctors or energy change, so the model says nothing about vertical stability or beam size. Each corrector is a linear kick behind one first-order lag, whereas a real chamber's eddy currents have several time constants and real supplies have slew and rate limits [18]. BPM noise is white and independent, with no gain error or drift of its own. The tolerances at the source points (± 3 µm, ± 1 µrad) and of the orbit (2 µm rms) are teaching values, while a real light source derives its own from its experiments' needs [15]. The lattice error and the tune are bounded together. The seeded error raises the tune by 0.067 at 2 %, and beyond the slider ranges the true tune would reach the half-integer, where these optics have no stable solution, so that resonance stays out of reach.

6.4 Strong and Weak Directions of the Response

The singular value decomposition of Section 1.4.7, $R_s = U\Sigma V^{\mathsf{T}}$, writes the response as 96 independent modes. Column m of V is a pattern of corrector kicks, column m of U the orbit pattern it makes, and σ_m its gain, in µm per µrad. For this ring $\sigma_1 = 379.17$ and $\sigma_{96} = 0.9357$, a condition number of

405.2, and the 96 values fall in two groups of 48 with a cliff between $\sigma_{48} = 8.740$ and $\sigma_{49} = 1.282$, as shown in Figure 6.3a. The cliff comes from the β pattern. Modes 1 to 48 put a mean 89 % of their corrector weight on the 48 correctors at the QFs ($\beta = 18.8$ m), and modes 49 to 96 the same share on those at the QDs ($\beta = 2.07$ m). With every $\sqrt{\beta}$ replaced by the ring's mean and the phases kept, the ratio at modes 48 and 49 falls from 6.82 to 1.29 and the condition number to 192.8. Within each group the gain falls with the distance of the mode's betatron harmonic from the tune. Modes 1 to 6 are harmonics 13 to 15, modes 45 to 48 harmonics 0 to 2, and modes 91 to 96 harmonics 45 to 48. The 24 fast correctors' response R_f runs from 157 to 5.84 µm/µrad.

The disturbances do not spread evenly over these modes, and that decides what an inverse can buy. Projected on the orbit modes U, the girder drift as the plant makes it, smoothed along the ring, puts 86.0 % of its power in modes 45 to 48, the lowest harmonics of the strong group, 6.4 % in modes 1 to 44, 0.34 % in modes 49 to 92, and 7.3 % in the four weakest, modes 93 to 96. An inverse truncated anywhere from 48 to 92 modes therefore removes the same 92 to 93 percent of the drift's power, and the rest lies in the directions where each micrometer of orbit costs about a microradian of kick.

6.5 Feedback in Mode Space as Implemented

Every controller reads only a measurement packet: the 96 BPM readings with the BPM system's declared noise rating, the three gap readbacks, and its own corrector commands. A test of the application fails if a controller file reads the plant, the truth, the lattice error, the failed BPM's identity, the device kick scale or the event schedule, and a second test replays 320,000 recorded packets into a controller with no plant behind it and obtains every command again.

6.5.1 Filter Factors on One Decomposition

The controllers design with the model response $\hat{R}_s = U\Sigma V^{\mathsf{T}}$ of the declared lattice at the present tune (or the measured $\tilde{R}$), restricted to the BPMs in use, decomposed by the one-sided Jacobi method [13]. A pair of columns is rotated when $|w_p \cdot w_q| > 10^{-15}\|w_p\|\|w_q\|$, and the sweeps stop when one rotates nothing, which takes 14 sweeps for R_s and about 12 ms in a server-side JavaScript runtime. The inverse in force is a filter factor on each mode, for the full inverse, a truncation at k modes and Tikhonov regularization in turn, as shown in Equation (6.7).

$$l_m = 1, \qquad l_m = [m \le k], \qquad l_m = \frac{\sigma_m^2}{\sigma_m^2 + \mu} \tag{6.7}$$

where μ is in $(\mu\text{m}/\mu\text{rad})^2$, so a mode with $\sigma = \sqrt{\mu}$ is corrected halfway, which is the form of the mode-space paper [8] and of Tikhonov regularization in the filter-factor form [2]. At $\mu = 10$ the strongest modes are corrected by 1.00 of the way, mode 48 by 0.88, mode 49 by 0.14 and the weakest by 0.08.

6.5.2 Internal Model Control Per Mode

Each mode is a sampled first-order plant. A command held over a period T reaches the kick through the lag, with $a = \mathrm{e}^{-T/\tau}$, and the slow loop reads the mean of $n = 10$ fast samples, so its mode seen at the BPMs is $\sigma_m z^{-1}(b_0 + b_1 z^{-1})/(1 - a z^{-1})$ with $b_0 = 1 - c$, $b_1 = c - a$, and c the mean of a_f^j over $j = 1$ to 10, $a_f = \mathrm{e}^{-1\,\text{ms}/\tau}$. The fast loop reads the latest sample, so $b_0 = 1 - a$ and $b_1 = 0$. The controller is internal model control in each mode [8, 16], as shown in Equation (6.8).

$$\begin{aligned} \hat{d} &= \bar{y} - \hat{R}_s \bar{\hat{\theta}}_s, \qquad v_{m,k} = -u_m^{\mathsf{T}} \hat{d}_k \\ \bar{u}_{m,k} &= \lambda\, \bar{u}_{m,k-1} + K_m\,(v_{m,k} - a\, v_{m,k-1}) \\ \theta_{cmd} &= \textstyle\sum_m v'_m\, \bar{u}_m \\ K_m &= \frac{l_m(1-\lambda)}{\sigma_m(b_0 + b_1)}, \qquad \lambda = \mathrm{e}^{-2\pi f_{bw} T} \end{aligned} \tag{6.8}$$

where u_m and v'_m are the columns of U and V, $\bar{y}$ is the mean BPM reading over the period, and $\bar{\hat{\theta}}_s$ is the controller's own model of the kicks over the same period, the applied (clamped) commands passed through the lag in closed form. Because the model is driven by the clamped commands, a corrector held at its limit winds nothing up. With the model right and $l_m = 1$ the mode's closed loop is $(1-\lambda)z^{-1}\beta(z^{-1})/(1-\lambda z^{-1})$, with $\beta = (b_0 + b_1 z^{-1})/(b_0 + b_1)$, a first-order response at the declared bandwidth whatever τ_s, and with $l_m < 1$ the mode settles on the fraction l_m of the full correction. The step decay of one simulated mode gives 0.5000, 1.0002 and 3.0002 Hz for slow bandwidths of 0.5, 1 and 3 Hz, and 50.000, 100.000 and 200.000 Hz for the fast loop. A mode with $l_m = 0$ keeps its command where it is. When a BPM is removed both decompositions are rebuilt on the remaining rows, the controllers' states are projected onto the new modes, and the part of each command that no remaining BPM can see decays at the loop's own rate λ.

The noise cost follows from the same loop. White BPM noise of σ_n per sample, averaged over n samples, moves the kicks with the variance of Equation (6.9).

$$\begin{aligned} \sum_j \operatorname{var}\theta_j &= \frac{\sigma_n^2}{n}\,\frac{1-\lambda}{1+\lambda}\sum_m \frac{l_m^2}{\sigma_m^2} \\ \frac{\operatorname{rms}\theta_{cmd}}{\operatorname{rms}\theta} &= \frac{\sqrt{1+a^2-2a\lambda}}{1-a} \end{aligned} \tag{6.9}$$

The commands lead the lag, so they jitter 2.01 times as much as the kicks at a slow bandwidth of 1 Hz, 3.09 times at 3 Hz and 4.92 times at 10 Hz. The full inverse's kicks therefore exceed those of a truncation at k modes by $(\sum_{all}\sigma_m^{-2}/\sum_{m\le k}\sigma_m^{-2})^{1/2}$, a root-sum-square of $1/\sigma$, which is 16.19 at 40 modes and 12.77 at 48. A ratio of plain sums of $1/\sigma$ would give 21.1 at 40 modes.

6.5.3 The Fast Set and Offloading

The fast loop applies Equation (6.8) to $\hat{R}_f = U_f \Sigma_f V_f^{\mathsf{T}}$ every 1 ms on the latest reading, with the same kind of inverse on its own 24 modes, a default bandwidth of 100 Hz and limits of ± 5 µrad. Without offloading, each loop integrates the same error, so the two share a pole at $z = 1$ and the fast set keeps whatever steady kick it took first. Offloading, the mid-ranging of process control [6] that Volume 4, Chapter 10 applies to a power cycle, here in the internal-model form proposed for two corrector arrays [9, 17], changes only the slow loop's input, as shown in Equation (6.10).

$$\hat{d}_s = \bar{y} - \hat{R}_f\,\bar{\hat{\theta}}_f - \hat{R}_s\,\bar{\hat{\theta}}_s \tag{6.10}$$

The slow loop then corrects the orbit the beam would have without the fast set, independent of both loops' kicks, and the fast loop corrects what remains, so the two form a cascade and the fast set's steady kick drains to the slow set wherever the slow inverse reaches.

6.5.4 Gap Feedforward

Each device's end coils kick $-\hat{\kappa}_d(g_d)$ from a declared table read at the gap readback. The table is the kick law of Equation (6.5) scaled by $1+\varepsilon_d$, with $\varepsilon_d = +12$, -10 and $+8$ percent for devices 1 to 3, a declared table error. It does not read the plant's hidden kick scale.

6.5.5 A Residual Test for a Failed Monitor

At every slow update the test compares the disturbance estimate $\hat{d} = \bar{y} - \hat{R}_s\bar{\hat{\theta}}_s - \hat{R}_f\bar{\hat{\theta}}_f$ with the 28 directions the declared sources can make, the 24 girder columns of G, the dipole pattern B and the three device columns of I, all from the model lattice. A monitor that fails reads a number at one place,

which no girder, dipole or device can make. The statistics are given in Equation (6.11).

$$
\begin{aligned}
r &= (I - QQ^{\mathsf{T}})\hat{d}, \qquad e_j = \frac{r_j^2}{1 - h_{jj}} \\
\bar{e}_j &\leftarrow \bar{e}_j + \alpha_e (e_j - \bar{e}_j) \\
z_j &= \frac{\bar{e}_j}{\max(\bar{n},\ \mathrm{median}_i\, \bar{e}_i)} \\
\bar{n} &= \max(\sigma_n, 0.05\,\mu\mathrm{m})^2 / 10
\end{aligned}
\tag{6.11}
$$

where Q is an orthonormal basis of the 28 directions on the rows in use, h_{jj} the leverage of row j, and $\alpha_e = 1 - \mathrm{e}^{-10\,\mathrm{ms}/0.5\,\mathrm{s}}$. A monitor is removed from both inverses when it has the largest z_j and z_j exceeds the threshold $z = 100$ for 0.2 s. The test then starts a fresh window and cannot act for 1 s. The median guards against a model that is wrong everywhere, and the floor $\bar{n}$ keeps a start-up residual from flagging a healthy monitor when the noise rating is zero. With the test off the statistic still runs and the Ring view still shows where the residual points.

6.5.6 Hand Kicks, a Helper and Three Models

With the loops in manual the slow commands hold, and one corrector, chosen by number, takes the hand kick within $\pm 50\,\mu$rad. A helper applies one correction step from the last 10 ms mean, $\theta \leftarrow \theta - V\,\mathrm{diag}(l_m/\sigma_m)\,U^{\mathsf{T}}\bar{y}$, with 48 modes if the slow feedback is off. In automatic mode the hand station follows the command requested of the chosen corrector, and in manual the loops track the held commands, so a transfer either way is bumpless. The plant uses the true lattice, with any quadrupole error and, in case 10, one corrector's column reversed. The controllers' model is the declared lattice at the present tune, which knows neither, and the identified model is the measured response of Equation (6.6). With the model wrong, $\hat{d}$ contains $(R - \hat{R})\theta$ and each mode's loop closes through the eigenvalues η of $L(M - I)$, with $L = \mathrm{diag}(l_m)$ and $M = \Sigma^{-1}U^{\mathsf{T}}RV$, whose eigenvalues are those of $R\hat{R}^{+}$, and each eigenvalue gives a loop whose poles solve Equation (6.12).

$$
z^2 - \lambda z + \eta\,(1 - \lambda)\,\frac{b_0 z + b_1}{b_0 + b_1} = 0 \tag{6.12}
$$

For the fast loop ($b_1 = 0$) the root is $z = 1 - (1 - \lambda)e$, with $e = 1 + \eta$, stable for $0 < e < 2/(1 - \lambda)$. A loop gain e that is negative diverges at every bandwidth, and one far above 1 diverges only at a high one.

6.6 Ten Cases, Measured

The ten cases are listed in Table 6.3. Each runs on the case's seed, 20261030, with any setting its row names entered before Run, as Section 1.6.4 describes. Every number in this section comes from the simulator's own engine, run without a browser, and the scorecard judges the true orbit, sampled at the BPMs' 1 kHz. The orbit rms is over the 96 BPM places, the slow jitter is the rms of each slow command about its own 2 s mean, averaged over the 96, and the fast-set DC is the rms over the 24 fast correctors of each one's 1 s mean. The headline numbers are collected in Table 6.4.

6.6.1 One Kick and the Full Inverse

In case 1, 5 µrad at corrector 37 moves the orbit by up to 74.8 µm at BPM 85 and by 48.8 µm at a source point, and only 2 of the 96 BPMs read less than 1 % of the peak. The kick makes a closed-orbit distortion all the way round the ring, not a closed bump. At corrector 38, where β is 2.07 m instead of 18.8 m, the largest excursion is 3.0 times smaller, the ratio $\sqrt{18.8/2.07}$ that Equation (6.3) predicts. In case 2 the full inverse at 3 Hz holds the orbit at 0.09 µm rms while the slow commands jitter 0.198 µrad, and the noise model of Equation (6.9) predicts 0.199 µrad. Truncated to 48 modes for the next 20 s, the jitter falls 12 times, to 0.0164 µrad, against 12.77 predicted, and the orbit grows to

Table 6.3 The Ten Cases

#	Settings and events	Run	Teaching point
1	Manual, all feedback off, kick by hand	20 s	a column of R
2	Full inverse at 3 Hz, $\sigma_n = 1$ µm, drift 0.05	40 s	noise amplification
3	Truncated at 40 modes, drift 0.3	60 s	where to cut
4	Tikhonov, $\mu = 10$, drift 0.3	60 s	a filter, not a cut
5	Slow set only, 0.5 µm at 10 Hz	20 s	the slow bandwidth
6	Fast set added, ground step at 5 s	40 s	offloading
7	Device 2 closes to 12 mm at 3 s	20 s	feedforward
8	Fast set, offloading, BPM 37 stuck at 3 s	30 s	a residual that points
9	Challenge, full inverse, slow set only	600 s	the whole problem
10	Full inverse, corrector 38 reversed, drift 0.15	30 s	infeasible: a sign error

Drift rates q in $\mu\text{m}/\sqrt{\text{s}}$. Unless a row says otherwise the slow loop is truncated at 48 modes at 1 Hz, and the fast set, offloading, the feedforward and the BPM test are off.

Table 6.4 Measured Outcomes of the Cases, on the True Orbit Over Each Case's Run Unless Stated

#	Configuration	Measured
1	5 µrad at corrector 37 / 38	largest 74.8 / 24.8 µm, rms 40.0 / 13.4 µm
2	full / 48 modes, 20 s each	jitter 0.198 / 0.0164 µrad, orbit 0.09 / 0.21 µm
3	k = 20 / 40 / 48 / 96	2.17 / 2.14 / 0.61 / 0.09 µm, jitter 0.0009 / 0.0043 / 0.0278 / 0.0875 µrad
4	μ = 1 / 10 / 100	0.33 / 0.61 / 1.32 µm, jitter 0.0484 / 0.0256 / 0.0123 µrad
5	slow bandwidth 1 / 10 Hz	10 Hz line kept 105 / 111 %, in tolerance 92.8 / 89.1 %
6	fast set, offloading off / on at 10 s	line 22 %, fast DC at 40 s 0.18 / 0.0022 µrad
7	feedback / with feedforward	peak at source 2: 2.68 / 0.30 µm (×8.8)
8	test off / on at 5 s	41.8 µm at 30 s / within 2 µm 5.9 s after, 0.18 µm
10	full at 0.1–10 Hz / k 48, 60 / k 70	132.6–136.0 µm / 0.38, 0.38 / 16.28 µm
	μ 10 / μ 1 / measured $\tilde{R}$	0.38 / 25.45 / 0.05 µm

Case 2's orbit and case 8's and 10's values are the true orbit's rms through a 1 s filter at the end of the run. Case 9 is in Table 6.5.

0.21 µm. Measured on noise alone, the ratio of the full inverse to a truncation at 40 modes is 16.14, against 16.19 predicted.

Pause. Try Exercise 6.1 before reading on.

6.6.2 Where to Cut

Case 3 adds a minute of girder drift ($q = 0.3\,\mu\text{m}/\sqrt{\text{s}}$, about 7.3 µm per girder in ten minutes) and asks where to cut. Each point of Figure 6.3b is one run of that minute, and the measured trade falls in steps where the spectrum and the drift's content put them. Up to 44 modes the orbit stays at 2.11 to 2.17 µm rms, because the drift lies in modes 45 to 48. At 48 modes it falls to 0.61 µm with a jitter of 0.0278 µrad, and from 48 to 88 modes the orbit does not move while the jitter rises to 0.0430 µrad, because modes 49 to 88 carry almost none of the drift. Only the four weakest modes buy the rest, from 0.59 µm at 92 modes to 0.09 µm at all 96, for a jitter of 0.0875 µrad. The 48-mode cut is therefore the corner of the trade on this ring, and it sits at the cliff because the β pattern puts both the cliff and the drift there.

Tikhonov regularization in case 4 draws a smooth curve through the same plane. At $\mu = 10$ it leaves 0.61 µm for 0.0256 µrad of jitter, the residual of the 48-mode truncation with 8 % less jitter, because its filter falls where the singular values fall. A lower μ buys more of the weak modes, 0.33 µm for 0.0484 µrad at $\mu = 1$ and 0.11 µm for 0.0798 µrad at $\mu = 0.1$, and $\mu = 100$ leaves 1.32 µm for 0.0123 µrad. At every jitter the sweep covers, the Tikhonov points lie below the truncation curve, since a filter corrects a fraction of every weak mode while a cut corrects each one fully or not at all.

6.6.3 A Line Faster Than the Chamber

Case 5 vibrates the girders by 0.5 µm at 10 Hz with only the slow set on. At a slow bandwidth of 1 Hz the 10 Hz line at source point 2 is 105 % of the disturbance's own, and at the 10 Hz maximum it is 111 %, while the share of samples with all three source points in tolerance falls from 92.8 to 89.1 %.

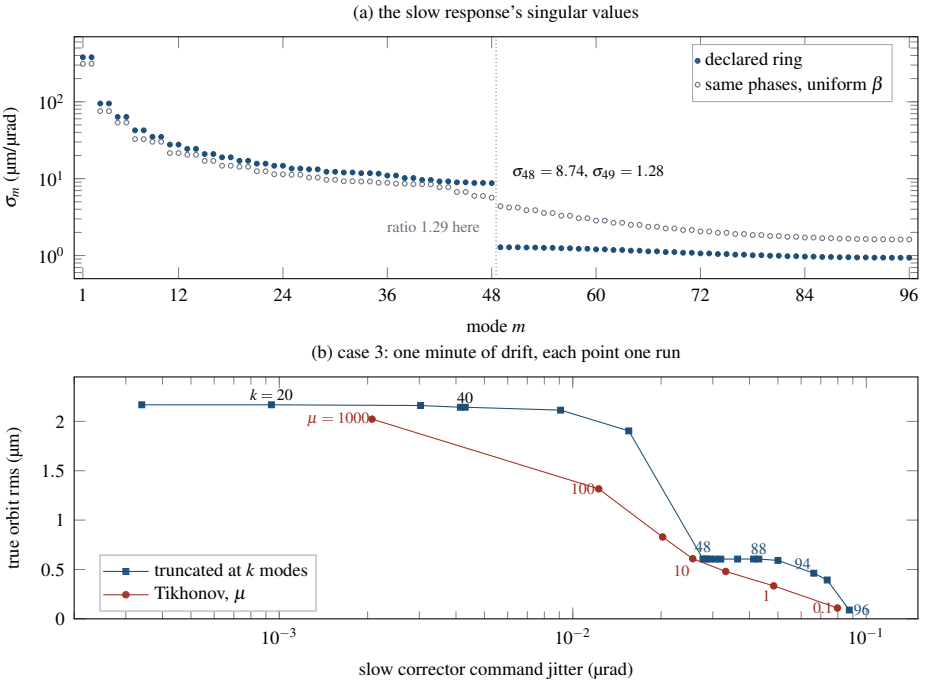

Figure 6.3 The spectrum and the trade it sets. (a) The 96 singular values of the slow response (filled dots), and of the same ring with β made uniform (open circles), where the cliff between modes 48 and 49 disappears. (b) Case 3's minute of drift, each point a run: the true orbit's rms against the slow command jitter for truncation at k modes (squares) and for Tikhonov regularization at every μ the simulator offers (dots).

The cause is the loop and not the chamber. The internal model inverts the 50 ms lag, so what limits the loop is its 10 ms period and 10 ms mean, and its sensitivity at 10 Hz derived from Equation (6.8) is 1.049 at a 1 Hz bandwidth and 1.114 at 10 Hz. An ideal first-order loop would keep 71 % of a line at its own bandwidth, and the same loop run every 5 ms would keep 92 %. The chamber is the reason the slow supplies stop at 10 Hz, and the fast set is the remedy. In case 6 the fast set at 100 Hz keeps 22 % of the line at source point 2. After the ground step at 5 s it holds the beam, and with offloading off its 1 s mean kick is still 0.18 µrad rms at 40 s. Switched on at 10 s, offloading brings that below 0.05 µrad in 1.7 s and to 0.0022 µrad by 40 s, which leaves the fast set in the middle of its ±5 µrad range. Offloading reaches only where the slow inverse reaches. Truncated at 40 or 47 modes the fast set keeps 0.25 or 0.11 µrad, from 48 modes 0.002 µrad, and every μ from 1 to 100 drains it, to 0.002 to 0.036 µrad.

6.6.4 What Feedback Cannot See

In case 7 device 2 closes from 40 to 12 mm, where it kicks 1.92 µrad. With feedback alone source point 2 reaches 2.68 µm and is still 2.63 µm off at 20 s, while every BPM reads close to zero. Ninety-six correctors can null 96 readings, but the device kicks between two BPMs, and the bump it leaves inside its own straight is invisible to all of them. The gap table cancels the kick where it occurs, and the peak falls to 0.30 µm, 8.8 times less. The rest, 0.25 µm at the end, is about the 10 % error of device 2's table applied to its 2.6 µm, which the BPMs cannot see either, so the reduction is measured with an imperfect table and not with a perfect one.

Table 6.5 The Challenge on Seven Paired Records. In Tolerance Is the Share of Samples With All Three Source Points Inside, and Jitter the Slow Commands' at the End

Variant	Pass	In tolerance (%)	Orbit rms (µm)	Jitter (µrad)
shipped: full inverse, slow set only, no test	0/7	67.70–97.48	0.72–1.75	0.126–0.153
suggested: $\mu = 10$, fast set, offloading, feedforward, test	6/7	98.89–100.00	0.51–1.13	0.011–0.015
suggested, full inverse	6/7	99.89–100.00	0.36–0.78	0.047–0.056
suggested, truncated at 48	6/7	96.75–100.00	0.52–1.21	0.012–0.017
suggested, no offloading	7/7	99.51–100.00	0.51–1.14	0.010–0.013
suggested, no feedforward	2/7	60.79–99.26	0.52–1.14	0.011–0.015
suggested, no fast set	0/7	83.63–98.55	0.99–2.00	0.011–0.015
suggested, no BPM test	0/7	86.48–99.56	1.00–1.41	0.025–0.030
shipped with the fast set and test	5/7	97.75–99.89	0.36–0.86	0.047–0.056

TRY IT IN THE APP **7 · The Insertion-Device Gap: Feedback, Then Feedforward**

Load the scenario, press Run and watch the source-point trend. Device 2 closes its gap at 3 s, and with feedback alone source point 2 reaches 2.68 µm while every monitor reads close to zero. Switch Gap feedforward On, press Reset and run again: the peak falls to 0.30 µm.

Pause. Try Exercise 6.2 before reading on.

6.6.5 A Monitor That Lies

When BPM 37 freezes 60 µm off in case 8, both loops believe it. A corrector reaches its limit 0.01 s later and the true orbit runs to 41.8 µm rms by 30 s, while the residual that no girder, dipole or device can make points at BPM 37 with 4,238 times the median share. Switched on at 5 s, the test removes BPM 37 at once, since it has stood above the threshold for 2 s, and the orbit is back within 2 µm rms 5.9 s later and at 0.18 µm by 30 s. With the test on from the start the removal comes 0.20 s after the failure. Before any failure the largest healthy share on the challenge's seven records is 2.1 times the median, against a threshold of 100. The test trusts the model lattice, and a 2 % quadrupole error raises the largest healthy share to 55, without a removal.

6.6.6 Ten Minutes, Graded

The challenge runs ten minutes with drift at $q = 0.15\,\mu\text{m}/\sqrt{\text{s}}$, a 0.3 µm girder vibration at 17 Hz, 0.05 µrad of dipole ripple, device 1 closing at 40 s, bursts at 150 and 520 s, device 2 closing at 220 s, BPM 37 turning to noise at 300 s, the ground step at 380 s and device 3 closing at 440 s. It passes if all three source points are within ± 3 µm and ± 1 µrad on at least 99 % of the 1 kHz samples, the true orbit stays within 2 µm rms over the run, no corrector command touches its limit at any sample, and BPM 37 and no other is removed within 2 s of its failure. The verdict is judged on teaching truth, since no monitor sits at a source point, and the attempt is void once any model parameter or disturbance slider moves. The variants on seven paired records, the case's seed and seeds 1 to 6, each seeing the same disturbances, are compared in Table 6.5.

The shipped settings fail every record, at 67.70 to 97.48 % in tolerance. The suggested settings pass 6 of 7, and on the case's own record they hold 99.73 % in tolerance and 0.78 µm rms and remove BPM 37 0.21 s after it fails, as shown at the moment before the removal in Figure 6.2. The seventh record, seed 5, reaches 98.89 %, and 6,580 of its 6,682 lost samples are source point 2's angle after the ground step and device 3's gap closing, the drift a regularized slow inverse leaves in its weak modes at a place no monitor sees.

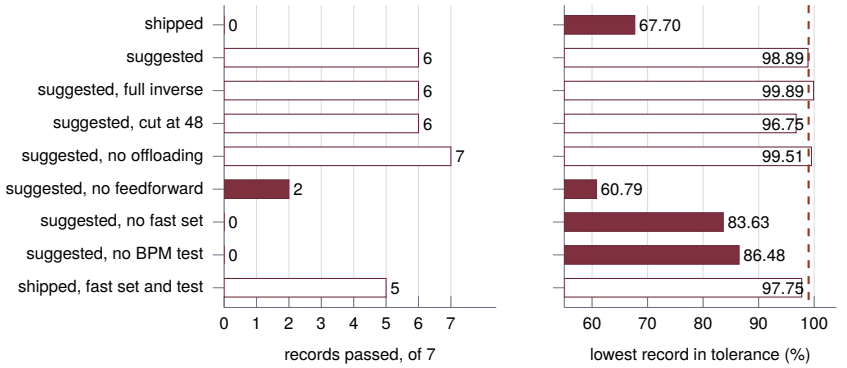

Figure 6.4 The challenge variants of Table 6.5 on its seven paired records: (left) how many records pass, (right) the lowest share of samples in tolerance among them, with the challenge's 99 % as a dashed line. Dark bars pass 2 of 7 or fewer, white bars 5 or more. Notice that the full inverse passes as often as the suggested Tikhonov setting, while removing the fast set, the feedforward or the BPM test drops the passes to 2 or none.

> ✦ **DID YOU KNOW? Quieter Magnets, Not a Pass**
>
> From the simulator: in the challenge, the suggested settings with the full inverse in place of Tikhonov regularization pass the same 6 of 7 records. What the regularization changes is the slow jitter, 0.011 to 0.015 µrad against 0.047 to 0.056 µrad (Table 6.5).

Taking the suggested settings apart one piece at a time shows what the ring needs. Without the fast set no record passes (83.63 to 98.55 %, the bursts), and without the BPM test none does either, because the noisy monitor is chased for five minutes. Without the feedforward only 2 of 7 pass (60.79 to 99.26 %), since feedback alone leaves each source point up to about 2.6 µm off once its gap has closed, as in case 7. The regularization is not what passes the challenge. The full inverse in its place passes the same 6 of 7, failing seed 5 on one 1 ms command sample at its limit at 300.149 s, before the noisy monitor is removed, and dropping offloading passes all 7. What Tikhonov regularization buys is quieter magnets, a slow jitter of 0.011 to 0.015 µrad against 0.047 to 0.056 µrad for the full inverse, and what offloading buys is a centered fast set, a mean fast kick of 0.02 to 0.10 µrad against 0.07 to 0.29 µrad without it. An independent referee written for the review of the simulator, judging the true orbit and the commands at every 1 ms, agrees with the simulator's verdict on all ten records it checks. Figure 6.4 draws the passes and the worst records of Table 6.5.

6.6.7 A Sign No Bandwidth Fixes

The plan the simulator was built from expected its infeasible case from a lattice error without a measured response. On this ring that does not happen. A 1 % rms quadrupole error beats β by 11 % rms and keeps the real part of every loop gain e of the full inverse above 0.82 and its magnitude below 1.17. The slider's 2 % beats β by 23 % rms (57 % at one BPM) and keeps them above 0.69 and below 1.41, so by Equation (6.12) every mode stays stable at every bandwidth, and on 30 s of the challenge's drift the full inverse at 10 Hz holds the orbit at 0.09 µm. A tune error from -0.15 to $+0.15$ with the β pattern right keeps $\mathrm{Re}\, e \geq 0.58$ in the slow loop. Case 10 uses a model error of another kind, corrector 38 wired backward. Its reversed column makes $M = I - 2V^{\mathsf{T}} e_{38} e_{38}^{\mathsf{T}} V$, with one eigenvalue $e = -1$, so the full inverse diverges at every slow bandwidth. It reaches a corrector limit 1.6 s in and an orbit of 132.6 to 136.0 µm at bandwidths from 0.1 to 10 Hz. Truncated at k modes the eigenvalue becomes $1 - 2\|V_k^{\mathsf{T}} e_{38}\|^2$, which holds while the reversed corrector keeps less than half its weight in the kept modes. That weight is 0.107 at 48 modes, 0.347 at 60 and 0.556 at 70,

and the runs hold at 0.38, hold at 0.38 and diverge to 16.28 µm. Under Tikhonov the smallest loop gain is 0.61 at $\mu = 10$, which holds at 0.38 µm, and -0.18 at $\mu = 1$, which diverges to 25.45 µm. The full inverse on a measured response holds at 0.05 µm. With the reversed column it tolerates element errors up to about 0.1 µm/µrad, 40 times the measurement's, where its smallest loop gain is still 0.61, and fails at 0.3 µm/µrad. The BPM test misreads the case. On from the start it removes BPMs 39 and 41, both healthy and beside the reversed corrector, after which the full inverse holds at 0.32 µm only because their rows are gone, and switched on at 5 s it removes BPMs 39, 41 and 40.

> **TRY IT IN THE APP** **10 · A Corrector Wired Backward: The Full Inverse Diverges (Infeasible)**
>
> Load the scenario and press Run: the full inverse drives a corrector to its limit 1.6 s in. Lower Slow bandwidth to 0.1 Hz, press Reset and run again, and the orbit still runs away. Then set Slow feedback to Truncated with Modes kept at 48, press Reset and run once more: the orbit holds at 0.38 µm.

Pause. Try Exercise 6.3 before reading on.

6.6.8 Checks, Agreement and the Control Rate

Code verification asks whether the simulator solves its equations. Closed orbits tracked independently with transfer matrices agree with Equation (6.3) at all 11,520 entries of R_s and R_f to 1.3×10^{-12} µm/µrad. The decomposition reconstructs R_s to 2.6×10^{-14} of its largest entry, and its singular values agree with the eigenvalues of $R_s^{\mathsf{T}} R_s$ from a QR algorithm to 9.9×10^{-15} of σ_1 and with the monitoring studio's own Jacobi solver to 2.2×10^{-15}. One simulated slow mode reproduces the kick standard deviation implied by Equation (6.9) within 1 % at 1, 3 and 10 Hz. Halving the physics step to 0.5 ms moves the rms orbit of cases 5, 6 and 7 by at most 1.4×10^{-11} percent, because every state is updated exactly over a held command and every disturbance keeps its own clock. The refinement that matters is the control rate. In a copy of the simulator whose BPMs and fast loop run at 2 kHz and slow loop at 200 Hz, with the noise per sample raised by $\sqrt{2}$ to keep its density, case 5 keeps 102.3 and 91.3 % of the line, case 6 21.1 %, and the challenge ranks the variants as before, the suggested settings passing 7 of 7 (the lowest at 99.04 %), without the feedforward 3 of 7, and without the fast set or the test none. All 40 of the simulator's checks pass, 22 written for this ring. In a desktop browser at 20 times real time the page draws 58.9 to 60.1 frames a second, and the ten minutes of the challenge take 30.0 s. Agreement with published relations covers the lattice and the inverse. The thin-lens values of Equation (6.1) and the integrated optics agree to four digits, the response is Wiedemann's closed form [1], and the filter factors are the mode-space and filter-factor forms of Tikhonov [2, 8]. No parameter is fitted to a real ring, so the model is verified and not validated. No study of learning has been made with the simulator, so no educational gain is claimed.

6.7 The Ring on the Multivariable Ladder

The ring ends a sequence in this series that runs from a two-by-two temperature laboratory (Volume 1, Chapter 4), through the identification of a many-by-many model (Volume 1, Chapter 12) and a paper machine's profile, to this hundred-by-hundred inverse. It is best taken once multivariable feedback and the singular value decomposition of Section 1.4.7 are familiar. A simulation can draw beside every monitor the truth that no real ring measures [19], here the orbit between the BPMs and at the source points. Three cases carry most of the lesson. Case 3 places the cut on a measured trade, case 7 shows a disturbance that feedback cannot see, and case 10 a model error that no tuning repairs. The residual test of case 8 echoes the gross-error tests of data reconciliation, which the monitoring workbench of Volume 1, Chapter 13 applies to stored records. A corrector wired backward, or a monitor that freezes 60 µm off in one sample, would be rare on a real ring, but each isolates one mechanism. At top speed the ten-minute challenge plays in about 30 s in a browser tab, and the export puts the true orbit and source points beside the readings under column names that begin with `truth`.

Each case is credited by a completion rule that needs the ring's response to the reader's own action (Section 1.6.4). Case 3 asks for a cut between 44 and 52 modes and a full minute run with it, and case 10 asks for the full inverse run to a corrector limit, then a truncation at 60 modes or fewer or a measured response, and an orbit held within 1 µm rms for 5 s after Reset. The challenge is graded by the function the scorecard shows.

> **OPEN CHALLENGE** **Stability at a Tenth of the Beam**
>
> The upgrade of one light source expects beam sizes of the order of 4 µm in both planes, and holding the orbit to a tenth of that pushes the state of the art in fast orbit feedback, in space and in time. Its planned system runs twelve times faster than the present one with a target bandwidth of 1 kHz, and its designers look beyond the classic PID regulator to model-based methods that use the known spectra of orbit motion, which first need a reliable identified model [3]. Progress would be such a controller shown holding a running ring.

6.8 Technology in Practice

What Measures It

A ring reads its orbit only where it has monitors. At one ring two button monitors are planned close to each new magnet so that the beam's position in it can be measured accurately, and their sensitivity and linearity are worked out from the shape of the chamber and the buttons [20]. The monitors are sampled and averaged into fast data at a fixed rate, 10 kHz at one ring, and that average delays the reading a little behind the beam's motion, an effect seen both in simulation and on the machine [21]. The chapter's BPMs are read every 1 ms, and the slow loop works on the mean of the ten readings since its last update.

> **TECHNOLOGY** **Sensor: The Button Beam Position Monitor**
>
> A button monitor sets four small electrodes into the wall of the vacuum chamber. Each passing bunch induces an image charge on them, the front-end electronics filter out a single harmonic of the bunch train's spectrum, and the beam's coordinates follow from the normalized differences of the four button voltages. The relation is linear near the center, and an electrostatic model of the exact chamber and buttons corrects its nonlinearity farther out [20]. At one storage ring the monitors read to about 100 to 250 nm in fast acquisition at 10 kHz [21]. The chapter's 96 monitors read the true orbit plus white noise of 0.5 µm per 1 ms sample, with no gain error or drift of their own.

What Moves It

Two corrector sets act on the beam. At one ring two versions of fast supplies run side by side, the newer with a bandwidth of 6.3 kHz in the horizontal plane and the original with 3.2 kHz. Their converters take a 16-bit command over a range of 1 A, a step of about 15 µA, and the feedback there tolerated supply noise up to about 0.1 mA without a change in the integrated spectrum of the orbit [21]. The chapter's correctors are limited at 5 A, which is 50 µrad for the slow set and 5 µrad for the fast one.

> **TECHNOLOGY** **Final Element: The Corrector and Its Supply**
>
> A fast corrector is a small dipole magnet fed by a bipolar power supply. One supply built for an upgraded ring is a four-quadrant converter, an H-bridge of four transistor switches under pulse-width modulation, whose output current is read by a Hall-effect sensor and held by a proportional-integral regulator. Its key requirement, a small-signal bandwidth of 10 kHz, is hard to meet because the magnet's inductance is high and the supply's input voltage low [22]. The field must then pass the vacuum chamber, whose eddy currents act at one ring as a delay of about 70 to 80 µs [21]. The chapter lumps supply, magnet and chamber into one lag, 50 ms for the slow set and 0.5 ms for the fast one.

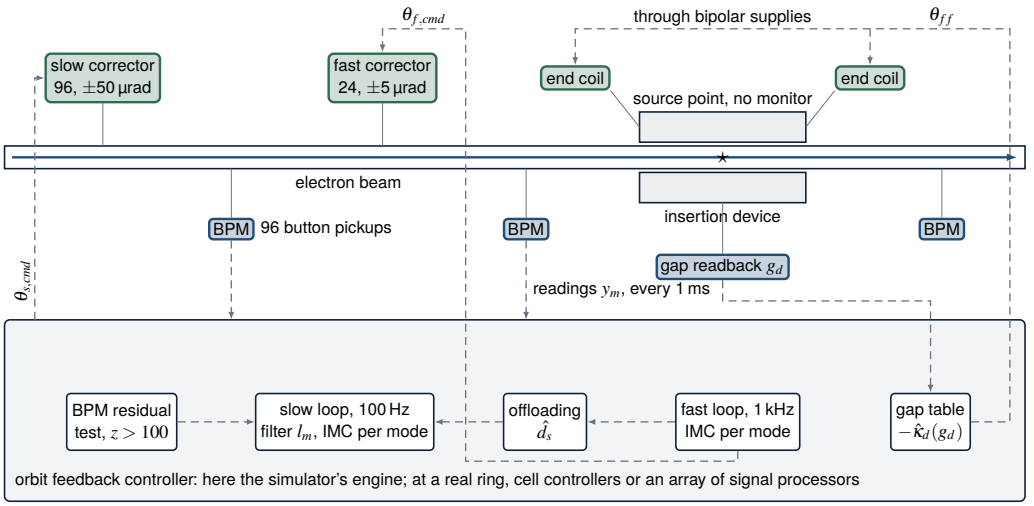

Figure 6.5 The loops of the chapter on their devices, drawn for one stretch of the ring. The 96 button BPMs (blue) feed both loops, the gap readback feeds the table, and the slow correctors, fast correctors and end coils (green) act on the beam through their supplies. Notice that the source point between the end coils carries no monitor, so only the gap table acts there, and that the slow loop reads the fast set's share through offloading.

Where the Controller Runs

At one ring the chain runs from the monitors to cell controllers, which compute the inverse response and a PID compensation, then to an interface that converts each command to an analog signal and through a current limiter to the supply, all at a correction rate of 10 kHz [21]. At another, a distributed array of digital signal processors computes the corrections at 1.6 kHz with an integral-only regulator for a unity-gain bandwidth of 60 to 80 Hz [3]. The chapter's structure is of the same kind at lower rates: the slow loop every 10 ms and the fast loop every 1 ms, each an internal-model controller per singular mode, with offloading between them, the gap table ahead of them and the residual test beside them. Figure 6.5 draws the loops on their devices.

> **IN PRACTICE** **Measure the Response, Then Choose the Cut**
>
> A ring's response matrix is measured by reading the change of beam position while the correctors are varied one by one [7]. Measuring how well the closed loop suppresses motion needs an excitation larger than the beam's own noise, yet a large one disturbs the gain it measures, so its amplitude is a compromise. At one ring with 90 modes, the first few did most of the noise suppression and about ten were enough [21]. The threshold that drops the weak modes is set mainly by the correction needed and the corrector limits, usually as small as keeps every supply out of saturation [7].

What Is Modernizing It

Three directions stand out. Faster and larger systems are one: the feedback planned for one upgraded ring runs twelve times faster than its predecessor, and its model is identified from beam measurements that include the supply, the magnet, the vacuum chamber and the monitors [3]. Supplies that switch an H-bridge at high frequency are the second, built to reach a 10 kHz bandwidth into an inductive magnet [22]. The third brings the experiment into the loop. At one nanoimaging beamline, vibrations at 27 and 120 Hz came from the beamline optics, where the electron monitors cannot see them, and a data link fed the beamline's own position monitor into the electron orbit feedback, whose bandwidth

of about 400 Hz reduced the vibration markedly [23]. As in case 7, feedback corrects only what some monitor sees.

6.9 Limitations

The ring is a declared teaching lattice of thin lenses in the horizontal plane, without coupling, dispersion or energy change, so it says nothing about vertical stability, beam size or a real machine's optics. The BPMs and the fast loop run at 1 kHz and the slow loop at 100 Hz, where the plan asked for 10 kHz and 1 kHz and real fast feedback runs at several kilohertz [8]. The plan's structure cost 103 % of each wall second at 20 times real time in a server-side runtime, and the effect of doubling the rates is the one measured in Section 6.6. Each corrector is one lag without rate limits. The measured response is the true matrix plus element noise of the right size, not a simulated 48 s measurement that would also record the drift. The feedforward table carries a declared error of 8 to 12 %, and case 7's factor of 8.8 belongs to that table. The BPM test trusts the model lattice, removes healthy monitors under a model error, as in case 10, and needed a noise floor to stay quiet at zero noise. The tune and lattice-error sliders stop at 14.35 and 2 %, which keeps the half-integer resonance out of reach. One recorded run per case is graded, seven records are checked for the challenge, and the suggested settings fail one of them. The note on what the reader changed appears 3 s after the change and describes the state at that later moment. A simulated pass is not a specification for any light source.

> **OPEN CHALLENGE** **Using the Weak Modes Without Their Noise**
>
> The small singular values of a measured response are dominated by measurement errors, yet they still carry information that orbit feedback could use. Tikhonov regularization, used at many light sources, lowers their gain instead of dropping them, and at one ring a simulation that matched its measurements pointed to a smoother spatial gain from it, with the dynamics of each mode still to be optimized [21]. Progress would be a way to correct each weak direction as far as it is known. In this chapter the four weakest modes hold 7.3 % of the drift's power.

6.10 Conclusion

A storage ring's orbit feedback inverts a matrix whose weak directions turn noise into magnet motion, and this simulator lets each choice about that inverse be made and measured on a declared ring of 96 monitors and 120 correctors. One 5 µrad kick moves the orbit by up to 74.8 µm round the whole ring, the singular values fall from 379 to 0.94 µm/µrad with a cliff at mode 48 that the β pattern makes, and the full inverse makes the slow correctors jitter 12 times as much as a 48-mode truncation. The girder drift puts 86 percent of its power in modes 45 to 48 and 7.3 percent in modes 93 to 96, so the cut at the cliff leaves 0.61 µm and only the four weakest modes buy the rest. The slow loop amplifies a 10 Hz line by its own period, a gap table corrects what no feedback sees at a point without a monitor, and the residual test finds a frozen monitor at 4,238 times the median share. In the graded ten minutes the fast set, the BPM test and the feedforward decide the verdict, while regularization and offloading buy quieter and centered magnets. No realistic lattice error breaks the full inverse here, and one corrector wired backward breaks it at every bandwidth. The ring is a thin-lens teaching model run at reduced control rates. A fuller model would add the vertical plane and coupling, rate-limited supplies, chambers with several time constants, a response measured through the drifting plant, and the rates real rings run. The simulator shows that how far to trust a model's weak directions is a design decision with a measurable price.

KEY POINTS

- A storage ring's orbit feedback inverts a matrix whose weak directions turn monitor noise into magnet motion.
- The full inverse makes the slow correctors jitter 12 times as much as a 48-mode truncation.
- The girder drift lies mostly in modes 45 to 48, so the cut at the cliff leaves 0.61 µm and only the four weakest modes buy the rest.
- A gap table corrects what no feedback sees at a point without a monitor, and one corrector wired backward breaks the full inverse at every bandwidth.

Where else this applies. Volume 4, Chapter 10 keeps a fast handle centered by a slow one, as offloading does here for the fast correctors.

References

[1] H. Wiedemann, *Particle Accelerator Physics*, ser. Graduate Texts in Physics. Springer International Publishing, 2015.

[2] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. Society for Industrial and Applied Mathematics, 1998.

[3] P. Kallakuri, D. Paskvan, J. Carwardine, N. Sereno, and N. Arnold, "Modeling the fast orbit feedback control system for APS upgrade," in *Proceedings of IBIC2017*. Grand Rapids, MI, USA: JACoW, 2018, pp. 196–198.

[4] S. Skogestad, M. Morari, and J. C. Doyle, "Robust control of ill-conditioned plants: high-purity distillation," *IEEE Transactions on Automatic Control*, vol. 33, no. 12, pp. 1092–1105, 1988.

[5] G. E. Stewart, D. M. Gorinevsky, and G. A. Dumont, "Feedback controller design for a spatially distributed system: the paper machine problem," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 5, pp. 612–628, 2003.

[6] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[7] Y. Chung, G. Decker, and K. Evans, Jr., "Closed orbit correction using singular value decomposition of the response matrix," in *Proceedings of International Conference on Particle Accelerators*. IEEE, 1993, pp. 2263–2265.

[8] S. Gayadeen, S. R. Duncan, and M. T. Heron, "Fast orbit feedback control in mode space," in *Proceedings of ICALEPCS2013*, San Francisco, CA, USA, 2013, pp. 1082–1085, paper THMIB07.

[9] S. Gayadeen, S. R. Duncan, and W. P. Heath, "Design of multi-array controllers for electron beam stabilisation on synchrotrons," in *2013 American Control Conference*. IEEE, 2013, pp. 1201–1206.

[10] E. D. Courant, M. S. Livingston, and H. S. Snyder, "The strong-focusing synchroton—a new high energy accelerator," *Physical Review*, vol. 88, no. 5, pp. 1190–1196, 1952.

[11] E. D. Courant and H. S. Snyder, "Theory of the alternating-gradient synchrotron," *Annals of Physics*, vol. 3, no. 1, pp. 1–48, 1958.

[12] G. Golub and W. Kahan, "Calculating the singular values and pseudo-inverse of a matrix," *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, vol. 2, no. 2, pp. 205–224, 1965.

[13] J. Demmel and K. Veselić, "Jacobi's method is more accurate than QR," *SIAM Journal on Matrix Analysis and Applications*, vol. 13, no. 4, pp. 1204–1245, 1992.

[14] J. Safranek, "Experimental determination of storage ring optics using orbit response measurements," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 388, no. 1–2, pp. 27–36, 1997.

[15] R. O. Hettel, "Beam stability at light sources (invited)," *Review of Scientific Instruments*, vol. 73, no. 3, pp. 1396–1401, 2002.

[16] C. E. Garcia and M. Morari, "Internal model control. a unifying review and some new results," *Industrial & Engineering Chemistry Process Design and Development*, vol. 21, no. 2, pp. 308–323, 1982.

[17] S. Gayadeen and W. Heath, "An internal model control approach to mid-ranging control," *IFAC Proceedings Volumes*, vol. 42, no. 11, pp. 542–547, 2009.

[18] S. Gayadeen and S. R. Duncan, "Discrete-time anti-windup compensation for synchrotron electron beam controllers with rate constrained actuators," *Automatica*, vol. 67, pp. 224–232, 2016.

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

[20] V. Smaluk, B. Kosciuk, and B. Bacha, "A button-type beam position monitor for the complex bend test in NSLS-II storage ring," Brookhaven National Laboratory, Upton, NY, USA, Tech. Rep. NSLSII-ASD-TN-446 (BNL-230229-2026-TECH), 2026.

[21] S. Kongtawong, Y. Tian, L. Yu, X. Yang, G. Wang, K. Ha, and T. Shaftan, "Numerical simulation of NSLS-II fast orbit feedback system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 997, p. 165175, 2021.

[22] J. Wang and G. Sprau, "A high bandwidth bipolar power supply for the fast correctors in the APS upgrade," in *Proceedings of NAPAC2016*. Chicago, IL, USA: JACoW, 2017, pp. 96–98.

[23] S. Kongtawong, Y. Hidaka, H. Yan, G. Wang, Y. Tian, K. Ha, J. Mead, D. Padrazo, Y. Chu, and T. Shaftan, "Stabilization of x-ray beam for nanoimaging using source electron fast orbit feedback," in *X-Ray Nanoimaging: Instruments and Methods VII*. San Diego, CA, USA: SPIE, 2025.

6.11 Exercises

6.1 Concept check · 2 min

In scenario 2 the full inverse makes the slow correctors jitter 0.198 µrad about their own mean. Truncated to 48 modes on the same BPM noise and the same 3 Hz bandwidth, they jitter 0.0164 µrad, twelve times less. What makes the difference?

(a) The full inverse divides the noise in the 48 weak modes by singular values near 1 µm/µrad

(b) The full inverse corrects the weak modes, which carry most of the girder drift

(c) The 50 ms chamber amplifies the commands of the full inverse more than those of the truncated one

(d) The full inverse runs the slow loop every millisecond instead of every 10 ms

6.2 Concept check · 1 min

In scenario 7 device 2 closes its gap with feedback alone. After 20 s every BPM reads close to zero, but source point 2 is still 2.63 µm off. Why does the slow loop leave it there?

(a) The 1 Hz slow loop has not yet caught up with a gap that closes over 5.6 s

(b) The device kicks between two BPMs, and the bump inside its straight is invisible to all of them

(c) The 48-mode truncation drops the modes in which the device's kick lies

(d) The gap table is off by 10 % for device 2

6.3 Concept check · 2 min

In scenario 10 corrector 38 is wired backward and the full inverse drives the orbit to more than 130 µm rms. Why does lowering the slow bandwidth to 0.1 Hz not stop it?

(a) The correctors reach their 50 µrad limit, and the integrators wind up behind it

(b) A 0.1 Hz loop is too slow for the ground drift of the case

(c) One eigenvalue of the loop gain is -1, and a negative loop gain diverges at every bandwidth

(d) The BPM test removes two healthy monitors and leaves the inverse without rank

6.4 Theory and calculation · 8 min

The teaching ring has a tune $\nu = 14.2$, 48 thin-lens FODO periods and drifts of $L = 5\,\mathrm{m}$. Use Equations (6.1) and (6.3).

(a) Find the phase advance μ of one period, the focal length f and the amplitude function $\beta_{\pm}$ at a QF and at a QD. Why does each 20 m cell carry two periods? *Checkpoint:* 106.5°, 3.12 m, 18.79 m and 2.07 m; one period per cell would need 213°.

(b) A BPM 0.4 m after a QF sees $\beta = 16.46\,\mathrm{m}$. Taking $|\cos\Delta_{ij}| = 1$, find the largest orbit move at such a BPM for a 5 µrad kick at a corrector at a QF and at one at a QD. *Checkpoint:* 14.96 and 4.97 µm/µrad, so 74.8 and 24.8 µm.

(c) Find the ratio of the two moves and say what sets it. *Checkpoint:* 3.0, $\sqrt{18.79/2.07}$.

6.5 Theory and calculation · 7 min

The slow loop runs every $T = 10\,\mathrm{ms}$ behind a chamber lag of $\tau_s = 50\,\mathrm{ms}$. Its singular values include $\sigma_1 = 379.17$, $\sigma_{48} = 8.740$, $\sigma_{49} = 1.282$ and $\sigma_{96} = 0.9357\,\mu\mathrm{m}/\mu\mathrm{rad}$. Use Equations (6.7) and (6.9).

(a) Find the Tikhonov filter factor l_m at $\mu = 10$ for these four modes. *Checkpoint:* 1.00, 0.88, 0.14 and 0.08.

(b) At what singular value is a mode corrected halfway, and where does it fall in the spectrum? *Checkpoint:* $\sigma = \sqrt{10} = 3.16\,\mu\mathrm{m}/\mu\mathrm{rad}$, inside the cliff between modes 48 and 49.

(c) Find the ratio of command jitter to kick jitter at slow bandwidths of 1 and 10 Hz. *Checkpoint:* 2.01 and 4.92.

6.6 Theory and calculation · 6 min

In scenario 10 corrector 38 is reversed. With the inverse truncated at k modes, the reversed mode's loop gain is $e = 1 - 2\|V_k^{\mathsf{T}} e_{38}\|^2$. The weight $\|V_k^{\mathsf{T}} e_{38}\|^2$ is 0.107 at 48 modes, 0.347 at 60 and 0.556 at 70.

(a) Find e at each truncation and predict which runs hold. *Checkpoint:* 0.786, 0.306 and -0.112: 48 and 60 hold, 70 diverges.

(b) For the fast loop ($b_1 = 0$, $T = 1\,\mathrm{ms}$) at 100 Hz, find λ, the range of e for which the loop is stable, and the pole for $e = -1$. *Checkpoint:* $\lambda = 0.533$, $0 < e < 4.29$, $z = 1.47$.

6.7 App challenge · 8 min

Load scenario 2, The full inverse: the magnets chase the noise, at the default seed and a speed of 1. The scenario sets Slow feedback to Full inverse at a slow bandwidth of 3 Hz and a BPM noise of 1 µm.

(a) Run to 20 s and pause. Read the Slow corrector jitter tile, the noise prediction it shows beside the measurement, and the RMS orbit through the 1 s filter. *Checkpoint:* 0.198 µrad against 0.199 µrad predicted, orbit 0.09 µm.

(b) Set Slow feedback to Truncated (Modes kept stays at 48) and run to the end of the case at 40 s. Read the same tiles. *Checkpoint:* 0.016 µrad against 0.016 µrad, orbit 0.21 µm.

(c) Compare the jitter ratio with the 12.77 that Equation (6.9) predicts, and say what the truncation costs. *Checkpoint:* About 12; the orbit doubles, from 0.09 to 0.21 µm.

Accept jitters within 0.002 µrad and orbits within 0.03 µm. A switch up to half a second from 20 s stays inside these tolerances.

6.8 App challenge · 8 min

Load scenario 3, Truncation at 40 modes: residual against noise, at the default seed and a speed of 5. The case runs one minute of girder drift and pauses at 60 s. For each run read the RMS orbit over the run and the Slow corrector jitter at the end.

(a) Run the case with its 40 modes. Then set Modes kept to 48, press Reset and run it again. *Checkpoint:* 2.14 µm and 0.004 µrad; 0.61 µm and 0.028 µrad.

(b) Repeat with 88 and with 96 modes kept. *Checkpoint:* 0.61 µm and 0.043 µrad; 0.09 µm and 0.088 µrad.

(c) Explain why 40 extra modes from 48 to 88 buy nothing while the last eight buy most of the remaining residual. *Checkpoint:* The drift's power by mode: modes 45 to 48 and 93 to 96.

Accept orbits within 0.02 µm and jitters within 0.002 µrad.

6.9 App challenge · 7 min

Load scenario 7, The insertion-device gap: feedback, then feedforward, at the default seed and a speed of 2. Device 2 closes from 40 to 12 mm starting at 3 s, and the case pauses at 20 s.

(a) Run the case with feedback alone. Read the Worst source point tile, and the position of source point 2 at 20 s from the column `truth_src2_um` of the exported file. *Checkpoint:* 2.69 µm; 2.63 µm.

(b) Set Gap feedforward to On, press Reset and run again. Read the same two values. *Checkpoint:* 0.31 µm; 0.25 µm.

(c) Account for the 0.25 µm that the feedforward leaves. *Checkpoint:* The table's declared −10 % error for device 2.

Accept positions within 0.03 µm.

6.10 Application area · 4 min

Scenario 10 can be held by inverting a measured response instead of the model. A real measurement kicks each of the 96 slow and 24 fast correctors by ±10 µrad and averages the BPMs over 200 samples each way at 1 kHz.

(a) How long does the measurement take, and what element error does Equation (6.6) give at a BPM noise of 0.5 µm? *Checkpoint:* 48 s; 0.0025 µm/µrad.

(b) The full inverse on the reversed ring tolerates element errors up to about 0.1 µm/µrad. Name two things a measurement on a working ring would add that the simulator leaves out. *Checkpoint:* Drift during the measurement, and beam time.

6.11 Application area · 3 min

A beamline scientist reads that the suggested settings keep all three source points within ±3 µm and ±1 µrad on 99.73 % of the samples of the graded challenge. She asks whether the ring would meet her experiment's needs. Give three reasons from the chapter why this pass is not a specification.

6.12 Review the agent · 5 min

An AI assistant reviews the graded challenge and writes: "The shipped full inverse fails because it amplifies BPM noise through the weak modes. The fix is Tikhonov regularization at $\mu = 10$. It is the setting that makes the suggested configuration pass." Find the flaw in the reasoning and say how you would check it on the app.

CHAPTER 7

Switchable Control Structures and the Snowball Effect in a Reactor–Column Recycle Plant

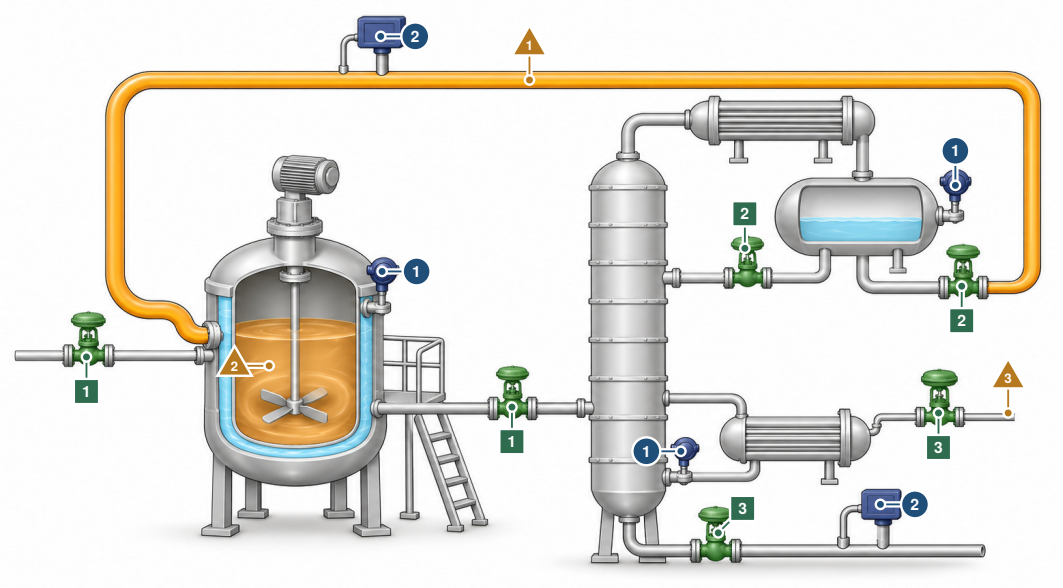

The controller moves 1 fresh feed valve and reactor outlet valve 2 reflux valve and recycle valve under the drum 3 product valve and reboiler steam valve

The controller reads 1 levels of the reactor, the reflux drum and the column sump 2 purity analyzers on the recycle and on the product

What disturbs it 1 a production step, which the recycle multiplies 2 loss of catalyst activity 3 a limit on the reboiler steam

The fresh feed enters at the far left while the much wider recycle returns over the top, and at a fixed reactor holdup a production step swells that recycle and with it the vapor the column must boil. Illustration, not to scale.

IN PLAIN TERMS A Loop That Turns a Nudge Into a Shove

The main idea. A reactor turns a liquid into a product but leaves part of it unchanged. A tall column boils that leftover off and pipes it back to the reactor, so the same liquid goes round and round.

The challenge. Ask for a little more product and the returning stream grows far more: ten percent more feed swells it by more than half. The column must boil all of it, and it can only boil so fast.

Why it matters. Which flow the operators set, and where, decides whether the reactor, the returning stream or the column takes the strain.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on the windup in Section 7.1, the history at the end of Section 7.2, the try-its and the fact on a catalyst loss in Section 7.6, and the key points at the end of Section 7.9
Study	40 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 7.6), Technology in Practice (Section 7.7), Limitations (Section 7.8) and the open challenge

The control question A reactor that leaves part of its feed unconverted sends that reactant round a column and back, and a small production change can become a large load on the column. Which flow of the recycle loop should be held fixed, where should production be set, and what does each choice do to the recycle, the reactor and the column's vapor limit?

Learning objectives

- Derive the steady recycle of the reactor-column plant from two balances, and its elasticity and pole.
- Distinguish the five quantities called a snowball gain by their definitions and units.
- Explain why every structure that holds the holdup and both purities lands on one curve, so the structure chooses only the production handle.
- Compute the reactor holdup that moves production on CS2 and the recycle purity a fixed recycle would need on CS3.
- Show on the app why a purity loop cannot recover a column at its vapor limit, and pass the graded day by choosing structure and holdup.

Prerequisites Sections 1.4.1, 1.4.3 and 1.4.4 (the ideal-separator example, the throughput manipulator and the snowball), 1.6.2 (velocity-form PI) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/recycle-snowball/

See also Volume 13, Chapter 2 for a recycle with a converter inside it and inventory pairing, and Volume 13, Chapter 7 for a recycle that changes a loop's gain rather than its speed. Volume 13, Chapter 8 for a column whose spare handle moves only the energy, and Volume 13, Chapter 11 for setpoints moved by a steady-state optimizer.

7.1 Introduction

A reaction that stops short of completion leaves reactant that has to be separated from the product and sent back, and the flowsheet that results has a loop in it. In the generic plant simulated here, 200 kmol/h of fresh reactant A meets 500 kmol/h of recycled distillate in a stirred reactor holding 400 kmol, which converts 196 of the 675 kmol/h of A it receives, 29.0 % per pass. A column of twenty trays boils 1115 kmol/h, 67.6 % of its flooding rate, to return the unconverted A at 95 mol % and to send the product B out at 98 mol %. With the reactor holdup and both purities held, 10 % more fresh feed takes the recycle from 500 to 812 kmol/h in 12 h, on its way to a steady 817 kmol/h, and leaves the column 3.1 % of vapor margin.

Three beliefs about such a plant fail on the running simulator. The first is that a recycle stream grows in proportion to production. The steady recycle here grows 63.4 % for a 10 % step, because the reactor can convert more A in the same holdup only by running richer in A, and the column returns that richer effluent's A. The second is that fixing a flow inside the loop removes the problem. With the reactor effluent fixed and production moved by the holdup, the recycle falls 4.2 % while production rises 10.4 % by 6 h, but the holdup has to rise 14 % for each 10 % of production, and the same structure moved by its effluent lands back on the curve of the first. The third is that the snowball has one size. Five different quantities carry the name snowball gain on this one design point, in two units, and they read 4.61, 11.5, 11.55, 63.4 and 9.4 to 11.1, as Table 7.3 sets out.

The shape of the problem is general. A recycle is a feedback path inside the plant, so it moves the plant's poles and with them its speed and possibly its stability [1], and some recycle is a common feature of most plants [2]. Wherever a unit inside a recycle loop is pushed toward its capacity, the loop amplifies the push. An esterification plant whose extraction step limited it showed the same windup, with flooding of its extractor and columns when the throughput was set too high [3]. The ammonia synthesis loop of Volume 13, Chapter 2 circulates several times its makeup flow for the same reason, a low conversion per pass, and the gas recycle of Chapter 8 carries an inert that only a purge removes. A second decision is attached to every such loop. The point where production is set, the throughput manipulator, fixes how every inventory upstream and downstream of it has to be held [3–6], as Section 1.4.3 sets out, and whether that point lies inside or outside the recycle decides which flows move when production does.

> **DID YOU KNOW? Half a Day Before the Windup Shows**
>
> On that esterification plant the recycle windup, which its authors liken to the snowball effect, took 12 to 24 hours to engage fully. A rate that was stable one day could wind up the next, so operators set production below the optimum [3].

The reactor, column and recycle plant of the plantwide control literature is usually met as steady-state curves and as simulation studies reported in papers [7, 8]. A reader of those seldom switches the plant's control structure during a run, measures the plant against the curve, or watches the column's vapor rate reach its limit. This chapter places that plant in a browser with the control structure as a switch, and measures on it what each structure does with one production change, a catalyst loss and a vapor limit. The balances, the analytical steady state, which is the one Larsson, Govatsmark, Skogestad and Yu derived [8], and the proportional-integral (PI) loops are established methods. The measurements on one generic design show three things. At a fixed holdup every structure that holds both purities settles on one curve, to 0.005 %, so a structure decides only which handle moves production. The holdup handle removes the snowball from the recycle and moves it into the reactor. Switching the recycle purity loop off raises the recycle's 12 h rise from 62.3 to 93.5 %. Behind the measurements are a sixty-nine-state model whose mole and A balances close to 1×10^{-11} of the throughput, a column module with a closed-form bubble point that other columns of the series reuse, and an analytical overlay computed apart from the dynamic model. Each of the nine cases carries a rule that credits its completion (Section 1.6.4).

Section 1.4.4 derived the snowball on a plant with an ideal separator. This chapter adds the column, its vapor limit and the line between them. Section 7.2 traces plantwide control of recycle plants from 1964. Section 7.3 builds the reactor, the line and the column, Section 7.4 derives the steady state and compares the three handles, and Section 7.5 describes the loops as coded. Section 7.6 measures the nine cases, Section 7.7 describes the real meters, valves and controllers, Section 7.8 states what the model leaves out, and Section 7.9 sums up.

7.2 Recycle, Inventory and the Throughput Handle Since 1964

The control of a whole plant, rather than of its units one by one, was first discussed at length by Page Buckley in a chapter on overall process control in his 1964 book [4]. A later review finds there the material balance control that runs in the direction of flow or against it, production rate control, surge tanks as low-pass filters, and the observation that inventory and flow cannot be set independently at one point of a plant, because the material balance ties them [2]. The same review credits Buckley with the point that where the production rate is set determines the structure of the inventory loops, and it judges that industrial practice had not moved far from his chapter in the thirty-five years that followed. In the same year Gilliland, Gould and Boyle used a plant with a reaction, a separation and a recycle to show how the positive feedback of the recycle changes the dynamics and the steady state [9], as Larsson and coauthors summarize that study [8].

The name came three decades later. In a 1993 series on recycle systems, Luyben took up the high sensitivity of the recycle flow to the fresh feed and named it the snowball effect [8, 10], then gave the reactor and separator case a paper of its own [7]. His remedy was to let the reactor holdup vary and, as a generic rule, to put one flow somewhere in the recycle loop on flow control [8]. The rule became part of the nine-step plantwide procedure of Luyben, Tyreus and Luyben, which was applied to a vinyl acetate process, the Eastman plantwide control process and a hydrodealkylation plant [2, 11]. Price and Georgakis built a tiered procedure on Buckley's ideas and, after 318 simulations of a reactor, separator and recycle, advised that an internal flow be preferred as the throughput manipulator and that the inventory loops be consistent with it [2, 5].

The rule was soon tested on the same plant. Wu and Yu proposed balanced structures that hold the reactor composition and found that Luyben's structure moves the snowball from the recycle into the reactor holdup, while a fixed flow at a constant holdup could take only small throughput

changes [8, 12]. Larsson and Skogestad's review then argued that the snowball occurs because the reactor is too small for the feed, which makes it a steady-state design problem [2]. Skogestad framed the choice of what to control as self-optimizing control, a search for controlled variables that keep the economic loss acceptable at constant setpoints [13]. Applied to this plant, the argument places a liquid reactor at its largest holdup, an active constraint, where the conventional structure gives less recycle than Luyben's structure at every feed rate [8]. Skogestad's procedure for complete plants then sets the production rate at the bottleneck, which for this plant is the reactor holdup at steady state and the column's vapor capacity dynamically [14]. Aske and Skogestad stated the requirement on the inventory loops as local consistency [6].

Industry had moved the handle before these procedures were published, as the history segment at the end of this section tells, and one justification offered for Luyben's rule is that a fixed flow in the loop avoids the recycle's slow dynamics [1, 3]. The column that sits in the loop is modeled here the way the distillation tutorial literature models it, with constant relative volatility, constant molar flows and linearized liquid dynamics on the trays [15], while a real column's flooding limit comes from tray evaluation procedures built on Fair's flooding and entrainment data of 1961, which were refitted as algebraic correlations as recently as 2025 [16].

> ⧖ HISTORY **Late 1970s: Production Set Away From the Feed**
>
> In the late 1970s and early 1980s one chemical company moved the throughput manipulator of numerous plants away from the process feeds. On an esterification plant whose units had been designed with the throughput set at the feed, the handle went first to the distillate of the first column and later to the feed of the extraction step, once that step had become the bottleneck. Early adoption was hard for three reasons: a mindset that the process feeds must stay constant, the operators' grasp of an inventory-to-feed strategy, and pneumatic hardware that made the control decision difficult to reverse. Setting the handle somewhere other than the feed is now common there, driven by variability propagation and ease of operation [3].

7.3 The Reactor, the Column and the Line as Modeled

The plant has three units, as shown in Figure 7.1. Pure A enters a stirred tank, where it reacts to B. The tank's effluent feeds tray 9 of a twenty-tray column, whose condensed overhead collects in a reflux drum. The drum's liquid is split between the reflux and the distillate, and the distillate returns to the reactor through a pipe as the recycle. A partial reboiler with its sump boils the column's bottom liquid and releases the bottoms as the product. Every size, rate constant and limit is a generic teaching value chosen for a consistent design point, not a measurement on any plant, and exports carry the label synthetic educational model.

Table 7.1 lists the variables by role. The clock is in minutes and the flows in kmol/h, so every balance below carries the factor $1/60$ in the code, which the equations omit. Sixty-nine states are integrated, four in the reactor (holdup, composition, temperature and coolant temperature), ten in the recycle line, two on each of the column's twenty-two stages (drum, trays and sump), six flow states and five ledgers that count the moles and the A entering, leaving and reacting.

7.3.1 Reactor Holdup and Composition

The irreversible reaction $A \rightarrow B$ conserves moles, so the reactor holdup changes only with the flows, and its A changes with the flows and the reaction, as shown in Equation (7.1).

$$\begin{aligned} \frac{\mathrm{d}M_R}{\mathrm{d}t} &= F_0 + D - F, \qquad M_R = \rho\, V_{vessel}\, \frac{L_R}{100} \\ \frac{\mathrm{d}(M_R z)}{\mathrm{d}t} &= F_0 z_0 + D x_{rec} - F z - k M_R z \end{aligned} \tag{7.1}$$

where z is the mole fraction of A in the reactor, z_0 in the fresh feed and x_{rec} in the recycle as it reaches the reactor, $\rho = 4\,\mathrm{kmol/m^3}$ is the molar density and V_{vessel} the vessel volume. The design

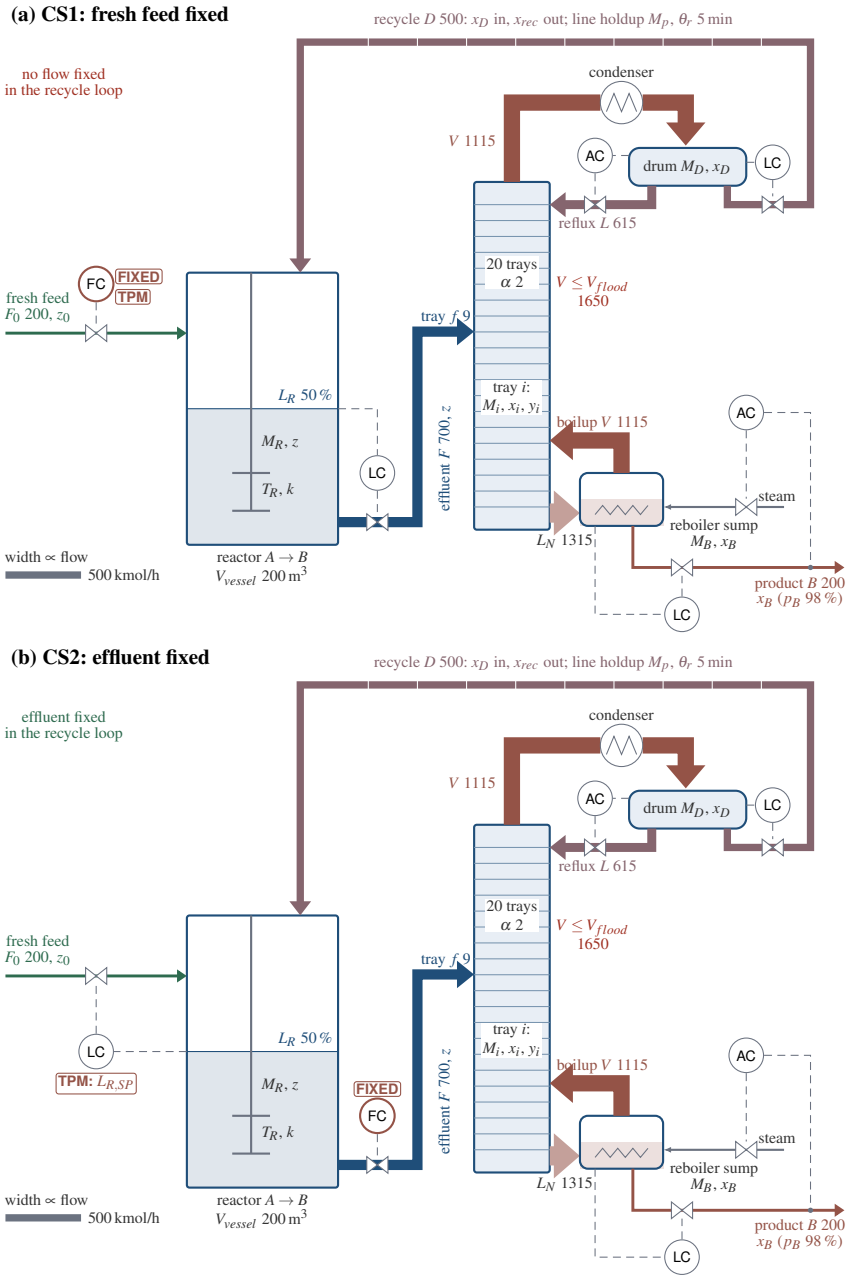

Figure 7.1 The plant at its design point, drawn like the application's Plant view. Every stream's line width is proportional to its flow on one scale (500 kmol/h in the key), flows are in kmol/h, and each valve carries the controller that moves it. (a) In CS1, the conventional structure, the fresh feed F_0 is on flow control and is the throughput manipulator (TPM), and the reactor level controller moves the effluent F. (b) In CS2 the effluent is fixed, the reactor level controller moves the fresh feed, and production is set by the reactor level setpoint. In both, the drum level moves the recycle D, the recycle purity x_D moves the reflux L, the product purity p_B moves the reboiler's steam and so the boilup V, and the sump level moves the product B. CS3 is described in Table 7.4. Not to scale.

Table 7.1 Plant Variables, Their Roles and Design Values

Symbol	Role	Design	Range
L_R	CV, reactor level (%)	50	SP 10–90
x_D	CV, recycle purity (mol % A)	95.00	SP 85–99.5
p_B	CV, product purity (mol % B)	98.00	SP 90–99.9
L_D, L_B	CV, drum and sump levels (%)	50, 50	SP 20–80
F_0	MV, fresh feed (kmol/h)	200	0–400
F	MV, reactor effluent (kmol/h)	700.19	0–2400
D	MV, recycle (distillate, kmol/h)	500.19	0–2000
L	MV, reflux (kmol/h)	614.76	0–3000
V	MV, boilup (kmol/h)	1114.95	0–V_{cap}
B	MV, product (bottoms, kmol/h)	200	0–500
a	DV, catalyst activity (%)	100	50–120
z_0	DV, A in the fresh feed (mol %)	100	80–100
V_{steam}	DV, boilup the steam allows (kmol/h)	2000	800–2000
M_R, z	states, reactor holdup (kmol), A (mol %)	400, 68.44	—
T_R, T_c	states, reactor, coolant (°C)	70, 59.6[a]	—
x_j	states, A in the line's cells (mol %)	95.00	$j = 1 \ldots 10$
M_i, x_i	states, stage holdup (kmol), A	4, profile	$i = 1 \ldots 20$
M_D, M_B	states, drum, sump (kmol)	50, 60	cap 100, 120

[a]The coolant matters only with the jacket option.

holds 100 m^3 of liquid in a 200 m^3 vessel, so one percent of level is 8 kmol and the holdup can rise as a handle. The residence time at design is $M_R/F = 34.3$ min and the reaction's own time constant $1/k$ is 83.8 min.

7.3.2 Kinetics and the Optional Jacket

The rate constant follows the Arrhenius law about the design temperature, scaled by the catalyst activity, as shown in Equation (7.2).

$$
\begin{aligned}
k &= k_0\, a \exp\Big[-\frac{E_a}{R}\Big(\frac{1}{T_R}-\frac{1}{T_0}\Big)\Big] \\
\tau_T \frac{\mathrm{d}T_R}{\mathrm{d}t} &= T_{R,SP} - T_R \qquad \text{(isothermal default)} \\
M_R c_p \frac{\mathrm{d}T_R}{\mathrm{d}t} &= (-\Delta H_r)\, k M_R z + c_p (F_0 + D)(T_f - T_R) \\
&\quad - UA\,(T_R - T_c), \qquad \tau_J \frac{\mathrm{d}T_c}{\mathrm{d}t} = T_{c,SP} - T_c
\end{aligned}
\tag{7.2}
$$

where R is the gas constant, temperatures in the exponent are absolute, a is the catalyst activity, $k_0 = 0.716\,\mathrm{h}^{-1}$ at $T_0 = 70\,^\circ$C and $E_a = 60$ kJ/mol, so near the design one kelvin raises k by 6.3 %. In the energy balance the flows are taken per second and c_p is per mole. By default an ideal temperature loop holds the reactor with $\tau_T = 5$ min. The jacket option replaces it with the energy balance, in which the reaction releases 2.18 MW at design, the feeds enter at $T_f = 50\,^\circ$C and a temperature controller sets $T_{c,SP}$, which the coolant temperature T_c follows with $\tau_J = 3$ min. The coolant sits at 59.6 °C at the design steady state.

7.3.3 A Line That Carries Composition Late and Flow at Once

The recycle line is ten well-mixed cells of fixed holdup, as shown in Equation (7.3).

$$
\begin{aligned}
\frac{M_p}{N_p}\frac{\mathrm{d}x_j}{\mathrm{d}t} &= D\,(x_{j-1} - x_j), \qquad j = 1 \ldots N_p \\
x_0 &= x_D, \qquad x_{rec} = x_{N_p}
\end{aligned}
\tag{7.3}
$$

where $N_p = 10$ and $M_p = \theta_r D_{ref}/60$, so the line holds 41.7 kmol, the content of a $\theta_r = 5$ min transit at $D_{ref} = 500$ kmol/h. Because the holdup is fixed, the line runs full of liquid and delivers the recycle

at the rate it leaves the drum at every instant. Only the composition arrives late. The minimum θ_r is 1 min, which keeps the explicit 2 s step on these cells stable up to the 2000 kmol/h valve limit (a step size times eigenvalue of 1.33 against the 2.79 that the fourth-order Runge-Kutta method tolerates on the real axis).

7.3.4 Twenty Trays, a Drum and a Reboiler Sump

The column follows the binary model of the distillation tutorial literature [15], with constant molar overflow, constant relative volatility and a linearized weir on each tray, as shown in Equation (7.4).

$$\begin{aligned} \frac{\mathrm{d}M_i}{\mathrm{d}t} &= L_{i-1} + F\,\delta_{if} - L_i, \qquad L_i = \bar{L}_i + \frac{M_i - \bar{M}}{\tau_L} \\ \frac{\mathrm{d}(M_i x_i)}{\mathrm{d}t} &= L_{i-1}x_{i-1} + V y_{i+1} + F z\,\delta_{if} - L_i x_i - V y_i \\ y_i &= \frac{\alpha x_i}{1 + (\alpha - 1)x_i} \\ T_i &= \Lambda \Big/ \ln\Big(\frac{P_0}{P}\sum_j x_{ij}\,\mathrm{e}^{\Lambda/T_{b,j}}\Big) \end{aligned} \tag{7.4}$$

where trays are numbered $i = 1 \ldots 20$ from the top, the feed enters as saturated liquid on $f = 9$, δ_{if} is one on the feed tray, the vapor V is the same on every tray, $L_0 = L$ and $x_0 = x_D$, and $y_{21} = y_B$ is the vapor leaving the sump, x_{ij} is the fraction of component j on tray i, $T_{b,j}$ its normal boiling point, and $P = P_0 = 1.01325$ bar. The weir's reference flows $\bar{L}_i$ are those of a 615 kmol/h reflux and a 700 kmol/h feed, $\bar{M} = 4$ kmol and $\tau_L = 6$ s, and $\alpha = 2$. Temperatures come from one common latent heat, $\Lambda = 3600$ K in units of the gas constant, which makes α exactly constant and the bubble point closed form, with normal boiling points of 80.0 and 105.75 °C for A and B. A total condenser feeds the drum, and the partial reboiler is an equilibrium stage with its sump, as shown in Equation (7.5).

$$\begin{aligned} \frac{\mathrm{d}M_D}{\mathrm{d}t} &= V - L - D, \qquad M_D\frac{\mathrm{d}x_D}{\mathrm{d}t} = V\,(y_1 - x_D) \\ \frac{\mathrm{d}M_B}{\mathrm{d}t} &= L_{20} - V - B \\ \frac{\mathrm{d}(M_B x_B)}{\mathrm{d}t} &= L_{20}x_{20} - V y_B - B x_B \end{aligned} \tag{7.5}$$

where x_B is the A in the sump and the product, $p_B = 100(1 - x_B)$ is the product purity in mol % B, and y_B is in equilibrium with x_B. At design the drum boils at 80.9 °C and the sump at 105.0 °C.

7.3.5 Flows, Pumps and Instruments

Each manipulated flow follows its setpoint q_{SP} through a first-order loop, $\tau_Q\,\mathrm{d}q/\mathrm{d}t = q_{SP} - q$ with $\tau_Q = 0.2$ min, and the boilup through the reboiler with $\tau_V = 1$ min. A pump loses suction as its vessel empties, so the delivered flow is the flow times $\mathrm{clip}(M/(0.02M_{cap}), 0, 1)$ of the vessel it draws from. Level transmitters lag 0.1 min with noise of 0.05 % of span (0.1 % on the drum and sump), flow meters lag 0.05 min with 0.3 kmol/h of noise, and the reactor thermowell lags 0.5 min. The two purity analyzers are continuous by default, a 1 min sample line with 0.02 mol % of noise, and the Advanced options replace them with a chromatograph that draws a sample every 2, 5 or 10 min through the same 1 min sample line and reports it one cycle later. A measurement packet carries only these readings to the controllers, and a packet holding anything beyond them fails the application's own test suite. The logged flow columns are the delivered flows rather than the meter readings.

7.3.6 Parameters, Starting State and Integration

Table 7.2 gives every parameter. A case begins at the steady state of its own settings. The analytical steady state of Section 7.4 supplies the reactor and the flows, a damped Newton solution of the

Table 7.2 Parameters (A Adjustable in the Parameters Panel, H Hidden From It, S Written Into the Source)

Symbol	Meaning	Value	Unit		Symbol	Meaning	Value	Unit	
V_{vessel}	reactor vessel volume	200	m^3	A	α	relative volatility, A to B	2	—	A
ρ	liquid molar density	4	$kmol/m^3$	H	$T_{b,B}$	normal boiling point of B	105.75	°C	H
k_0	rate constant at T_0	0.716	h^{-1}	A	Λ	latent heat over R	3600	K	H
T_0	design reactor temperature	70	°C	H	N	theoretical trays	20	—	S
E_a	activation energy	60	kJ/mol	A	f	feed tray, from the top	9	—	H
ΔH_r	heat of reaction (jacket)	−40	kJ/mol	A	$\bar{M}$	tray holdup at the reference	4	kmol	A
c_p	liquid heat capacity	0.16	kJ/(mol K)	H	τ_L	tray hydraulic lag	0.1	min	A
UA	jacket conductance	150	kW/K	A	$\bar{L}, \bar{F}$	weir reference reflux, feed	615, 700	kmol/h	S
T_f	feed temperature (jacket)	50	°C	H	$M_{D,cap}$	drum capacity	100	kmol	A
τ_J	coolant response	3	min	H	$M_{B,cap}$	sump capacity	120	kmol	A
τ_T	isothermal temperature loop	5	min	S	V_{flood}	boilup at flooding	1650	kmol/h	A
$L_{R,max}$	reactor high-level alarm	90	%	A	θ_r	line transit at D_{ref}	5	min	A
τ_Q	flow loop response	0.2	min	H	D_{ref}, N_p	line reference recycle, cells	500, 10	kmol/h, —	S
τ_V	reboiler response	1.0	min	H	P	column pressure	1.01325	bar	S

column at those flows supplies every stage's composition, and a bisection on the reflux makes the distillate meet its purity. The largest derivative at $t = 0$ is 4.9×10^{-9} in the state's units per minute. Settings with no steady state start from the design. Integration uses a fixed 2 s step of the classical fourth-order Runge-Kutta scheme, control executes every 6 s, and the historian writes one row per 5 min, on the engine of Section 1.6.1. Halving the step moves the recycle, the boilup, both purities and the reactor's A in case 1 at 6 h by 3.5×10^{-11} relative. Over a disturbed test run with structure switches, an activity loss, a feed composition change, a chromatograph and the steam limit, the moles and the A held plus the ledgers close to 1.0×10^{-11} of the throughput.

7.3.7 Assumptions

One irreversible first-order reaction at constant molar density stands for the chemistry, and a reversible or second-order reaction would move the steady curve without changing its shape. The column is binary with constant relative volatility and constant molar overflow, which is the same assumption as one latent heat, and the pressure, the condenser and the reboiler's heat transfer are ideal. Flooding is a fixed ceiling on the boilup rather than a hydraulic model, so the model has no flooded state. At the ceiling the trays keep their full separation and their weir holdups, and what the plant loses is the vapor the product loop asks for. A real column loses efficiency gradually as it approaches flooding, and its limit moves with the liquid load. The recycle line's ten cells add some mixing to the delay, where a real line is closer to plug flow. The flow loops are fast first-order lags whose own tuning is outside the lesson, and the reactor temperature is held ideally unless the jacket is chosen.

7.4 One Steady State and Three Handles

7.4.1 The Curve From Two Balances

With the holdup M_R and both purities held, the steady state follows from the balances alone. Overall, $F_0 = B$, because the reaction conserves moles. The A in the fresh feed either reacts or leaves in the product, and the column's feed leaves as recycle or product, as shown in Equation (7.6).

$$z^* = \frac{F_0(z_0 - x_B)}{kM_R}, \qquad D^* = F_0 \frac{z^* - x_B}{x_D - z^*}$$
$$F_0^{pole} = \frac{kM_R x_D}{z_0 - x_B} \tag{7.6}$$

These are equations 1 to 3 of Larsson, Govatsmark, Skogestad and Yu [8], which give z^*, the column feed $F = F_0 + D^*$ with z^* eliminated, and the holdup below which no steady state exists, here rearranged as the fresh feed at which z^* would reach x_D. With an ideal separator, $x_D = 1$ and $x_B = 0$ with $z_0 = 1$, the recycle reduces to Equation (1.3) and the pole to $F_0 = kM_R$. Their own design,

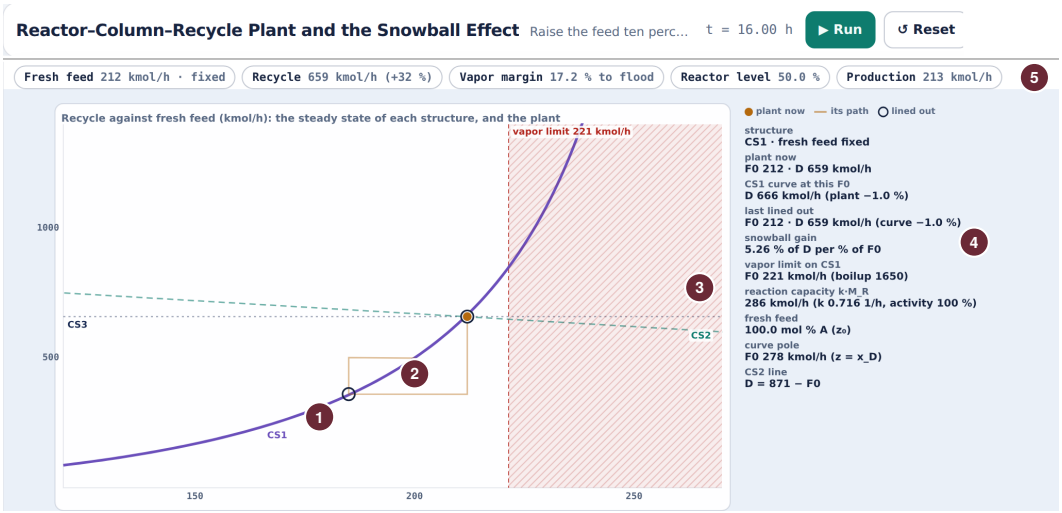

Figure 7.2 The snowball curve of case 2 at $t = 16$ h on the Snowball view, CS1, a composite of the application's header above its Snowball view, with the control panel and the trends cut. The CS1 curve (1), CS2's line $D = 871 - F_0$ and CS3's flat line are drawn with the region beyond the vapor limit at 221 kmol/h hatched (3). The plant's path (2) runs from the design down to 185 kmol/h and up to 212 kmol/h (Table 7.5), with a ring where the fresh feed and the recycle held within 1 % for an hour. The run places the rings at 360.8 kmol/h at 185, 0.84 % above the curve, and 659.3 at 212, 1.00 % under it and still climbing. The panel (4) reads the curve's 666, the gain 5.26 % per percent, the reaction capacity kM_R of 286 kmol/h and the pole at 278 kmol/h. The badges (5) read a recycle 32 % above its start, a vapor margin of 17.2 % to flooding and 213 kmol/h of production.

Table 7.3 Five Quantities Called a Snowball Gain on One Plant

Definition	Value	Where it appears
local elasticity $\mathrm{d}\ln D^*/\mathrm{d}\ln F_0$ on the CS1 curve, % per %	4.61 at design, 4.03 and 5.26 at 185 and 212, 2.63 with the vessel doubled	scorecard tile, Snowball view, cases 2 and 6
local slope $\mathrm{d}D^*/\mathrm{d}F_0$, kmol/h per kmol/h	11.5 at design	Model dialog
secant after ±2 % production steps lined out 32 h, kmol/h per kmol/h	CS1 11.55, CS2 by the holdup −1.00 (−0.40 % per %)	Instructions, application tests
finite step, steady $\Delta D/D$ for +10 % of F_0	+63.4 % (6.3 % per %)	case 1, Model dialog
secants between points not yet lined out, kmol/h per kmol/h	9.4 down, 11.1 up (curve: 9.5, 11.4)	case 2

460 kmol/h of feed at 0.9 mole fraction A, $x_B = 0.0105$, $k = 0.341\,\mathrm{h}^{-1}$ and 2800 kmol, gives their $z^* = 0.43$ (0.4285). At the app's design, $kM_R = 286.4\,\mathrm{kmol/h}$, $z^* = 68.44\,mol\,\%$, $D^* = 500.19\,\mathrm{kmol/h}$, so $D/F_0 = 2.50$, and the pole is at 277.6 kmol/h. The dynamic model is not built from Equation (7.6). Stepped from the design and run 32 h, its recycle lands within 0.005 % of D^* at seven fresh feeds from 185 to 220 kmol/h, as shown in Figure 7.3a. The application draws the curve on its Snowball view with the plant's path over it, as shown in Figure 7.2.

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

7.4.2 Five Quantities Called a Snowball Gain

Differentiating Equation (7.6) at a fixed holdup, where $\mathrm{d}\ln z^*/\mathrm{d}\ln F_0 = 1$, gives the elasticities of Equation (7.7).

$$\begin{aligned}\frac{\mathrm{d}\ln D^*}{\mathrm{d}\ln F_0} &= 1 + \frac{z}{z - x_B} + \frac{z}{x_D - z}\\ \frac{\mathrm{d}\ln D^*}{\mathrm{d}\ln k} &= -\left(\frac{z}{z - x_B} + \frac{z}{x_D - z}\right)\end{aligned} \tag{7.7}$$

At design the first is 4.61 and the second −3.61. With the ideal separator of Section 1.4.1 the first is the elasticity of Equation (1.3). The application and its texts quote five different numbers under the name snowball gain, and Table 7.3 names each by its definition. The local elasticity is a percent per percent, and the local slope $\mathrm{d}D^*/\mathrm{d}F_0 = 11.5$ is the same tangent in kmol/h per kmol/h, larger by D/F_0. A finite step reads more than either, because the curve steepens on the way, and 10 % of fresh feed is 63.4 % of recycle at steady state. The secant over ±2 % steps lined out, 11.55, sits just above the tangent because the curve is convex. The recycle purity enters through the same denominator, and each mol % of x_D lost at design costs 18.8 kmol/h of recycle at steady state.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.4.3 The Other Two Structures

A structure that holds the holdup and both purities lands on Equation (7.6) whatever flow it fixes, so the structures differ in which handle moves production. Fixing the effluent (CS2) or the recycle and the effluent (CS3) gives the relations of Equation (7.8).

$$\begin{aligned}&\text{CS2:}\quad D = F - F_0, \qquad F_0 = F\,\frac{x_D - z}{x_D - x_B}\\ &\text{CS3:}\quad F_0 = F - D, \qquad x_D = \frac{Fz - F_0 x_B}{D}\end{aligned} \tag{7.8}$$

On CS2 the recycle is the fixed effluent less the product, so production moved by the holdup moves the recycle the other way, −1 kmol/h per kmol/h, or −0.40 percent per percent at design. Lined out, five holdups from 46.6 to 57.2 % of level put the plant on that line to 0.018 kmol/h. The effluent at a fixed holdup is a different handle. Three effluents of 670, 730 and 760 kmol/h settle the plant at fresh feeds of 197.5, 202.3 and 204.5 kmol/h with recycles within 0.016 kmol/h of the steady state the balances predict, and those steady states lie on the CS1 curve, as Figure 7.3a shows. This measurement is why CS2's production handle in the app is the reactor level setpoint rather than the effluent. On CS3 production is the difference of two flow setpoints, and the recycle purity floats to whatever closes the column's balance, which on a lined-out run at 704 kmol/h of effluent it does to 3×10^{-4} mol %.

7.4.4 Where the Column Stops the Snowball

The boilup the column needs grows with the recycle. From the column's own steady state at both purities, the CS1 steady states at 216, 220, 221.4, 223, 226 and 230 kmol/h of fresh feed need 1485, 1604, 1649, 1703, 1812 and 1974 kmol/h of boilup, so the largest fresh feed the column can boil under its 1650 kmol/h flooding rate is 221.4 kmol/h, 10.7 % above the design. The boilup controller never asks for more than $V_{cap} = \min(V_{flood}, V_{steam})$, which is where the model holds the boilup when the snowball reaches the wall. The design sits at 67.6 % of flooding on purpose. At a design of 80 % the wall would come at 6.3 % more fresh feed, and the 10 % step of case 1 could not show the snowball's size before reaching it. Beyond the wall the product loop has no vapor left, A breaks into the product, and that A no longer returns, which is the end Larsson and coauthors describe for a real snowball [8].

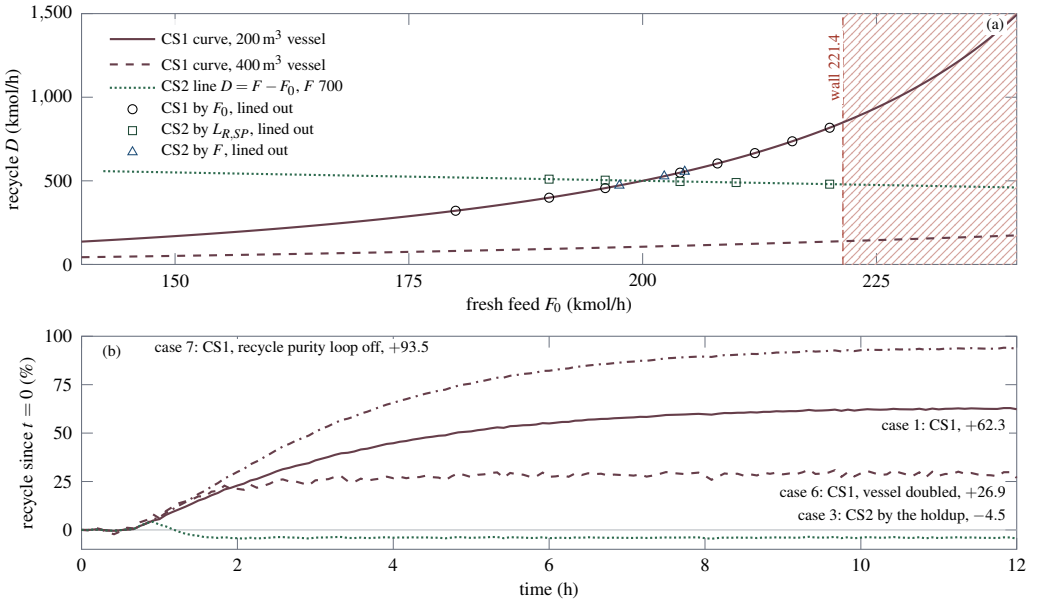

Figure 7.3 The steady state and the path to it. (a) Recycle against fresh feed from Equation (7.6) at the design holdup (solid) and with the vessel doubled (dashed), CS2's line (dotted), the vapor wall at 221.4 kmol/h (thin dashed, hatched beyond), and the plant lined out 32 h after steps from the design under three handles (circles, squares and triangles). (b) The recycle since $t = 0$ for the same production increase in cases 1 (solid), 3 (dotted), 6 (dashed) and 7 (dash-dotted) as shipped, with the values at 12 h.

Table 7.4 What Each Structure Fixes and Pairs

Loop	CS1	CS2	CS3
flow control	F_0	F	F and D
reactor level →	F	F_0	F_0
drum level →	D	D	L
recycle purity x_D →	L	L	floats
product purity p_B →	V	V	V
sump level →	B	B	B
production handle (TPM)	$F_{0,SP}$	$L_{R,SP}$	F_{SP}

7.5 The Loops as Coded

7.5.1 Pairings by Structure

The three structures are listed in Table 7.4. CS1 is the conventional structure, with nothing fixed inside the recycle loop. CS2 fixes the effluent and moves production by the reactor holdup, which is the structure the literature calls Luyben's [8]. CS3 fixes the recycle itself, which leaves no flow to hold the drum except the reflux, and so lets the recycle purity float. Each structure has its own production setpoint, marked TPM on the Plant view, rather than one setpoint routed to a different valve. The structure can change during a run. One PI controller exists per manipulated flow, the reactor level having one on F and one on F_0 and the drum level one on D and one on L, and every controller that is not driving its flow tracks it at every execution, so a switch is bumpless. A test makes each of the six possible switches at 60 min and finds no flow moving more than 1 % of its design value over the next control period. Manual mode holds the five level and purity controllers at their stations, while the structure's flow setpoints stay in force.

7.5.2 The Shared PI Law

Every loop uses the velocity-form PID of Section 1.6.2, Equations (1.11) to (1.14), without derivative action, with the clamp on the emitted output only and back-calculation whose tracking time equals τ_I. Its bias u_b is the flow when the controller starts, and the control period is $\Delta t = 0.1$ min. The action is $a = -1$ for a level held by an outflow (the reflux included) and $a = +1$ for the reactor level held by the fresh feed and for both purities. Integral action is dropped when the integral time is set to 600 min or longer. The reactor level moves 50 kmol/h per percent with $\tau_I = 30$ min, the drum and sump levels 30 kmol/h per percent with 30 min, both purity loops 40 kmol/h per mol % with 15 min, and the jacket's temperature loop 2 K of coolant per kelvin with 20 min. The level controllers see their transmitter through a 0.5 min first-order filter, which cuts the production jitter the level noise causes. The boilup's upper limit is V_{cap}, and the valve ranges of Table 7.2 limit the rest.

7.5.3 Why the Default Analyzer Is Continuous

A sampled analyzer is the realistic instrument for a purity, and the Advanced options offer one. With a 5 min cycle at the default purity tuning, case 1's product swings between 65.6 and 100 mol % B after 4 h, and the boilup sits at the flooding rate for 279 min from 2.08 h, so the snowball cases would spend hours off specification for a reason unrelated to the structure. Both purity loops at 20 kmol/h per mol % and 30 min hold the product between 96.3 and 97.8 mol % on the same chromatograph without reaching the flooding rate. The continuous analyzers are therefore the default, and the chromatograph stays an Advanced option. Volume 5, Chapter 3 measures what an analyzer's lag costs a single composition loop on a binary column.

7.5.4 The Line Delay and the Reactor Level

On CS1 the reactor level loop acts through the recycle, because more effluent is more column feed and the product purity loop boils most of it back up. With the level loop out of the way, nine tenths of a 20 kmol/h effluent step return as recycle within 11.9 min, at once in flow because the line runs full. The line's composition delay therefore barely touches the level loop. Case 1's largest reactor level error is 0.850, 0.853 and 0.870 % with a 1, 5 and 30 min line. At three faster tunings, up to 200 kmol/h per percent with 5 min, the largest error moves by less than 0.003 points between a 1 and a 30 min line, and at all four tunings the level's standard deviation from 8 to 12 h stays between 0.007 and 0.022 points.

7.5.5 Status Flags and the Three Models

FLOODING marks the boilup controller at the flooding rate and STEAM LIMIT at a lower steam supply. OFF SPEC marks the true product more than 0.5 mol % under its setpoint, and LEVEL HIGH the reactor above its 90 % alarm. These are annunciators, not trips, and the flooding rate, the alarm and the 98 mol % specification are teaching values, while real ones come from a tray rating [16], the vessel's design and the product's sales specification. The application keeps three models apart. The plant is Section 7.3. The analytical model of Section 7.4 is an overlay computed from the true activity and temperature, drawn on the Snowball view and logged, and no controller reads it. The controllers carry no model at all.

7.6 Nine Cases on the Running Plant

The nine cases are listed in Table 7.5. The numbers below come from the engine of Section 1.6.1 run without a browser at the default seed, each case loaded as the application loads it and each move made at the instant the text names. Cases 4, 6, 7 and 8 compare two runs, the second started by Reset after the change, because a structure, vessel or setpoint chosen before a run acts from $t = 0$ as a starting condition. The scorecard the application shows for each case at its reading instant is collected in Table 7.6, and Figure 7.4 draws its recycle and vapor margin columns. Every case loads identically on first selection, from the list after another case, and after Reset.

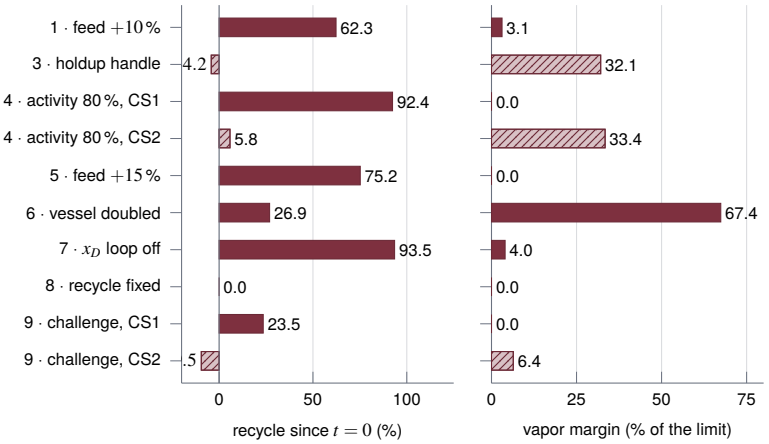

Figure 7.4 The scorecard of Table 7.6 for ten of its rows, at each row's reading instant: dark bars run on CS1, hatched bars on CS2 and the white bar on CS3. The margins of case 9 are to the steam limit, the others to flooding. Notice that the dark bars of cases 1, 4, 5, 7 and 9 end at or near a vapor margin of zero, while on the hatched bars the recycle barely moves.

Table 7.5 The Nine Cases

#	Structure and events	Teaching point
1	CS1, fresh feed 200 to 220 kmol/h at 30 min, 12 h	the snowball
2	CS1, fresh feed to 185 at 0.5 h, to 212 at 7.5 h, 16 h	the curve against the plant
3	CS2, level 50 to 57.2 % at 30 min	the holdup as the handle
4	CS1, activity to 80 % at 30 min, then CS2	who absorbs a kinetic loss
5	CS1, fresh feed +15 % at 30 min	the vapor wall, infeasible
6	CS1, 400 m^3 vessel, +10 %	holdup against recycle
7	CS1, recycle purity loop off, +10 %	the column's loops are plantwide
8	CS3, effluent 700 to 720 kmol/h over 0.5–2.5 h	what a fixed recycle costs
9	Challenge: steam 1200, target 240 at 15 min, activity 90 at 3 h, 8 h	the structure and the holdup

7.6.1 The Snowball and the Curve

Case 1 raises the fresh feed by 10 % at 30 min on CS1. The reactor's A climbs from 68.4 to 75.2 mol % by 12 h, the recycle from 500 to 812 kmol/h, still short of the curve's 817, and the boilup from 1115 to 1599 kmol/h, which takes the vapor margin from 32.4 to 3.1 % of the flooding rate, as shown in Figure 7.3b. The product loop falls behind the climb at first. The product is under 97.5 mol % B from 0.95 h, reaches 97.17 mol % at 1.8 h, and OFF SPEC clears at 4.7 h. At the new steady state the margin would be 2.8 %.

> **TRY IT IN THE APP** **1 · Fresh Feed +10 % on the Conventional Structure**
>
> Load the scenario, press Run and watch the Recycle since $t = 0$ tile after the fresh feed steps up by a tenth at 30 min. By 12 h the recycle has risen 62.3 %, still climbing, while the Vapor margin badge has fallen from 32.4 to 3.1 % of the flooding rate.

Case 2 takes the fresh feed down to 185 and then up to 212 kmol/h (Table 7.5) on the Snowball view, as shown in Figure 7.2. At 7.4 h the recycle is at 360 kmol/h, 0.6 % above the curve's 358 and still closing, and at 16 h it is at 659, 1.0 % under the curve's 666 and still climbing. The recycle stays within 2 % of the curve from 5.0 h after the first step and from 7.4 h after the second, because each new point is approached at the pace of the reactor and the level loop. Between the points the plant reached, 15 fewer kmol/h of fresh feed took 140 kmol/h off the recycle and 27 more added 299,

Table 7.6 The Application's Scorecard at Each Case's Reading Instant, With the Lowest True Product Purity of the Run

Case	Settings and instant	Recycle since $t = 0$ (%)	Production (kmol/h)	Vapor margin (% of the limit)	Lowest p_B (mol % B)	Also
1	as shipped, 12 h	+62.3	220	3.1 to flood	97.17	OFF SPEC 0.95–4.7 h
2	as shipped, 16 h	+31.8	213	17.2 to flood	96.98	OFF SPEC 7.8–11.5 h
3	as shipped, 6 h	−4.2	221	32.1 to flood	97.86	production dips to 178.7 at 0.92 h
3	effluent to 760 kmol/h at 6 h, 12 h	+6.4	225	26.5 to flood	96.79	60 more effluent, 4.7 more product
4	as shipped (CS1), 12 h	+92.4	200	0.0 to flood	91.67	FLOODING from 3.80 h
4	CS2 and Reset, 6 h	+5.8	170	33.4 to flood	97.88	production −14.9 %
5	as shipped, 12 h	+75.2	230	0.0 to flood	95.01	FLOODING from 3.96 h, 11.4 kmol/h of A in B
5	fresh feed to 220 at 6 h, 12 h	+64.8	220	1.9 to flood	95.28	crosses 97.5 mol % at 7.62 h; one hour above it off flooding by 9.0 h
6	as shipped, 12 h	+26.9	219	67.4 to flood	97.58	starts at 106 kmol/h of recycle
7	as shipped (loop off), 12 h	+93.5	220	4.0 to flood	97.28	x_D 91.83 mol %
8	as shipped, 8 h	0.0	219	0.0 to flood	92.68	x_D 99.32 mol %, FLOODING from 4.47 h
8	level 56.5 % and Reset, 8 h	0.0	219	30.7 to flood	97.58	x_D at most 95.79 mol %
9	as shipped (CS1), 8 h	+23.5	239	0.0 to steam	79.57	steam limit from 0.75 h, fails
9	suggested, CS2 at 75 %, 8 h	−9.5	245	6.4 to steam	97.87	passes

secants of 9.4 and 11.1 against the curve's own 9.5 and 11.4. In percent the curve's slope is 4.03 at 185 and 5.26 at 212 kmol/h.

7.6.2 The Holdup as the Handle

Case 3 runs on CS2 and raises the reactor level setpoint from 50 to 57.2 %, the holdup that the balances give for 10 % more product. Production first dips to 178.7 kmol/h at 0.92 h, 25 min after the step, because the fresh A that fills the reactor reaches the column first, and it is back above 200 by 1.22 h. By 6 h production is 220.8 kmol/h, the reactor's A has fallen from 68.4 to 65.8 mol % and the recycle from 500 to 479 kmol/h, with the boilup almost unchanged at 1120 kmol/h. The handle has a price. The holdup rose 14 % for 10 % of production, the snowball in the reactor holdup that Wu and Yu found on Luyben's structure [8, 12]. Raising the effluent to 760 kmol/h at 6 h instead brings the CS1 curve back. The recycle is at 536 kmol/h within half an hour, and at 12 h production is 225.4 kmol/h with 532 kmol/h of recycle, 60 more kmol/h of effluent for about 5 of product.

7.6.3 Activity Lost on Two Structures

Case 4 cuts the catalyst activity to 80 %. On CS1 the fresh feed stays at 200 kmol/h, so the recycle has to make up the lost conversion. The product slips under 97.5 mol % at 1.2 h, the boilup reaches the flooding rate at 3.8 h, and by 12 h the product is at 91.76 mol % with 962 kmol/h of recycle, while the steady state would need 1767 kmol/h. On CS2, after Reset, the same loss costs production instead, 170.3 kmol/h at 6 h (−14.9 %), with the recycle at 529 kmol/h and the boilup at 1099. Switched to CS2 during the flooded run without Reset, the structure fixes the effluent where the snowball left it, about 1130 kmol/h, and the column stays at its flooding rate to 12 h.

> **+ DID YOU KNOW? The Free Flow Takes the Catalyst Loss**
>
> From the simulator (Table 7.6). With a fifth of the catalyst activity gone, the fixed fresh feed of CS1 leaves the loss to the recycle, which has risen 92.4 % by 12 h, with the column flooding from 3.80 h. On CS2, after Reset, the recycle has risen only 5.8 % at 6 h, and production has fallen 14.9 % instead.

7.6.4 The Vapor Wall

Case 5 asks for 15 % more fresh feed, beyond the 10.7 % the wall allows. The boilup reaches the flooding rate at 3.96 h with the recycle at 856 kmol/h, and from then the column sends no more A overhead. A goes into the product instead, 11.4 kmol/h of it at 12 h, where the product is at 95.04 mol %. That A no longer returns, so the recycle stalls near 877 kmol/h, far short of the 1082 kmol/h the curve gives at 230. Doubling the product purity gain leaves the product at the same 95.04 mol %, because the loop asks for vapor that is not there. With the fresh feed cut to 220 kmol/h at 6 h, the product first crosses 97.5 mol % at 7.62 h while still flooding, and by 9.0 h has held above that value off the flooding rate for an hour.

Pause. Try Exercise 7.3 before reading on.

7.6.5 Vessel Size and the Recycle Purity Loop

Case 6 doubles the vessel to 400 m^3 at the same level. The reactor starts at 34.2 mol % A with 106 kmol/h of recycle and 453 kmol/h of boilup, far from the curve's pole, and the same step raises the recycle 26.9 % by 12 h (28.9 % at steady state), against 62.3 % in case 1. The local elasticity at 200 kmol/h falls from 4.61 to 2.63. This is the design side of the snowball that the literature points to [2, 8]. Case 7 repeats case 1 with the recycle purity loop off and the reflux held. More A arriving at a fixed reflux lets B into the distillate, whose purity falls to 91.83 mol %, and since the reactor still needs the same z, more kmol of a leaner recycle must come back. The recycle rises 93.5 % against 62.3 % with the loop on, where the reflux rose 171 kmol/h to hold 95 mol %.

> **TRY IT IN THE APP** **7 · The Column's Purity Loops Are Part of the Recycle**
>
> Load the scenario, in which the recycle purity loop is off, run to 12 h and read the Recycle since $t = 0$ tile: the same step as case 1 now raises the recycle 93.5 %, because the recycle purity sags to 91.83 mol %. Switch the loop On, press Reset and run again to compare with the 62.3 % of case 1.

7.6.6 A Fixed Recycle

Case 8 runs on CS3 and ramps the effluent from 700 to 720 kmol/h over two hours. Production follows, 221.0 kmol/h at 6 h, and the recycle stays at 500 kmol/h. By Equation (7.8), at the design holdup 10 % more production would need a recycle of 107.5 mol % A, purer than pure A. The recycle purity climbs to 97.89 mol % at 3 h, the boilup reaches the flooding rate at 4.47 h, and at 8 h the product is at 92.68 mol % with the recycle at 99.32 mol %. With the level setpoint at 56.5 %, the holdup that keeps x_D at 95, the recycle purity peaks at 95.79 mol % and the boilup at 1154 kmol/h. A fixed recycle protects the column only while something else raises the conversion with production.

7.6.7 A Graded Day

The challenge limits the reboiler's steam to 1200 kmol/h of boilup, 7.6 % above the design's, raises the production target to 240 kmol/h at 15 min and takes a tenth of the catalyst activity away at 3 h. Over 8 h it judges production as twelve half-hour means of at least 235.2 kmol/h from 2 h, the true product at or above 97.5 mol % and the true recycle at or above 94.0 mol % at every step, the boilup never at its limit, and the reactor never above its alarm. Table 7.7 lists the settings tried. On CS1 the steady state at 240 kmol/h would need 2496 kmol/h of boilup, and as shipped the boilup reaches the steam limit at 0.75 h and the product falls to 79.57 mol %. On CS2 at a 75 % level the day passes, with 245.9 kmol/h in the worst half hour and the recycle falling from 500 to 453 kmol/h while production rises, and at seven other noise seeds the worst half hour stays between 245.7 and 245.9 kmol/h. The lowest passing CS2 level is 70.5 %, against the 72.3 % that makes 240 kmol/h at 90 % activity at steady state.

The size of the loss was chosen by measurement. With a fifth of the activity lost, CS2 needs 81.3 % of level at steady state, and the day passes only between 79 and 87.5 %, because at 88 % the overshoot breaks the 90 % alarm. With a tenth lost, the passing band is 70.5 to 87.5 %. Both teach the

Table 7.7 The Challenge Over 8 h, Default Seed

Settings	Worst (kmol/h)	Min p_B (mol %)	Min x_D (mol %)	At limit (min)	Max L_R (%)	Verdict
CS1, level 50 (shipped)	238.6	79.57	94.69	434.9	51.3	fails
CS2, level 75 (suggested)	245.9	97.87	94.55	0	77.4	passes
CS2, level 50	185.2	97.95	95.00	0	50.2	fails
CS2, level 70	234.9	97.87	94.58	0	72.3	fails
CS2, level 70.5	236.0	97.87	94.58	0	72.8	passes
CS2, level 85	266.5	97.87	94.53	0	87.4	passes
CS1, level 75	236.2	96.80	76.18	0	77.9	fails
CS1, level 75, slow loop[a]	198.7	97.73	92.37	0	78.3	fails
CS3, F 745, level 70	242.7	85.49	93.89	324.2	71.9	fails

[a] 10 kmol/h per percent and 120 min on the reactor level.

same lesson, that the structure and the holdup decide the day rather than the tuning, and the wider band leaves a first estimate room to pass.

7.7 Technology in Practice

What Measures It

Every production handle of this plant is a flow setpoint: the fresh feed in CS1, the effluent in CS2 and CS3, and the recycle as well in CS3. Each sits on a flow loop, often metered by a pressure-based element whose signal is characterized by a square root [17]. The chapter's flow meters lag 0.05 min with 0.3 kmol/h of noise and read as well at the bottom of their range as at the top, which an orifice meter does not. The two purities come from analyzers. A process chromatograph separates species by time, so cycle times of minutes are not uncommon even on line, and some measurement dead time is inevitable [17]. The chapter's analyzers are continuous by default, with the chromatograph an Advanced option.

> **TECHNOLOGY Sensor: The Orifice Meter Under Each Flow Loop**
>
> The most common pressure-based flow element is the orifice plate, a metal plate with a hole, sandwiched between two pipe flanges. The pressure drop it develops grows with the square of the flow, so the transmitter's signal is characterized by a square root before an indicator or controller uses it as flow. Accuracy falls toward the bottom of the range: the practical turndown of pressure-based meters is often no more than 3 to 1. A dulled orifice edge or uneven liquid columns in the impulse tubes then cause large errors at low flow, even with a perfectly calibrated transmitter [17].

What Moves It

In the model each manipulated flow follows its setpoint through a first-order loop of 0.2 min, which stands for a control valve under its own flow controller, and the boilup follows through the reboiler in 1 min. The boilup controller never asks for more than V_{cap} and uses back-calculation, so the vapor wall of case 5 arrives as a saturated output without windup, as in a real controller that receives back-calculation feedback from its final element [17].

> **TECHNOLOGY Final Element: The Reboiler Steam Valve at Its Limit**
>
> Boilup is moved by the steam valve of the reboiler. When the steam supply cannot deliver what the loop asks for, a controller with integral action winds up, driving its output to saturation in a futile attempt to open the valve further, and when steam returns the saturated output overshoots. On a trend the output sits at its maximum while the measurement stays off setpoint. Back-calculation from the final element stops the controller from driving the valve farther than it is able to respond. A steam flow loop under the temperature or composition loop takes up steam-side upsets first, and it must respond faster than the loop above it [17].

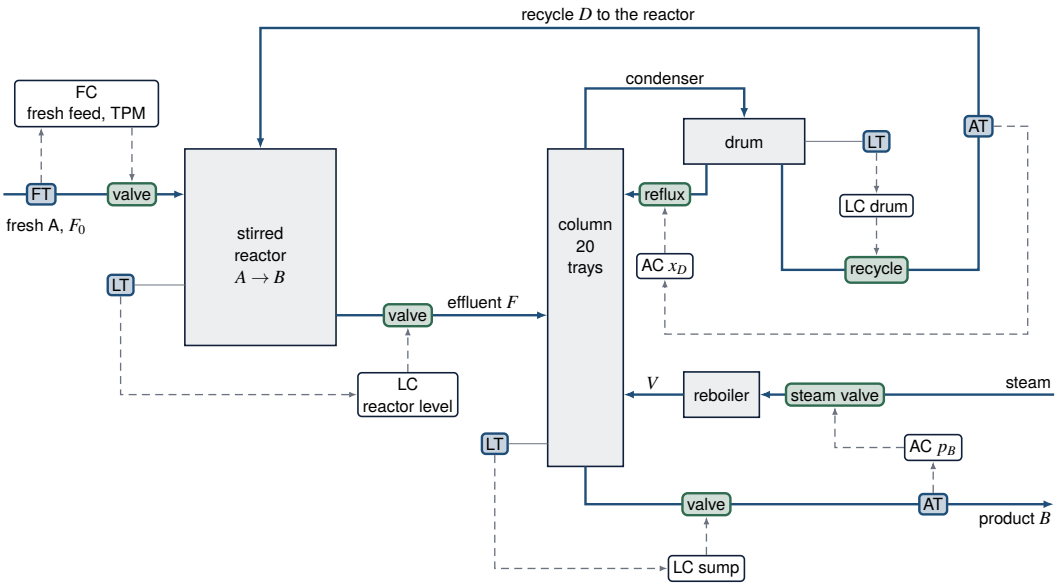

Blocks FC, LC and AC: velocity-form PI in the distributed control system; the app executes all of them every 6 s. CS2 and CS3 re-pair the same blocks (Table 7.4).

Figure 7.5 The loops of CS1 on plant hardware. The fresh feed is the throughput manipulator, held by FC through a flow transmitter on an orifice, and the reactor level is held by the effluent valve. The drum level moves the recycle, the recycle analyzer moves the reflux, the product analyzer moves the reboiler steam valve and the sump level moves the product valve, each through a flow loop on its valve that is not drawn. Notice that no setpoint of the recycle is fixed in CS1: it is whatever the drum level controller sends back.

Where the Controller Runs

In a plant these loops run in a distributed control system, whose supervisory level oversees localized controllers. Loops are often nested or cascaded, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [18]. Supervisory control in the process industry uses advanced regulatory control, a setpoint-based model predictive controller on top of the PID layer, or both, and usually no optimization layer exists [19]. The app integrates with a 2 s step and executes its controllers every 6 s, as Figure 7.5 shows for CS1. A structure is not a block of its own: one PI controller exists per manipulated flow, and those not driving their flow track it. Output tracking is also how a real controller picks up from its last manual value without a bump [17].

> **IN PRACTICE** **Throughput First, Then the Bottleneck**
>
> After a new plant starts up, the operating staff usually turn first to raising throughput, because the extra profit there is usually much larger than from running more efficiently at a given rate. Yet the design phase usually makes no effort to give the plant a control system that runs at maximum throughput. At that point some unit is the bottleneck, and its maximum flow becomes an active constraint, which gives the rule of locating the throughput manipulator at the bottleneck unit. The design team should also say where the bottleneck is expected to appear when throughput rises [3].

What Is Modernizing It

Two directions stand out. The first designs self-optimizing behavior into simple, mostly single-loop regulatory strategies from the start. Against a centralized real-time optimizer, such strategies give simplicity, robustness and reliability, and they have found particular appeal in industry [3]. The

second moves the throughput manipulator automatically, where a handle moved to a new bottleneck would otherwise need several inventory loops rearranged. In a bidirectional structure each inventory has two controllers, one with a high setpoint on its inflow and one with a low setpoint on its outflow, and min-selectors choose between them. With unreachable flow setpoints it finds the bottleneck by itself and maximizes throughput [19]. The chapter's structures are switched by hand, and the bottleneck structure that would set production at the column's vapor rate is not offered (Section 7.8).

7.8 Limitations

Every size and setting here is a generic teaching value, chosen to give a consistent design point at 67.6 % of flooding, not a fit to any plant. The column is binary with constant relative volatility and constant molar overflow, its pressure is fixed, and flooding is a fixed ceiling on the boilup, so the model cannot show the gradual loss of efficiency near flooding, entrainment or the liquid limit of a downcomer. One irreversible first-order reaction stands for the chemistry, and side reactions, an equilibrium limit and an energy balance by default are absent. The recycle line mixes as well as delays.

The structures are three fixed pairings, so the bottleneck structure that sets production at the column's vapor rate [14], the balanced structures of Wu and Yu [12] and the reflux ratio as a self-optimizing variable [8] are not offered, and no supervisory layer moves the holdup or the temperature with production. A real-time optimizer that moves the setpoints of a reactor's loops is the subject of Volume 13, Chapter 11. The PI loops are tuned for continuous analyzers, the flow loops are ideal, and the status flags do not trip anything. The disturbances arrive as steps and ramps. A flow setpoint that the structure in force does not use follows the flow it names while the plant runs in Auto. In Manual the panel does not refresh it, so it keeps the value it held when Manual was selected.

> ⚠ **OPEN CHALLENGE** **A Structure Designed, Not Inherited**
>
> Heuristics, rules from case studies and optimization algorithms have all contributed to plantwide design over some forty years, yet much of that research has not gained enough acceptance in industry to have a profound influence. Process design and control design are still commonly segregated, and plants are run unit by unit with simple loops that operators understand. On an older plant even setting the throughput somewhere other than the feed can take years of work to win acceptance. What is missing is a formal yet simple procedure that ties the inventory structure to the plant's economics [3].

7.9 Conclusion

This chapter modeled the reactor, column and recycle plant of the plantwide control literature from balances that close to 1×10^{-11} of the throughput, with the control structure as a switch and the analytical steady state of Larsson and coauthors drawn beside the running plant. With the fresh feed fixed, 10 % more feed raises the recycle 62.3 % in 12 h toward 63.4 %. The five quantities called a snowball gain on this one design point, 4.61, 11.5, 11.55, 63.4 and 9.4 to 11.1, are consistent once each is named by its definition. Lined out, every structure that holds the holdup and both purities lands on one curve, so a structure chooses the production handle. The holdup moves the recycle -1 kmol/h per kmol/h at the price of 14 % more holdup for 10 % more product, the effluent at a fixed holdup brings the curve back, and a fixed recycle pushes the column's purity past pure A. The column's vapor limit ends the snowball by putting A into the product, and a recycle purity loop left open raises the recycle's rise to 93.5 %. The model's flooding is a ceiling and its chemistry one reaction. A fuller model would add the bottleneck structure with production set at the boilup, a reflux-ratio loop, a supervisory layer that moves the holdup with production, and a hydraulic flooding model. Where production is set in a recycle plant decides where the snowball goes.

KEY POINTS

- With the reactor holdup held and the fresh feed fixed, ten percent more feed raises the recycle 62.3 % in 12 h.
- Lined out, every structure that holds the holdup and both purities lands on one curve, so a structure chooses the production handle.
- The holdup handle moves the recycle -1 kmol/h per kmol/h at the price of 14 % more holdup for 10 % more product.
- The column's vapor limit ends the snowball by putting A into the product, and a stronger purity loop cannot recover it.

Where else this applies. A column that runs out of steam returns in Volume 10, Chapter 10, where a steam header trip pushes out of contract the one product with no controller on it.

References

[1] J. Morud and S. Skogestad, "Dynamic behaviour of integrated plants," *Journal of Process Control*, vol. 6, no. 2-3, pp. 145–156, 1996.

[2] T. Larsson and S. Skogestad, "Plantwide control - A review and a new design procedure," *Modeling, Identification and Control: A Norwegian Research Bulletin*, vol. 21, no. 4, pp. 209–240, 2000.

[3] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[4] P. S. Buckley, *Techniques of Process Control*. New York: Wiley, 1964.

[5] R. M. Price and C. Georgakis, "Plantwide regulatory control design procedure using a tiered framework," *Industrial & Engineering Chemistry Research*, vol. 32, no. 11, pp. 2693–2705, 1993.

[6] E. M. B. Aske and S. Skogestad, "Consistent inventory control," *Industrial & Engineering Chemistry Research*, vol. 48, no. 24, pp. 10 892–10 902, 2009.

[7] W. L. Luyben, "Snowball effects in reactor/separator processes with recycle," *Industrial & Engineering Chemistry Research*, vol. 33, no. 2, pp. 299–305, 1994.

[8] T. Larsson, M. S. Govatsmark, S. Skogestad, and C. C. Yu, "Control structure selection for reactor, separator, and recycle processes," *Industrial & Engineering Chemistry Research*, vol. 42, no. 6, pp. 1225–1234, 2003.

[9] E. R. Gilliland, L. A. Gould, and T. J. Boyle, "Dynamic effects of material recycle," in *Preprints of the Joint Automatic Control Conference*, Stanford, CA, 1964, pp. 140–146.

[10] W. L. Luyben, "Dynamics and control of recycle systems. 1. Simple open-loop and closed-loop systems," *Industrial & Engineering Chemistry Research*, vol. 32, no. 3, pp. 466–475, 1993.

[11] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, "Plantwide control design procedure," *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

[12] K.-L. Wu and C.-C. Yu, "Reactor/separator processes with recycle—1. Candidate control structure for operability," *Computers & Chemical Engineering*, vol. 20, no. 11, pp. 1291–1316, 1996.

[13] S. Skogestad, "Plantwide control: the search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[14] ——, "Control structure design for complete chemical plants," *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[15] ——, "Dynamics and control of distillation columns," *Chemical Engineering Research and Design*, vol. 75, no. 6, pp. 539–562, 1997.

[16] A. R. C. Souza, M. J. Bagajewicz, and A. L. H. Costa, "Improved correlations for threshold flooding and entrainment in sieve trays in distillation/absorption columns," *Industrial & Engineering Chemistry Research*, vol. 64, no. 4, pp. 2256–2273, 2025.

[17] 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.

[18] 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.

[19] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

7.10 Exercises

7.1 Concept check · 2 min

In case 1 the fresh feed rises 10 % on CS1 while the reactor holdup and both purities are held, and the steady recycle rises 63.4 %. Which explanation is right?

(a) The reactor converts more A in the same holdup only by running richer in A, and the column returns that richer effluent's A

(b) The column separates less well under the larger load, so B circulates in the recycle and dilutes it

(c) The reactor level controller overshoots and lifts the effluent above what the production change needs

(d) The recycle line delays the returning A, so the recycle must grow to carry the A held in the line

7.2 Concept check · 1 min

At the design point the Snowball gain tile reads 4.61 % per percent, while the Model dialog gives a slope of 11.5 for the same curve at the same point. Why do the two differ?

(a) The tile is the secant of a 10 % step and the dialog is the tangent at the design point

(b) Both are the tangent of the CS1 curve, the tile in percent per percent and the dialog in kmol/h per kmol/h

(c) The tile is computed on CS2's line and the dialog on the CS1 curve, so they describe different structures

(d) The tile uses the analyzed purities and the dialog the setpoints, and the two differ by the analyzer noise

7.3 Concept check · 2 min

In case 5 the boilup reaches the 1650 kmol/h flooding rate and the product purity falls to 95.04 mol %. Doubling the Product purity gain leaves the product at the same 95.04 mol %. Why?

(a) The analyzer's sample line hides the improvement until after the 12 h reading

(b) Back-calculation freezes the integral at the limit, so only the small proportional term responds

(c) The boilup is already at its ceiling, so the loop asks for vapor that is not there and A leaves in the product

(d) The larger gain makes the loop oscillate about the ceiling, and the mean purity stays where it was

7.4 Theory and calculation · 6 min

The design reactor holds $M_R = 400\,\text{kmol}$ with $k = k_0 = 0.716\,\text{h}^{-1}$ at 70 °C and full activity. The fresh feed is pure A ($z_0 = 1$), the recycle purity is $x_D = 0.95$ and the product carries $x_B = 0.02$ of A. Use Equation (7.6).

(a) Find the reaction capacity kM_R, the reactor composition z^* and the steady recycle D^* at $F_0 = 200\,\text{kmol/h}$. *Checkpoint:* 286.4 kmol/h, 68.44 mol % and 500.2 kmol/h.

(b) Find the fresh feed at the curve's pole. *Checkpoint:* 277.6 kmol/h.

(c) Find the steady recycle at $F_0 = 220\,\text{kmol/h}$ and its rise over the design value. *Checkpoint:* 817.5 kmol/h, +63.4 %.

7.5 Theory and calculation · 6 min

Continue at the design point of the previous exercise, $z = 0.6844$, $x_D = 0.95$, $x_B = 0.02$, $F_0 = 200\,\text{kmol/h}$ and $D = 500.2\,\text{kmol/h}$.

(a) Evaluate both elasticities of Equation (7.7). *Checkpoint:* 4.61 and -3.61.

(b) Convert the first into the slope $\mathrm{d}D^*/\mathrm{d}F_0$ in kmol/h per kmol/h. *Checkpoint:* 11.5.

(c) Differentiate D^* of Equation (7.6) with respect to x_D at fixed F_0 and M_R, and find the recycle each mol % of recycle purity lost costs at steady state. *Checkpoint:* 18.8 kmol/h per mol %.

7.6 Theory and calculation · 6 min

Production is to rise from 200 to 220 kmol/h with both purities held ($x_D = 0.95$, $x_B = 0.02$, $z_0 = 1$). The design reactor converts $kM_R = 286.4\,\text{kmol/h}$ at a 50 % level.

(a) On CS2 the effluent stays at 700.19 kmol/h. Find the steady reactor composition and the reactor level setpoint that gives 220 kmol/h. *Checkpoint:* 65.78 mol % and 57.2 %.

(b) Find, for the CS2 move, the rise of the holdup in percent and the steady recycle. *Checkpoint:* $+14.4$ % and 480.2 kmol/h.

(c) On CS3 the recycle stays at 500 kmol/h and the effluent goes to 720 kmol/h at the design holdup. Find the recycle purity that would close the column's balance, and the level that keeps it at 95 mol %. *Checkpoint:* 107.5 mol % and 56.5 %.

7.7 App challenge · 8 min

Load scenario 1, Fresh feed +10 % on the conventional structure, and change no setting. Run to 12 h (720 min) at the default seed; the default 60 min/s plays it in 12 s.

(a) Read the Recycle since $t = 0$ and Vapor margin tiles at 12 h, and in the exported file the columns `recycle_kmolh` and `recycle_curve_kmolh` at 720 min. *Checkpoint:* $+62.3$ % and 3.1 % to flood; 811.9 kmol/h against 817.5 kmol/h, within 0.5 kmol/h.

(b) Compare the Snowball gain tile at $t = 0$ and at 12 h, and say why it changed. *Checkpoint:* 4.61 and 5.84 % per percent.

(c) Load scenario 7, The column's purity loops are part of the recycle, in which the recycle purity loop is off, run it to 12 h and read the Recycle since $t = 0$ tile and `truth_xD_pct` at 720 min. *Checkpoint:* $+93.5$ % and 91.83 mol %, within 0.02 mol %.

7.8 App challenge · 8 min

Load scenario 3, CS2: the same production increase by the reactor holdup, and use the default seed. The reactor level setpoint steps from 50 to 57.2 % at 30 min.

(a) Run to 6 h. Read the Production and Recycle since $t = 0$ tiles, and find in the exported file the lowest `production_kmolh` and its time. *Checkpoint:* 221 kmol/h and -4.2 %; lowest 178.7 kmol/h at 55 min, within 0.2 kmol/h.

(b) At 6 h raise the Reactor effluent setpoint to 760 kmol/h and run on to 12 h. Read the same two tiles. *Checkpoint:* 225 kmol/h and $+6.4$ %.

(c) Use the slope of the snowball curve to explain why 60 kmol/h more effluent buys only about 5 kmol/h more product. *Checkpoint:* $\Delta F_0 \approx \Delta F/(1 + \mathrm{d}D^*/\mathrm{d}F_0) \approx 60/12$.

7.9 App challenge · 9 min

Load scenario 9, the challenge, and use the default seed. Each run lasts 8 h; at 300 min/s it plays in under 2 s.

(a) Run as shipped (CS1). Read the three So far tiles at 8 h and say which criteria fail. *Checkpoint:* 238.6 kmol/h worst half hour; product 79.57 mol % B; boilup at its limit 434.9 min. Purity and limits fail.

(b) Choose CS2, set the Reactor level setpoint to 70 %, press Reset and run the 8 h. Repeat at 70.5 %. Read the production tile each time. *Checkpoint:* 234.9 kmol/h, which fails, and 236.0 kmol/h, which passes.

(c) Using the CS2 balances of the theory exercise above, find the level that makes 235.2 kmol/h, the production the verdict requires, at steady state with 90 % activity, and compare it with (b). *Checkpoint:* 70.1 %, so 70.5 % is the first setting on the slider's 0.5 % grid.

7.10 Application area · 3 min

The column was designed at 67.6 % of its 1650 kmol/h flooding rate. The CS1 steady states at 216, 220 and 221.4 kmol/h of fresh feed need 1485, 1604 and 1649 kmol/h of boilup.

(a) Find the vapor margin left at steady state after case 1's 10 % increase. *Checkpoint:* 2.8 % of the flooding rate.

(b) A project engineer proposes to size a new column of this plant at 80 % of flooding, as for a plant without recycle. What does the snowball do to that choice? *Checkpoint:* The wall moves to 6.3 % more fresh feed.

7.11 Application area · 4 min

Case 6 doubles the reactor vessel to 400 m^3 at the same level. The design recycle falls from 500 to 106 kmol/h, the boilup from 1115 to 453 kmol/h, and the snowball gain at 200 kmol/h from 4.61 to 2.63.

(a) The app's case 6 briefing prices a vessel at its volume to the 0.6 power. By what factor does doubling the vessel raise its cost? *Checkpoint:* About 1.52.

(b) Name the equipment whose size the larger reactor saves, and say what the lower snowball gain gives the operator. *Checkpoint:* Column, reboiler, condenser and recycle pumps; a smaller recycle swing per percent of production or activity.

7.12 Review the agent · 5 min

An AI assistant is asked how to keep the product on specification in case 5, where the fresh feed rises 15 % on CS1. It replies: "The product purity loop is too slow for the larger load. Double the Product purity gain to 80 kmol/h per mol % and shorten its integral time to 7 min; the product will return to 98 mol % within an hour of the FLOODING alarm." Find the flaw, and say how you would check the claim and what you would do instead.

CHAPTER 8

A Verified Browser Port of the Tennessee Eastman Plant and Its Shutdown

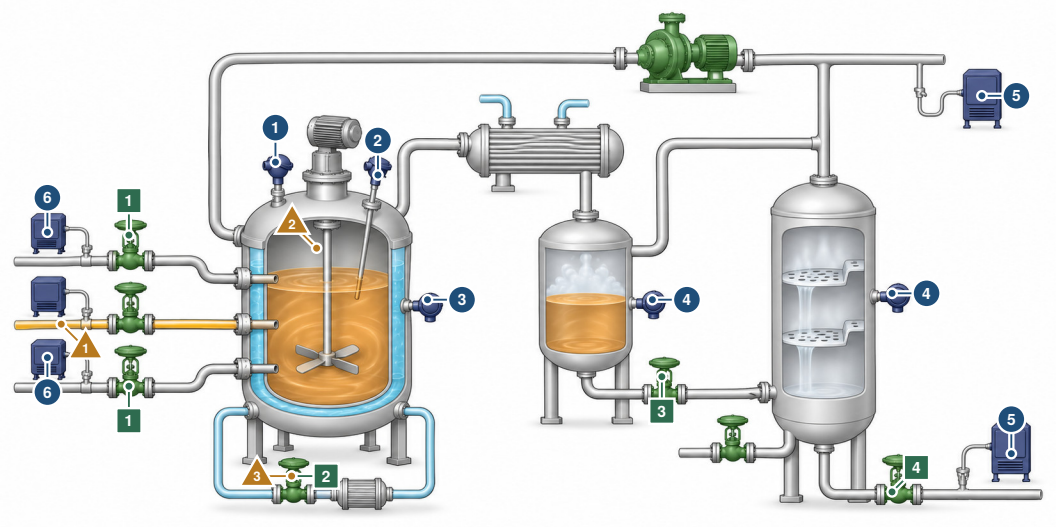

The controller moves 1 feed valves for A, D and E 2 reactor cooling water valve 3 separator liquid valve 4 product valve

The controller reads 1 reactor pressure 2 reactor temperature 3 reactor level 4 separator and stripper levels 5 purge and product analyzers 6 feed flow meters

What disturbs it 1 a feed lost or changing in makeup 2 slow drift of the reaction rates 3 a sticking cooling valve

Notice the gas that returns from the separator through the compressor to the reactor, which holds its pressure only while a loop opens the purge as the inert builds up. Illustration, not to scale.

IN PLAIN TERMS A Plant That Will Not Stand Still

The main idea. This is a model of a whole chemical plant. A reactor makes the product, a cooler and a tank split gas from liquid, and a column strips leftover gas out of the product. Gas that did not react is pumped back to the reactor, so each part feels the others.

The challenge. With every valve left fixed, gas builds up until the reactor pressure stops the plant within four hours. The controls must hold the gas and the liquid levels first, and only then the mix, the output and the cost.

Why it matters. The order in which the loops are closed decides whether such a plant keeps running.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on degrees of freedom in Section 8.1, the history in Section 8.2, the try-its and the fact on a lost feed in Section 8.6, and the key points at the end of Section 8.9
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the open challenge

The control question The published base case of the Tennessee Eastman plant is a steady state that is not stable, and with every valve fixed it trips on high reactor pressure within hours. Which loops must close first, what must a structure above them do for production, product mix and cost, and how far can a port of the public code be trusted?

Learning objectives

- Explain why a gas recycle with an inert and a purge on a fixed valve does not hold its own pressure.
- Order the layers of a plantwide structure, inventories before composition, production and cost, and support the order with the cases.
- Compute what the production index and its rate limit do to the flow setpoints and to the challenge's production score.
- Describe how the integrating pressure override trades production for margin, and tell a failure of a structure from a limit of the plant.
- Judge a reported trip time or cost as evidence about an implementation as well as about a plant.

Prerequisites Sections 1.4.2, 1.4.3 and 1.4.8 (degrees of freedom, inventories, levels as buffers), 1.6.2 and 1.6.4, with Section 1.6.3 for the predictive structure only. Chapter 7 treats recycle and the production handle on a simpler plant.

The app https://apmonitor.com/apps/tennessee-eastman/

See also Volume 4, Chapter 2 for a protective override that is idle in normal operation. Volume 1, Chapter 12 for step-test models of a multivariable plant and Volume 1, Chapter 9 for inventories judged by the time to a limit. Volume 13, Chapter 2 for an inert purged from a gas recycle under a pressure override, and Volume 13, Chapter 11 for optimization above the loops.



8.1 Introduction

Four gaseous feeds enter a stirred reactor, three reactions make the liquid products G and H, a condenser and a separator return the unreacted gas through a compressor, and a stripper recovers the light components from the product. At the base case of the public model this plant makes 22.95 m^3/h of product at a reactor pressure of 2705 kPa g and 120.4 °C, for 170.62 $/h by its own cost function. With all twelve valves frozen at that point and no disturbance, the reactor pressure still reads 2705 kPa g at 2 h, falls to 2690 at 3 h, and then climbs to the 3000 kPa g interlock at 3.59 h, which stops the plant.

That run corrects the belief that a plant left at a steady state stays there. A gas recycle with an inert, whose only exit is a purge through a fixed valve, does not hold its own pressure, and the liquid levels on pumped outlets are integrators. Section 1.4.3 treats such an inert as a component inventory, and the ammonia synthesis loop of Volume 13, Chapter 2 purges one in the same way. The plantwide procedures written since the benchmark appeared put these inventories first, before production rate and product quality, and they treat the class as general to plants with recycle, among them the reactor and separator flowsheets that the procedures were illustrated on [1–3]. A student who has tuned single loops has not met a plant in which the order of the loops decides whether it survives the afternoon.

> **DID YOU KNOW? Twelve Valves, Three Free Choices**
>
> The plant has 12 manipulated variables, yet a degree-of-freedom analysis finds only 8 degrees of freedom at steady state. At the economic optimum of the base operating mode 5 constraints are active, which leaves 3 unconstrained degrees of freedom for the control structure to settle by its choice of controlled variables [4].

The same plant corrects a second belief, that a port of a public benchmark reproduces the benchmark. The open-loop trip time above is set by the residue left when the constants are rounded. With the model's decimal constants read in double precision the trip comes at 3.59 h, with the single precision of a standard Fortran compiler at 4.06 h, and with an independent C translation's reading of five of them at 3.47 h, on the same one-second Euler step. However, the code itself was revised in

2015 because its results had stopped being repeatable in current simulation software [5], and much recent use of the plant is as a data set for fault diagnosis [6] rather than as a plant to hold.

This chapter carries the Tennessee Eastman model into a browser through a JavaScript port, states where the port came from and how closely it agrees with the code it was checked against, and measures what its control structures do with the plant. The model is Downs and Vogel's [7], and the decentralized structure follows Ricker's archived design [8]. The results are measurements, among them that this archived structure trips on the stripper level at 8.40 h after the A feed is lost, while Ricker's own recording of his 1996 strategy rides the same loss. The port is checked against three references, and its provenance, its licenses and the equations and settings from which it can be rebuilt are recorded here. Ten cases are provided, the tenth a two-day challenge redesigned so that its shipped structure fails and its suggested one passes at six seeds.

Section 8.2 traces the benchmark. Section 8.3 sets out the public model and the port, and Section 8.4 compares the port with the Python port, the C translation and the tabulated base case. Section 8.5 describes the inventory loops, the decentralized structure, the pressure override and the predictive controller. Section 8.6 measures the ten cases, Section 8.7 describes the real analyzers, valves and controllers, Section 8.8 states the limitations, and Section 8.9 concludes. The Exercises at the end of the chapter run the cases and the challenge.

8.2 A Challenge Problem From Kingsport

The problem came from industry. J. J. Downs and E. F. Vogel of the Process and Control Systems Engineering group of the Tennessee Eastman Company in Kingsport, Tennessee, wrote a simulation of an industrial process whose eight components are named only by the letters A to H, and distributed it as Fortran subroutines. The header of that code records a presentation at the AIChE Annual Meeting in Chicago on November 14, 1990, in a session on industrial challenge problems in process control, and a revision of April 4, 1991. The journal paper followed in 1993 [7]. The code defines 50 states, 12 manipulated variables (XMV), 41 measurements (XMEAS), of which 19 come from three sampled analyzers, 20 disturbances (IDV) and eight shutdown tests.

Control structures were proposed within two years. A base control layer was designed for the problem in 1994 [9], a plantwide structure and a control configuration design followed in 1995 [10, 11], and Ricker computed the plant's optimal steady operation [12] and, with Lee, applied nonlinear model predictive control to it [13]. His decentralized design of 1996 [8] is the one an implementer can rebuild without the paper, because he archived it as Simulink models with the code, his tabulated steady states and his closed-loop recordings at the University of Washington. The self-optimizing analysis of Larsson et al. recommends controlling the reactor pressure at its maximum and the reactor level at its minimum, closing the compressor recycle and stripper steam valves, and running the agitator at its maximum, and it finds that holding the inerts at a setpoint is a poor choice for this plant [14]. Ricker's tabulated optimized operating points all run at 2800 kPa g, against the base case's 2705.

> ⧖ HISTORY **1995 to 1996: Decentralized Loops Against Predictive Control**
>
> Ricker and Lee applied nonlinear model predictive control to the challenge plant and compared it with an earlier multiloop strategy. The decentralized loops performed adequately in many scenarios, but for the harder cases they would have needed multiple overrides, where the predictive controller might be preferred. In a second study Ricker designed decentralized control of his own. He controlled the variables that should optimally sit at their constraints, left out those whose best value varies a lot, and installed overrides [4]. He found little, if any, advantage in the nonlinear predictive controller: the decentralized strategy handled the constraints better, the very area where predictive control was reputed to excel. The problem, he wrote, had too many competing goals and special cases for a conventional predictive formulation [15].

The benchmark also shaped the general procedures. A nine-step heuristic procedure for plantwide control, ordered from energy management and production rate through inventories, component

balances and optimization, used the Eastman process as one of its three industrial examples [1], and control structure design for complete plants was later written as a systematic procedure [2]. One of the plant's authors has since described the industrial and academic views of plantwide control side by side [3]. The plant also serves as a source of process data, and comparisons of fault detection and diagnosis methods are run on it [6].

The code drifted as it moved between platforms. Ricker's archive carries a C translation of the Fortran produced by f2c and a MATLAB interface around it. In 2015 Bathelt, Ricker and Jelali reported that the original program flow no longer holds in current simulation software, so results were not repeatable, and gave a modified code with further measurements and a shorter run time [5]. Their revised code, in the same archive, gives the disturbances and the measurement noise separate random-number states and adds disturbances 21 to 28. The Fortran itself is redistributed in the Tennessee Eastman archive of R. D. Braatz's Large Scale Systems Research Laboratory at the University of Illinois, under that university's notice, and that archive is carried on GitHub in https://github.com/camaramm/tennessee-eastman-profBraatz, whose fork by J. Kitchin adds a pure-Python port.

8.3 The Public Model and Its Port

8.3.1 Provenance and Licenses

The application's model file, `teprob.js`, was translated routine by routine (TEFUNC, TEINIT and TESUB1 to TESUB8) from `tep/python_backend.py` in the fork https://github.com/jkitchin/tennessee-eastman-profbraatz at commit `9a6c8e5`, whose `pyproject.toml` and README declare the BSD-3-Clause license. Every block was then checked against `teprob.f` as the same repository distributes it from the Braatz archive, under the University of Illinois notice of 1998 to 2002. Where the Python and the Fortran differ, the Fortran was followed. Ricker's archive supplied the base-case table against which the initial state is checked, the decentralized structure in `MultiLoop_mode1.mdl`, the C translation `temexd.c` used as an independent check, and four of his closed-loop recordings. From the 2015 revision only the split random-number streams were taken, because the revision changes no plant equation, its solver corrections do nothing at a fixed Euler step, and its added disturbances lie outside this application. Both license notices are reproduced in full in the application's Model dialog, as their conditions ask of a redistribution in compiled form, and at the head of `teprob.js`. No Fortran, Python or WebAssembly compiled from either is shipped.

8.3.2 The Plant

Feeds of A (stream 1), D (2) and E (3) enter a header with the compressor discharge (8) and the stripper overhead (5), and the mixture (6) enters the reactor, as shown in Figure 8.1. The reactor effluent vapor (7) is cooled in a condenser and separated, the vapor returns through the compressor, whose recycle valve (XMV 5) spills discharge back to suction, and a purge (9) leaves from the compressor suction. The separator liquid (10) goes to the stripper, where the A and C feed (4) strips the light components back to the header and the product (11) leaves the foot, heated by steam. Table 8.1 classifies the variables. The code's internal units are pound-moles, cubic feet, millimeters of mercury and degrees Celsius, and the measurements leave it in the published units.

8.3.3 Balances

Each of the four vapor or liquid holdups carries eight component holdups and an internal energy. With $x_{k,i}$ the mole fraction of component i in stream k and F_k its molar flow, the component balances

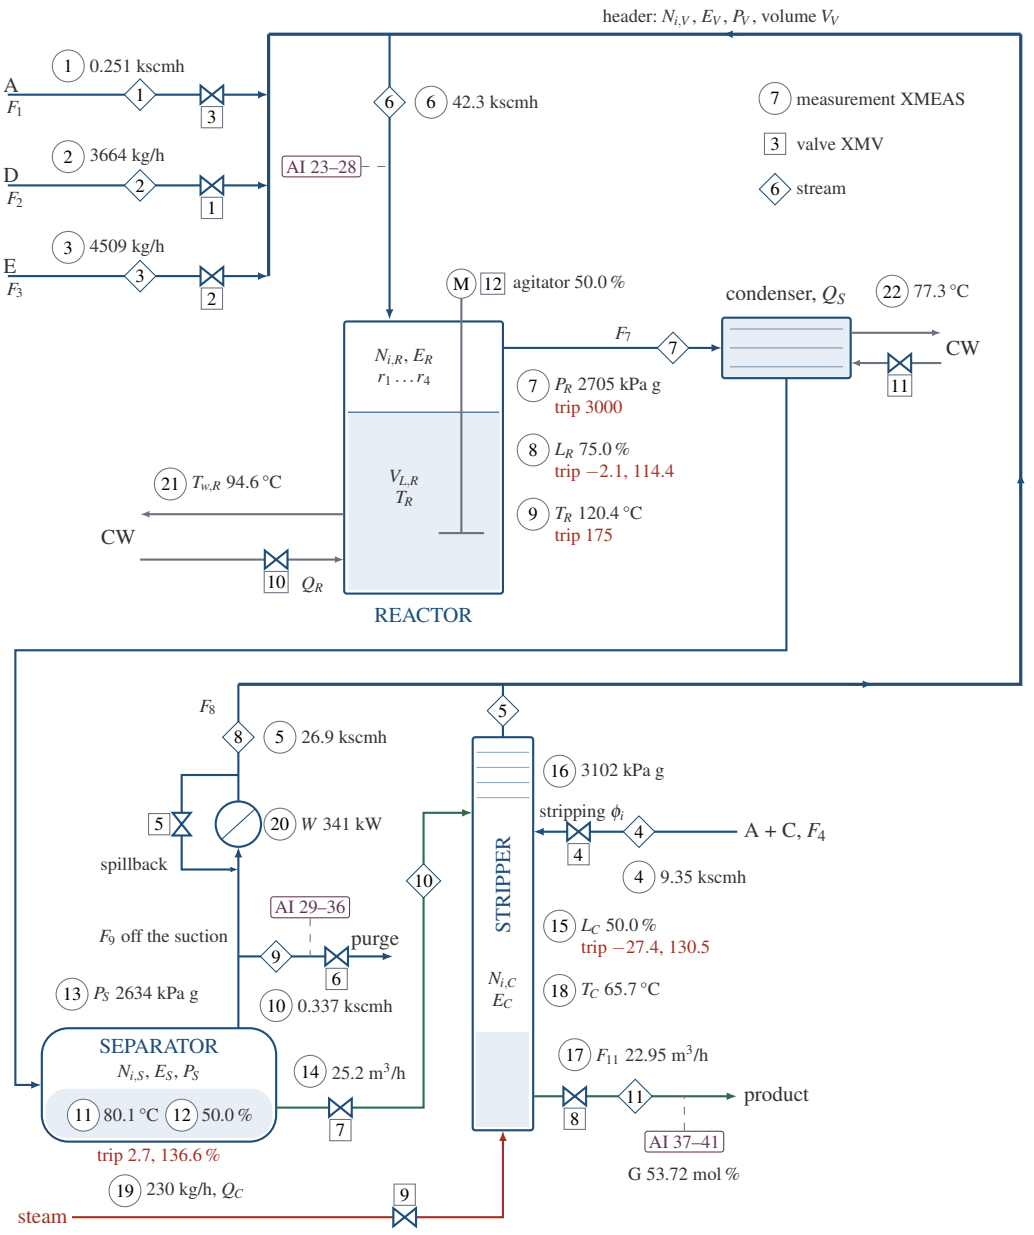

Figure 8.1 The Tennessee Eastman plant as modeled, after the application's Plant view of case 1 at $t = 0$, the published base case, drawn in two bands joined by the effluent from the condenser and by the header, which returns the compressor discharge (8) and the stripper overhead (5) up the right edge. Streams 1 to 11 carry their published numbers in diamonds, each circle is a continuous measurement (XMEAS 1 to 22) with its base-case value, each square a valve (XMV 1 to 12), and the three analyzer boxes give the XMEAS ranges they report. The lines marked trip give the interlock limits on the watched variables, the levels in their transmitters' percent. The holdups $N_{i,u}$ and energies E_u of the reactor (R), separator (S), stripper (C) and header (V) are the states of Equations (8.1) and (8.2). Process model only, not to scale.

Table 8.1 Variables by Role, With Base-Case Values

Variable	Role	Base case, range
F_{11}, XMEAS 17	CV, production	22.95 m^3/h, SP 15–32
x_G, XMEAS 40	CV, G in product (analyzer)	53.72 mol %, SP 5–95
P_R, XMEAS 7	CV, reactor pressure	2705 kPa g, trip 3000
L_R, L_S, L_C	CV, three levels (XMEAS 8, 12, 15)	75.0, 50.0, 50.0 %
T_R, XMEAS 9	CV, reactor temperature	120.4 °C, trip 175
y_A, y_{AC}	CV, feed analyzer (XMEAS 23, 25)	54.96 %, 58.57 mol %
$v_1 \dots v_{12}$	MV, XMV 1–12	Figure 8.1, 0–100 %
IDV 1–20	DV, steps, random walks, sticking	off, on or off
$N_{i,u}$, E_u	states, $u \in \{R,S,C,V\}$	36
$T_{w,R}$, $T_{w,S}$	states, cooling-water outlets	94.6, 77.3 °C
v_j	states, valve positions	12
F_p	controller, production index	100 %
J	output, operating cost	170.62 $/h

are Equation (8.1).

$$\begin{aligned}
\frac{dN_{i,V}}{dt} &= \sum_{k\in\{1,2,3,5,8\}} F_k x_{k,i} - F_6 x_{6,i} \\
\frac{dN_{i,R}}{dt} &= F_6 x_{6,i} - F_7 x_{7,i} + \sum_{j=1}^{4} \nu_{ij} r_j \\
\frac{dN_{i,S}}{dt} &= F_7 x_{7,i} - F_8 x_{8,i} - F_9 x_{9,i} - F_{10} x_{10,i} \\
\frac{dN_{i,C}}{dt} &= (1-\phi_i)(F_4 x_{4,i} + F_{10} x_{10,i}) - F_{11} x_{11,i}
\end{aligned} \tag{8.1}$$

where streams 6, 7, 8 and 9 carry the vapor compositions of the header, the reactor and the separator, 10 and 11 the liquid compositions of the separator and the stripper, and the stripper overhead is $F_5 x_{5,i} = \phi_i(F_4 x_{4,i} + F_{10} x_{10,i})$. The reactor and separator store A, B and C as vapor and D to H as liquid. The stripping factor is $\phi_i = \sigma_i \psi / (1 + \sigma_i \psi)$ with $\psi = (F_4/F_{10})\, g(T_C)$ and g a piecewise function of the stripper temperature, for D to H only, and A, B and C go overhead in fixed fractions. The energies follow the same streams, as shown in Equation (8.2).

$$\begin{aligned}
\frac{dE_V}{dt} &= \sum_{k\in\{1,2,3,5,8\}} H_k F_k - H_6 F_6 \\
\frac{dE_R}{dt} &= H_6 F_6 - H_7 F_7 + \Delta H_1 r_1 + \Delta H_2 r_2 + Q_R \\
\frac{dE_S}{dt} &= H_7 F_7 - \sum_{k\in\{8,9,10\}} H_k F_k + Q_S \\
\frac{dE_C}{dt} &= H_4 F_4 + H_{10} F_{10} - H_5 F_5 - H_{11} F_{11} + Q_C
\end{aligned} \tag{8.2}$$

where H_k is the molar enthalpy of stream k from cubic heat-capacity polynomials with a vaporization term for vapor, H_8 includes the compressor's work, Q_R and Q_S are the heats exchanged with the two cooling-water streams (negative in operation) and Q_C is the steam's. Only reactions 1 and 2 carry a heat of reaction. Each unit's temperature is found from its energy by Newton's method on the enthalpy polynomial, started from the previous temperature.

8.3.4 Pressure, Equilibrium and Reaction

The noncondensables obey the ideal gas law in each vessel's vapor space and the condensables Raoult's law with Antoine vapor pressures, as shown in Equation (8.3).

$$\begin{aligned} P_u &= \sum_{i\in\{A,B,C\}} \frac{N_{i,u} R T_u}{V_u - V_{L,u}} + \sum_{i\in\{D..H\}} x_{i,u}\, \mathrm{e}^{A_i + B_i/(T_u + C_i)} \\ P_V &= \frac{\left(\sum_i N_{i,V}\right) R T_V}{V_V}, \qquad V_{L,u} = \sum_{i\in\{D..H\}} \frac{N_{i,u}}{\rho_i(T_u)} \end{aligned} \tag{8.3}$$

where u is the reactor or the separator, T_u is in kelvin in the gas term, $R = 998.9\,\mathrm{mmHg\,ft^3/(lbmol\,K)}$ in the code's units, and the liquid molar densities ρ_i are quadratic in temperature. The four reactions run in the reactor's vapor volume, as shown in Equation (8.4).

$$\begin{aligned} r_1 &= \delta_1\,(V_R - V_{L,R})\,\mathrm{e}^{a_1 - E_1/(R_c T_R)}\, p_A^{n_A}\, p_C^{n_C}\, p_D \\ r_2 &= \delta_2\,(V_R - V_{L,R})\,\mathrm{e}^{a_2 - E_2/(R_c T_R)}\, p_A^{n_A}\, p_C^{n_C}\, p_E \\ r_3 &= (V_R - V_{L,R})\,\mathrm{e}^{a_3 - E_3/(R_c T_R)}\, p_A\, p_E \\ r_4 &= 0.767488334\,(V_R - V_{L,R})\,\mathrm{e}^{a_3 - E_3/(R_c T_R)}\, p_A\, p_D \end{aligned} \tag{8.4}$$

for A + C + D $\rightarrow$ G, A + C + E $\rightarrow$ H, A + E $\rightarrow$ F and 3D $\rightarrow$ 2F, with p_i the reactor partial pressures, $R_c = 1.987\,\mathrm{cal/(mol\,K)}$, and δ_1, δ_2 the kinetic factors that IDV 13 drifts. The stoichiometry gives $\nu_A = (-1,-1,-1,0)$, $\nu_C = (-1,-1,0,0)$, $\nu_D = (-1,0,0,-1.5)$, $\nu_E = (0,-1,-1,0)$, $\nu_F = (0,0,1,1)$, $\nu_G = (1,0,0,0)$ and $\nu_H = (0,1,0,0)$, and B is inert.

8.3.5 Flows, Heat Transfer and Valves

Pressure differences drive the gas between the vessels, and the compressor follows a curve in its pressure ratio less what its recycle valve returns, as shown in Equation (8.5).

$$\begin{aligned} M_6 F_6 &= 1937.6\sqrt{P_V - P_R} \\ M_7 F_7 &= 4574.21\sqrt{P_R - P_S}\,(1 - 0.25\,\delta_{20}) \\ M_8 F_8 &= F_{max}\left[1 + (1-\rho^3)/1.197\right] \\ &\quad - 53.349\,v_5\sqrt{P_V - P_S} \\ M_9 F_9 &= 0.151169\,v_6\sqrt{P_S - 760} \\ F_k &= v_{j(k)}\,F_{k,max}/100 \quad \text{(feeds, liquid draws)} \end{aligned} \tag{8.5}$$

where M_k is the stream's mean molar mass, pressures are absolute in mmHg, each square root is of a difference clipped at zero, $\rho = \min(1.3, \max(1, P_V/P_S))$ is the compressor's pressure ratio, δ_{20} is the flow walk of IDV 20, the A feed is cut to zero by IDV 6 and the A and C feed to 80 % by IDV 7. The heat removed from the reactor and the condenser and the steam's heat are

$$\begin{aligned} Q_R &= UA_R(V_{L,R}, a)\,(T_{w,R} - T_R) \\ Q_S &= UA_S(F_7)\,(T_{w,S} - T_R) \\ Q_C &= UA_C\,v_9\,\max(0,\, 100 - T_C) \end{aligned} \tag{8.6}$$

where UA_R rises linearly with the reactor liquid volume from 78 to 390 $\mathrm{ft^3}$ and is quadratic in the agitator speed $a = (v_{12} + 150)/100$, and UA_S saturates in the effluent flow. Each cooling-water outlet temperature is a state that balances the water's flow, set by v_{10} or v_{11}, against the heat it takes, and IDV 4, 5, 11, 12, 16, 17 and 18 move the inlet temperatures or the three heat-transfer coefficients. Every valve moves toward its command through a first-order lag, $\tau_j\,\mathrm{d}v_j/\mathrm{d}t = c_j - v_j$, with τ_j from 5 s to 9 s and 120 s for the steam valve, and under IDV 14, 15 and 19 a valve keeps its last command until the new one differs by more than 2 %.

Table 8.2 Model Constants. H: Hidden, Shown Read-Only in the Model Dialog. A: Adjustable in the Panel

Symbol	Meaning	Value	Mark
V_R, V_S, V_V	reactor, separator, header volumes (ft^3)	1300, 3500, 5000	H
V_C	stripper level span (ft^3)	156.5	H
a_1, a_2, a_3	pre-exponentials (ln)	31.5859536, 3.00094014, 53.4060443	H
E_1, E_2, E_3	activation energies (cal/mol)	40000, 20000, 60000	H
n_A, n_C	orders in p_A, p_C	1.1544, 0.3735	H
ΔH_1, ΔH_2	heats of reaction (code units)	0.06899381054, 0.05	H
F_{max}	compressor maximum flow (lb/h)	280275	H
ρ_{max}	compressor maximum ratio	1.3	H
$\tau_1 \ldots \tau_{12}$	valve lags (s)	8, 8, 6, 9, 7, 5, 5, 5, 120, 5, 5, 5	H
T_a	analyzer periods and dead times (h)	0.1, 0.1, 0.25	H
limits	P_R, T_R, liquid volumes	see text	H
v_j	valve stations (%)	base case	A
IDV	twenty disturbances	off	A
Δt	Euler step (s)	1	fixed

8.3.6 Measurements, Disturbances, Shutdown and Cost

The 22 continuous measurements carry additive noise of the published standard deviations σ_k, each draw the sum of twelve uniform deviates less six, times σ_k. The reactor feed and purge analyzers sample every 0.1 h and report what they sampled one period earlier, and the product analyzer does the same at 0.25 h, as shown in Equation (8.7).

$$\begin{aligned} y_k(t) &= h_k(\mathbf{x}(t)) + \sigma_k\Big(\textstyle\sum_{m=1}^{12} u_m - 6\Big), \quad k \le 22 \\ y_k(t) &= h_k\big(\mathbf{x}(t_s - T_a)\big) + \text{noise}, \quad t_s \le t < t_s + T_a \end{aligned} \tag{8.7}$$

where t_s is the last sample time and T_a the analyzer's period. The random disturbances are cubic splines whose knots are drawn at random intervals, and the step disturbances shift a feed composition, temperature or availability. The plant shuts itself down, freezing every state as the code does, when P_R exceeds 3000 kPa g, T_R exceeds 175 °C, or a liquid volume leaves its band: reactor 2 to 24 m^3, separator 1 to 12 m^3, stripper 1 to 8 m^3. The operating cost is the published function as Ricker's archived Simulink model codes it, as shown in Equation (8.8).

$$\begin{aligned} J = {} & 0.0536\,W + 0.0318\,S + 0.44791\,F_9 \textstyle\sum_i c_i\, y_{9,i} \\ & + \tfrac{211.3}{46.534}\, v_8\,(0.2206\,x_D + 0.1456\,x_E + 0.1789\,x_F) \end{aligned} \tag{8.8}$$

where W is the compressor work in kW (XMEAS 20), S the steam in kg/h (XMEAS 19), F_9 the purge in kscmh with its analyzed composition $y_{9,i}$ priced at c_i (2.209, 0, 6.177, 22.06, 14.56, 17.89, 30.44 and 22.94 for A to H), and the last term the D, E and F lost in the product, with the product rate taken from the stripper valve as 211.3 kmol/h at 46.534 % open. Ricker's Simulink block rounds that ratio to 4.541, which changes the cost of the tabulated base case by 0.0016 $/h.

8.3.7 What the Port Changes, and Its Numerics

Three deliberate changes lie outside the plant's arithmetic. The random numbers come from two seeded generators of the framework (Section 1.6.1), one for the disturbance walks and one for the measurement noise, as in the 2015 revision, instead of the single linear congruential generator of the code, so a run replays bit for bit after Reset while its noise realization differs from the original's. The analyzers fire on a count of steps (360 and 900 one-second steps) instead of comparing a floating-point time, and the state lives in one object per plant. The constants are the decimal values typed in the code, read in double precision as the Python port reads them, which is the subject of Section 8.4. The integrator is the model's own explicit Euler step of 1 s, as the driver distributed with the code uses, because TEFUNC advances the disturbance walks and the analyzers when it is called and therefore

cannot be placed inside a multistage integrator. A 48 h run of the challenge, controllers and logging included, took about 1.3 s headless in Node on the development host. In the application, at the default speed of 0.25 h per second the open-loop shutdown at 3.59 h arrives about 14 s after Run, and at the top speed of 4 h per second the two-day challenge plays in about 12 s. The constants with their marks are in Table 8.2.

8.3.8 Assumptions

The plant equations, constants, noise levels, analyzer periods and limits are those of the public code and are the benchmark's teaching values, not data of any operating plant, whose relief and interlock settings would come from its vessel design code and its safety instrumented system study under IEC 61511. The phases are in equilibrium at each vessel's temperature, the stripper separates by fixed factors rather than by stages, and the compressor and the flows between vessels are algebraic, so pressure waves and transport delays are missing, which makes fast transients the least trustworthy part of the model. The noise statistics are the published ones while the realizations are not. Disturbances 16 to 20, which the code implements although its own comments list them as unknown, are kept, and the revision's disturbances 21 to 28 and its extra measurements are not ported.

8.4 How Closely the Port Follows the Code

8.4.1 The Base Case

At $t = 0$ the 41 measurements agree with the base case tabulated in Ricker's archive to within 0.96 of half the last printed digit (worst at XMEAS 1), and the 12 valve positions to within 0.98 (XMV 11). Every derivative is at most 9.31×10^{-4} per hour with the constants in double precision, inside the 10^{-3} criterion that Ricker's archive states for a steady state, and 2.54×10^{-3} with the single-precision reading. This is why the port reads the constants in double precision.

8.4.2 Against the Python Port

For the point-by-point comparison the port runs in a validation mode that reproduces the Python port's choices: the original generator in one shared stream with exact integer arithmetic, its Newton starting temperatures, a floating-point analyzer clock and its clipping of the valve commands. The Python port was run once, in an isolated environment, with its trajectories recorded every minute, and every comparison below reads those records. Under IDV 1, 2, 4, 6, 8 and 13 in open loop, the analyzer readings differ by at most 1.76×10^{-11} mol % and the states by 9.57×10^{-12} relative, and with four PI loops closed for 6 h under nine disturbance sets the states differ by at most 1.59×10^{-11}, as shown in Figure 8.2b. The undisturbed open-loop run reaches 2.05×10^{-8}, because differences in the last bit grow along the unstable mode. Every shutdown comes at the same step, at 3.5856, 0.4783, 1.5992, 1.0564, 0.4422, 1.8878 and 2.0261 h in open loop (undisturbed, then IDV 1, 2, 4, 6, 8 and 13) and 2.8894 h closed under IDV 1.

8.4.3 Against an Independent C Translation

Ricker's f2c translation shares no code with the Python port. With the constants read as f2c reads them (the three ratios E_j/R_c folded in double precision and then rounded to single, the two exponents kept in double, the rest in single), the port agrees with it bit for bit at $t = 0$ in all 50 derivatives and 41 measurements, and to 3.98×10^{-10} relative over 6 h of open loop, with the shutdown at the same step, second 12485. Under IDV 1, 2, 4, 6 and 7, which are deterministic, the agreement over runs of 1 to 2 h is 3.95×10^{-12} or better. The shipped double-precision reading differs from the C code by 1.71×10^{-6} in the measurements at $t = 0$, which is the size of the rounding, not of an error.

8.4.4 An Open-Loop Trip Set by Rounding

The undisturbed open-loop shutdown is the plant's instability acting on that residue, so its time belongs to the reading of the constants, as shown in Figure 8.2a. On the same 1 s Euler step the trip comes at 3.47 h with the C translation's reading, 3.59 h in double precision and 4.06 h in single

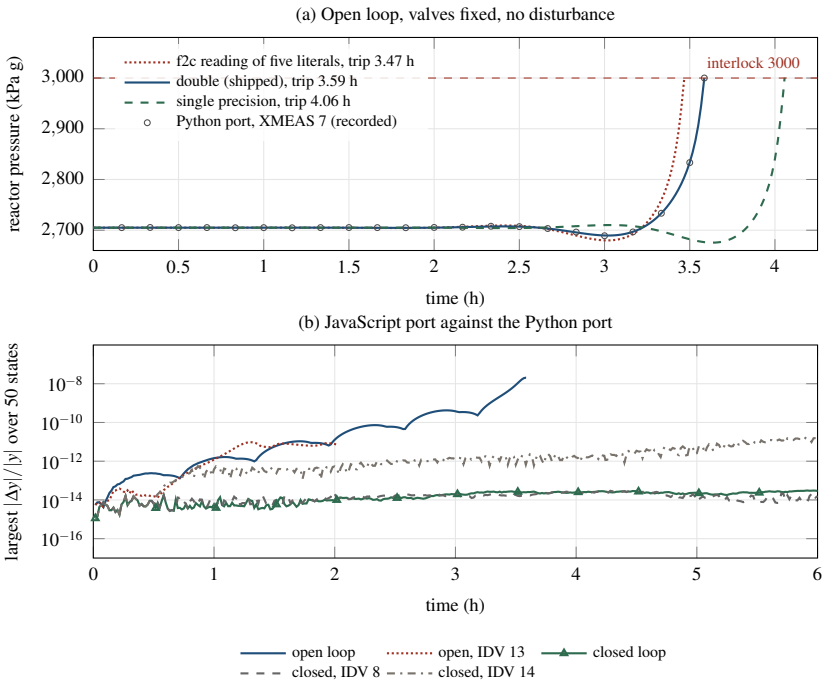

Figure 8.2 The port against its references. (a) Open-loop reactor pressure for three readings of the same decimal constants, the f2c reading dotted, the double-precision reading solid and the single-precision reading dashed, with the Python port's recorded measurement every 10 min as open circles and the interlock as a thin dashed line. (b) The largest relative state difference from the Python port, minute by minute, for two open-loop and three closed-loop runs: open loop solid, open loop under IDV 13 dotted, closed loop solid with triangles, and closed loop under IDV 8 dashed and under IDV 14 dash-dotted.

precision, always on high reactor pressure. The integration matters far less: half and a quarter of the step give 3.5863 and 3.5867 h, and the classical fourth-order method at 1 s gives 3.5872 h. The README of Ricker's Simulink example reports the same shutdown after about 2 h of simulated time. A trip time quoted for this plant therefore identifies an implementation as much as a dynamic.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.4.5 The Base-Case Cost

The port's cost at $t = 0$ is 170.62 \$/h, and the same function evaluated on Ricker's tabulated base case gives 170.83 \$/h. By term the table exceeds the port by -0.023 for power, -0.010 for steam, -0.045 for the purge and $+0.285$ \$/h for the product losses. Almost all of the gap is the table's two-decimal product composition (D, E and F printed as 0.02, 0.84 and 0.10 mol % against the port's 0.0179, 0.8357 and 0.0986), which alone accounts for 0.28 \$/h. Within its printed digits the table's cost can lie anywhere from 168.10 to 173.57 \$/h, and the port's value is inside. The 1993 paper was not available when the port was checked, so no printed base-case cost is quoted from it.

8.4.6 Python and Fortran Differences

Eight differences between the Python port and the Fortran were found while translating, none in the balances, rates, equilibrium, flows, cost or shutdown tests. They concern the generator's arithmetic, the precision of untyped literals, the Newton starting temperatures, where the valve commands are clipped, evaluation after a shutdown, guarded divisions and the driver's single-precision step. The

Table 8.3 The Seventeen PI Loops of the Archived Structure. K_c Acts on SP − PV, So a Negative Gain Raises the Output With the Measurement

Measured	Output	K_c	τ_I (min)	Limits
production, XMEAS 17	trim on F_p (%)	3.2	120	±30
P_R, XMEAS 7	purge ratio r_5	-10^{-4}	20	0, 100
L_R, XMEAS 8	separator T setpoint	0.8	60	0, 120
T_R, XMEAS 9	XMV 10	−8	7.5	0, 100
T_S, XMEAS 11	XMV 11	−4	15	0, 100
L_S, XMEAS 12	liquid ratio r_6	-10^{-3}	200	0, 100
L_C, XMEAS 15	product ratio r_7	-2×10^{-4}	200	0, 100
x_G, XMEAS 40	D/E shift E_{adj}	−0.4	100	none
y_A, XMEAS 23, 25	A trim u_A	2×10^{-4}	60	±1
y_{AC}, XMEAS 23, 25	A + C trim u_{AC}	3×10^{-4}	120	±1
seven flows	XMV 3, 1, 2, 4, 6, 7, 8	see text	0.001	0, 100

port follows the Fortran in each except the literals and the step, and its validation mode follows the Python.

8.5 Control Structures as Implemented

The controllers execute at every 1 s step and read only the 41 measurements, including the analyzers' late reports. The header's Manual gives all twelve valves to their stations. Auto runs the structure chosen in the panel, and a valve the structure does not drive stays on its station. Loops outside the running structure track the plant, so a change of structure or a transfer to Auto moves no valve. Every PI loop is the velocity-form PID of Section 1.6.2, with the tracking time of its back-calculation set to $T_t = \max(\tau_I, \Delta t)$ because Ricker's flow loops have integral times of 0.06 s, below the step. His discrete PI blocks are carried with his gains and integral times, his sign convention mapped to the loop's action, as listed in Table 8.3.

8.5.1 Inventory Loops

The inventory structure closes the reactor pressure on the purge, the reactor level on the separator temperature setpoint, the separator temperature on the condenser cooling water, the separator and stripper levels on their liquid draws, and the reactor temperature on its cooling water, with the four feeds, the recycle, the steam and the agitator on their stations. The reactor temperature loop belongs in this layer because without it the undisturbed plant trips on high pressure at 8.12 h.

8.5.2 The Decentralized Structure

The decentralized structure follows Ricker's archived Simulink model `MultiLoop_mode1.mdl`, run at the base-case setpoints. A production index sets every flow setpoint as a ratio times the index, as shown in Equation (8.9). The index is the plant's throughput manipulator in the sense of Section 1.4.3, and it moves every flow at once.

$$\begin{aligned} F_p &= \tfrac{100}{22.89}\,F_{SP} + \mathrm{PI}_F(F_{SP} - F_{11}) \\ F_{k,SP} &= r_k F_p, \qquad k = 1\ldots7 \\ r_2 &= p_D(x_{G,SP}) - 32\,E_{adj}/F_p + b_D \\ r_3 &= p_E(x_{G,SP}) + 46\,E_{adj}/F_p + b_E \\ r_1 &= \bar{r}_1 + u_A, \qquad r_4 = \bar{r}_4 - u_A + u_{AC} \end{aligned} \tag{8.9}$$

where F_{SP} is the production setpoint after a rate limit of 30 % of the base rate per day, p_D and p_E are Ricker's quadratic feedforward polynomials in the G setpoint (itself rate-limited to 50 mol % per day), E_{adj} is the output of the G loop, u_A and u_{AC} are the outputs of the two feed analyzer trims, and r_5 to r_7 are the outputs of the pressure, separator level and stripper level loops. The biases b_D,

b_E, $\bar{r}_1$ and $\bar{r}_4$ are set when a loop enters service so that each feed setpoint equals the feed the meter reads. The flow loops have gains of 0.01, 1.6×10^{-6}, 1.8×10^{-6}, 0.003, 0.01, 4×10^{-4} and 4×10^{-4} percent of valve per unit of their flows (kscmh for gases, kg/h for D and E, m^3/h for liquids).

In this archived model the production index, with the production loop's trim, scales all seven flow loops, the product draw among them. Ricker's archived recordings of his 1996 strategy show a different formulation. In the run of IDV 6 at 50 h his product setpoint is 20.877, equal to the product ratio times the index less the production trim, and not 21.143, the ratio times the full index, so his trim moved only the four feeds, and the same record shows his pressure override taking 8.71 points off the index. The application keeps the archived formulation because it is the one his code states. A one-element change toward the 1996 formulation, tried during the review of the application, neither reproduced his IDV 6 recording nor kept the suggested challenge running (it tripped at 5.88 h).

8.5.3 A Pressure Override on the Production Index

The archived model carries no override, so the application adds one of its own design. Like the surge override of Volume 4, Chapter 2, it is idle in normal operation and acts only near a limit. The reactor pressure passes through a first-order filter of 0.02 h, and while the filtered pressure P_f exceeds the threshold P_{ovr}, each step takes a cut from the index and writes the index applied back into the production loop's output, as shown in Equation (8.10).

$$\begin{aligned} P_{f,k} &= P_{f,k-1} + (1-\mathrm{e}^{-\Delta t/0.02})(P_{R,k}-P_{f,k-1}) \\ \delta_k &= K_{ovr}\,(P_{f,k}-P_{ovr})\,\Delta t_s, \qquad F_p \leftarrow F_p - \delta_k \\ u_{F,k} &\leftarrow \max(1,F_p) - \tfrac{100}{22.89}\,F_{SP} \end{aligned} \tag{8.10}$$

where Δt_s is the control period in seconds (1 at the shipped step), $K_{ovr}=2$ percent of F_p per kPa per second and $P_{ovr}=2950\,\mathrm{kPa}\,g$. Because each cut stays in the production loop's output, the override integrates until that output reaches its lower limit of -30 points, and beyond it the cut is in proportion to the present excess. With the production loop frozen at the override's first step in case 8 at 17 m^3/h, a filtered excess of at most 1.11 kPa took F_p from 83.8 to 42.1 % in 3 min. Over the last 4 h of that case the production loop sat at its limit 98 % of the time and the override acted 58 % of the time, with the filtered pressure on average 9.85 kPa above its threshold while it did. The measured pressure noise is 0.298 kPa rms against the code's 0.3, and the filter is there because an unfiltered override, tried while the application was built, cut production on that noise alone. In all ten cases as shipped (case 1 in Auto, the challenge on its suggested settings) at three seeds, the override acts only in case 8.

8.5.4 Model Predictive Control Above the Loops

In the predictive structure the production and mix loops stand down and a receding-horizon controller sets their two handles, the production target u_1 (m^3/h) that sets $F_p=(100/22.89)u_1$ and the G target u_2 (mol %) that feeds the polynomials of Equation (8.9). The program is the one of Section 1.6.3 with this plant's model and weights. Its model, identified by steps on this plant as Volume 1, Chapter 12 teaches, is two first-order responses, as shown in Equation (8.11).

$$\begin{aligned} 0.1\,\dot{y}_1 &= u_1 - y_1, \qquad 1.4\,\dot{z} = u_2 - z \\ y_2(t) &= z(t-0.65) \\ \min\ & \textstyle\sum_{k=1}^{80}\left[w_F(y_{1,k}-r_1)^2 + w_G(y_{2,k}-r_2)^2\right] \\ & \quad + \textstyle\sum_j\left[w_{\Delta u}\Delta u_{1,j}^2 + 0.1\,w_{\Delta u}\Delta u_{2,j}^2\right] \end{aligned} \tag{8.11}$$

in hours, discretized for a zero-order hold every 0.05 h over 80 steps (4 h) in 12 blocks, with a terminal weight five times the stage weights, $w_F=10$, $w_G=1$, $w_{\Delta u}=5$, bounds of 15 to 32 m^3/h and 5 to 95 mol %, and moves limited by the two ramp limits. The cross effects are left to output biases corrected toward the measurements with a gain of 0.3 at every decision. Each program is handed

Table 8.4 The Ten Cases

#	Case and events	Teaching point
1	Manual, every valve fixed, no disturbance	the open-loop shutdown
2	Inventory loops only, IDV 1 at 2 h	inventories first
3	Production −15 % over 2 h from 2 h	one production handle
4	IDV 1 at 1 h	analyzers that report late
5	IDV 8 throughout	the cost of noise and of pressure
6	IDV 13 at 1 h, 48 h	slow drift
7	IDV 14 at 1 h	a sticking valve
8	IDV 6 at 1 h	the override, a structure's limit
9	G setpoint to 11.7 mol % at 1 h	mix change, 50/50 to 10/90
10	Challenge: IDV 8, IDV 1 at 2, production +10 % at 24 h, inventory loops at 2705 kPa g as shipped	the whole problem

to the dual active-set QP solver of Goldfarb and Idnani [16] that the series shares, the first move is applied as industrial practice does [17], and a plan the solver rejects, or one with a non-finite value or outside its bounds, returns both handles to the PI loops without a bump.

8.5.5 Three Models and the Truth Boundary

The plant is the ported code. The identified model is the pair of responses in Equation (8.11), and the controllers carry no other model. The scorecard, the challenge verdict and the comma-separated export columns prefixed `truth_` use the noise-free, delay-free values, which no controller reads. The export also carries the 41 measurements, the twelve valve commands and positions and the controller states.

8.6 Ten Cases, Measured

The ten cases are listed in Table 8.4. Each starts from the base case on the decentralized structure with its override unless its row says otherwise. Every number below was measured on the engine headless at the default seed, with settings entered before Run, which the engine treats as starting conditions (Section 1.6.1). Each case logs a bit-identical run whether it is opened first, after another case, or again through Reset.

8.6.1 Inventories Before Everything

Case 1 is the run of the Introduction. The plant holds its published state for two hours with every valve fixed, then the pressure reaches the interlock at 3.59 h with the reactor at 96.6 °C. In Auto the same plant holds the pressure between 2702 and 2708 kPa g for 48 h. Case 2 runs the inventory loops alone and steps the A and C feed's composition at 2 h. The analyzer's A share of A plus C in the reactor feed falls from 55.0 to 37.0 %, the unreacted C accumulates in the recycle, and the pressure loop opens the purge as far as it helps. The pressure passes 2900 kPa g at 4.84 h and the interlock at 5.84 h. With the decentralized structure the A share never falls below 45.4 %, the pressure peaks at 2751 kPa g, and the plant runs on. The inventory loops make the plant stable, and only the composition loops above them make it survive a feed change.

Table 8.5 Measured Outcomes at the Default Seed

#	Setting	Outcome
1	as shipped (Manual)	trip, pressure, 3.59 h
	Auto	2702–2708 kPa g, 48 h
2	inventory loops	trip, pressure, 5.84 h
	decentralized	peak 2751 kPa g, no trip
3	as shipped	19.29 m^3/h at 12 h, stripper 28.7–81.1 %
	stripper level gain 0.08	stripper 42.3–63.6 %, production off by 0.55
4	as shipped	A share low 45.29 %, peak 2752 kPa g
	trims at zero gain	A share low 22.94 %, peak 2981 kPa g
5	at 2705 kPa g, 24 h	mean 176.1 $/h (170.7 undisturbed)
	at 2800 kPa g, 24 h	mean 146.3 $/h, peak 2827 kPa g
6	decentralized, 48 h	2566–2854 kPa g, mean 209.1 $/h
	inventory loops, 48 h	2543–2847 kPa g, mean 212.3 $/h
7	$K_c = 8$ (Ricker's)	±0.31 °C, period 4.8 min
	$K_c = 4$	±0.57 °C, period 8.1 min
	$K_c = 16$	±0.23 °C, period 3.2 min
8	as shipped	trip, stripper level low, 8.40 h
	17 m^3/h, override	no trip in 24 h, 5.88 m^3/h last 4 h
	17 m^3/h, no override	trip, pressure, 13.34 h
	inventory loops	no trip, 21.94 m^3/h, 408 $/h
9	ramp 2.08 mol %/h	within 2 points at 19.8 h, low 7.6 %
	ramp 10 mol %/h	within 2 points at 5.7 h, low 4.6 %
	MPC	within 2 points at 21.0 h, low 9.2 %

TRY IT IN THE APP 1 · Every Valve Fixed: The Open-Loop Shutdown

Load the scenario, leave the header on Manual, press Run and watch the reactor pressure. For two hours nothing seems to happen. Then the pressure climbs to the 3000 kPa g interlock at 3.59 h, and the plant stops. Switch the header to Auto and press Reset: the inventory loops now hold the pressure between 2702 and 2708 kPa g.

8.6.2 One Handle for Production

Case 3 lowers the production setpoint 15 %, to 19.51 m^3/h, ramped at 1.720 m^3/h per hour from 2 h. Because the index multiplies every flow setpoint, the ramp reaches all seven flow loops at once, and the production comes within 0.3 m^3/h of the new setpoint at 3.95 h. By 12 h the index reads 84.6 % and the production 19.29 m^3/h. The stripper level takes up the mismatch between what the reactor still makes and what the product valve draws, swinging from 28.7 to 81.1 %, because its loop is slow on purpose. At four times Ricker's gain the level stays between 42.3 and 63.6 %, but the production strays up to 0.55 m^3/h from its setpoint in force instead of 0.32, so a buffer held tight passes the disturbance on. This is the trade between a level and the flow it passes that Section 1.4.8 describes.

Pause. Try Exercise 8.3 before reading on.

8.6.3 Late Analyzers

In case 4 the A and C feed changes composition at 1 h, and the only measurement of the change is the reactor feed analyzer, reporting every 0.1 h what it sampled 0.1 h earlier. As shipped the A share bottoms at 45.29 % at 3.30 h and returns to 56.89 % by 12 h, with the A feed raised from 0.251 to 0.595 kscmh and the pressure peaking at 2752 kPa g. With both analyzer trims at zero gain the share falls to 22.94 % and the pressure reaches 2981 kPa g within 12 h. In case 9 the G setpoint moves from the 50/50 mode to the 10/90 mode of Ricker's table, 11.66 mol % there and 11.7 mol % in the application. At Ricker's ramp of 2.08 mol % per hour, G reaches within two points of 10 % by mass at 19.8 h and dips to 7.6 %. At 10 mol % per hour it arrives at 5.7 h and undershoots to 4.6 %, because the product analyzer reports a quarter of an hour late. The predictive controller, limited to the same ramp, arrives at 21.0 h with a lowest G of 9.2 %.

8.6.4 Noise, Pressure and Cost

Case 5 runs the random feed composition of IDV 8 for a day. The mean cost is 176.1 $/h against 170.7 undisturbed, so the noise costs about 5 $/h, while raising the pressure setpoint to 2800 kPa g lowers the mean to 146.3 $/h, because a higher pressure needs less purge to carry out the same inerts and loses less A, C, D and E with it. The price is margin, a peak of 2827 kPa g against 2756. The cost from minute to minute swings far wider than these differences, from 57 to 340 $/h between its 5th and 95th percentiles at 2705 kPa g, which is why the scorecard shows the mean. Moving such a setpoint toward its constraint is the work of an optimizer above the loops, as in Volume 13, Chapter 11. In case 6 the slow kinetic drift of IDV 13 runs 48 h without a trip, the pressure wandering over 287 kPa (2566 to 2854 kPa g) with the purge valve between fully shut and fully open, at a mean of 209.1 $/h. The inventory loops alone also hold the two days, at 212.3 $/h, because this drift is the pressure loop's to carry. In case 7 the reactor cooling valve sticks within a 2 % band, and with Ricker's gain the temperature cycles between 120.09 and 120.71 °C with a period of 4.8 min and the stem jumping between 38.7 and 43.4 %, against 120.39 to 120.41 °C without stiction. Halving the gain lengthens and enlarges the cycle and doubling it shortens and shrinks it, and none removes it, while the pressure stays between 2702 and 2708 kPa g, far from the override. That is why the override has its own case.

Pause. Try Exercise 8.2 before reading on.

> **TRY IT IN THE APP** **7 · Sticking Reactor Cooling Valve (IDV 14)**
>
> Run the scenario and watch the reactor temperature after the cooling water valve starts to stick at 1 h. With Ricker's gain of 8 it cycles by ±0.31 °C with a period of 4.8 min. Set the Reactor temperature gain to 4, press Reset and run again: the cycle grows to ±0.57 °C and slows to 8.1 min, and no gain removes it.

8.6.5 The A Feed Lost, and Ricker's Recording

Case 8 loses the A feed at 1 h. A still arrives with the A and C feed, whose 4.534 kscmh of A is 94.8 % of the A fed at the base case, but what is lost is the trim, since A now comes only in a fixed ratio to C. As shipped the structure trips on the stripper level at 8.40 h, and the override acts for only 2.0 min before it. The production loop chasing a rate that the reactor can no longer quite make scales the product draw with every feed and pulls the stripper empty. With the production setpoint cut to 17 m^3/h the override takes production off the index each time the filtered pressure passes 2950 kPa g, the pressure peaks at 2974 kPa g, and the plant settles at 5.88 m^3/h over its last four hours, a quarter of the base rate, alive. Without the override the same plant trips on high pressure at 13.34 h.

> **DID YOU KNOW?** **Fewer Loops, No Trip**
>
> From the simulator (Table 8.5). When the A feed is lost, the full decentralized structure trips on a low stripper level at 8.40 h. The inventory loops alone, with no production or composition loop above them, ride the same loss without a trip, at 21.94 m^3/h and 408 $/h.

The loss is not that deep for the plant. The inventory loops alone ride it at 21.94 m^3/h and 408 $/h, and Ricker's archive holds a recording of his 1996 strategy riding it at the base setpoint, ending at 20.82 m^3/h after 50 h with the pressure held near 2950 kPa g by his own override, as listed in Table 8.6. Under IDV 1, 8 and 13 the two structures agree within 68 kPa on the pressure extremes, 0.30 m^3/h on production and 0.81 mol % on G, and the stripper level is where they part, swinging two to three times as far here. Ricker's 1996 paper, whose own production-rate responses could not be consulted when the port was checked, is cited for the strategy [8], and the comparison rests on the recordings alone. Case 8 therefore teaches a property of this structure and not of the plant: a production loop that scales the product draw can empty a vessel that a gentler formulation would protect.

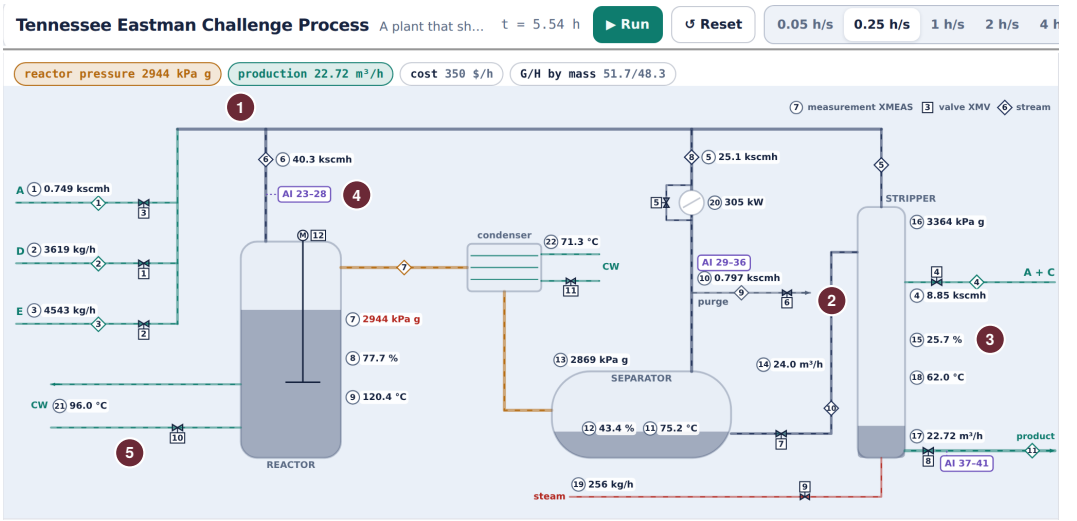

Figure 8.3 The suggested challenge at $t = 5.54$ h on the Plant view, decentralized structure with the override, reactor pressure setpoint 2800 kPa g, IDV 8 on and IDV 1 stepped at 2 h. The picture is a composite of the application's header above its Plant view, with the control panel, the trends and the Learn panel cut. The reactor pressure is at its peak of 2944 kPa g, 56 kPa below the interlock. The flowsheet carries every continuous measurement at its instrument, among them the A feed at 0.749 kscmh, the separator at 2869 kPa g, the stripper level at 25.7 % and the production at 22.72 m^3/h. Callout 1 marks the badge row of reactor pressure, production, operating cost and G/H split, 2 the purge valve (XMV 6) that the pressure loop moves, 3 the stripper level that case 8 pulls empty, 4 the reactor feed analyzer that reports late, and 5 the reactor cooling water valve (XMV 10) that sticks in case 7.

Table 8.6 Extremes Over 50 h, Disturbance at 1 h, Base Case: Ricker's Archived Recordings of His 1996 Strategy (R) and This Structure (A)

IDV		P_R (kPa g)	F_{11} (m^3/h)	L_C (%)	x_G (mol %)
1	R	2634–2747	22.25–23.37	36.0–65.8	51.97–55.78
	A	2599–2751	22.47–23.66	22.3–90.2	51.82–55.74
8	R	2617–2745	22.31–23.46	31.7–66.9	51.38–55.63
	A	2549–2783	22.09–23.77	−3.4–94.9	51.59–56.12
13	R	2532–2787	22.18–23.45	31.5–65.2	51.79–55.82
	A	2567–2854	22.41–23.63	19.9–91.8	50.98–55.82
6	R	2704–2961	20.55–23.21	40.3–56.4	52.32–55.11
	A	trip, stripper level low, at 8.40 h			

8.6.6 Infeasible at Maximum Production

An honestly infeasible case was also sought, the A feed lost at maximum production. At the base rate it is not infeasible, as the inventory loops show, so the claim is made at Ricker's tabulated maximum-production point, 36.04 m^3/h at 50/50 G/H. There every setting tried trips within a day: the decentralized structure with its override at four pairs of threshold and gain, without it, the inventory loops, the predictive controller, and production cuts to 8, 12, 17 or 18 m^3/h applied at the loss or earlier, sixteen settings in all, with trips between 5.08 and 18.92 h. Ricker's own strategy was not tested at that point, so the case says infeasible here, not infeasible for every structure.

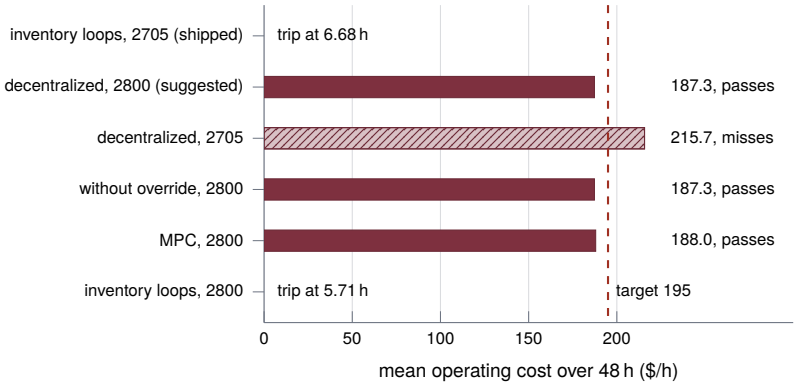

Figure 8.4 The challenge of Table 8.7 at the default seed, settings by reactor pressure setpoint in kPa g: mean cost of the runs that last 48 h (dark bars pass, the hatched bar misses), the target as a dashed line, and the trip times of the two runs on the inventory loops alone. Notice that the inventory loops trip at either pressure, and that the decentralized structure meets the cost target only when its pressure setpoint moves toward the interlock.

Table 8.7 The Challenge Over 48 h, Default Seed and Seeds 1 to 5

Settings	Default seed	Seeds 1–5
inventory loops, 2705 kPa g (shipped)	trip 6.68 h	trip 4.93–5.63 h
decentralized, 2800 kPa g (suggested)	passes, 187.3 $/h, G 2.33, 99.1 %	all pass, 156.7–179.0 $/h
decentralized, 2705 kPa g	misses, 215.7 $/h	4 miss (205.8–211.2), seed 1 passes at 186.9
without override, 2800 kPa g	passes, 187.3 $/h (identical)	—
MPC, 2800 kPa g	passes, 188.0 $/h	—
inventory loops, 2800 kPa g	trip 5.71 h	—

8.6.7 The Challenge, Redesigned

A first design of the challenge shipped the decentralized structure with its override removed and a detuned production loop. On the decentralized structure at the base pressure its disturbances raise the pressure no higher than 2899 kPa g, 51 kPa short of the override's threshold, and the override never acts, so the override was not what it tested. The challenge as built runs 48 h with IDV 8 throughout, IDV 1 at 2 h and production 10 % higher at 24 h, judged on true values: no shutdown, G within five points of 50 % by mass from 4 h, mean production at least 98 % of its setpoint, and a mean cost of at most 195 $/h. An attempt counts only with the model and the disturbances as given. It ships on the inventory loops at 2705 kPa g, and the suggested settings are the decentralized structure at 2800 kPa g. The results at six seeds are in Table 8.7, and Figure 8.4 draws the default seed.

An extended check separates instrument noise from the random feed path. With the challenge's physical IDV 8 walk fixed, the default realization and 30 additional measurement-noise streams all meet the target at both the one-second and half-second integration steps. Mean cost spans 187.08–187.88 $/h at one second and 168.42–169.94 $/h at half a second, while production remains above 99.06 % of its setpoint. The refined test keeps analyzer and disturbance clocks on physical time. Varying the physical walk as well is a different experiment: 30 of 31 days pass at each step, while seed 27 shuts down on low stripper inventory at 33.91 h with the native step and 33.82 h with the refined step. Thus the suggestion is reproducible for the fixed teaching day, including its instrument noise, but is not a robust operating policy for every feed trajectory. These wider results do not relax a shutdown or product target.

The inventory loops cannot follow the A/C step, as in case 2, and trip at 6.68 h. The decentralized structure keeps the plant running and the mix inside its band at either pressure, but at the base pressure

its mean cost of 215.7 $/h exceeds the target, while at 2800 kPa g the purge losses fall to a mean of 187.3 $/h with the worst G 2.33 points from 50 and the production at 99.1 % of its setpoint. The price is margin. After the A/C step the pressure peaks at 2944 kPa g at 5.54 h, as shown in Figure 8.3, and the straight-line estimate on the status strip reaches the interlock within the hour for 43.6 min, from 4.72 to 5.45 h. The override never acts in this run, so the structure without it scores identically. With half the integration step the shipped run still trips at 6.68 h and the suggested one still passes, at 169.3 $/h, the change coming from twice as many noise samples reaching the loops.

8.7 Technology in Practice

What Measures It

The inventory loops of this plant read pressures, levels, temperatures and flows, and the composition loops read three sampled analyzers. In the model the reactor feed and purge analyzers report every 0.1 h what they sampled one period earlier, and the product analyzer does the same every 0.25 h. In a plant such analyzers are process chromatographs, whose separation by time makes some measurement dead time inevitable, and whose sample line would delay the measurement grossly without a continuous flow to waste [18]. The 22 continuous measurements of the model carry only additive noise of the published standard deviations, so no measurement of the model ever fails or drifts.

> **TECHNOLOGY Sensor: The Process Chromatograph and Its Sample Line**
>
> The compositions of a plant like this one come from process chromatographs, which separate the species of a sample by the time each takes to pass through a column. That separation is why cycle times of minutes are common even on line, and some dead time is inevitable. Enough of it makes a feedback loop oscillate on its own. The sample must also reach the analyzer: a continuous flow from the process to waste keeps a fresh sample at the instrument, and without it the gas would take a very long time to travel the line. The sample conditioning, with its filtering, heating or cooling and drying, is one of many points of failure, and keeping analyzers working is treated as an advanced skill [18].

What Moves It

Each of the twelve valves follows its command through a first-order lag of 5 to 9 s, and 120 s for the steam valve. Under IDV 14, 15 and 19 a valve keeps its last command until the new one differs by more than 2 %. That rule is a dead band of the kind that stem friction gives a real valve [18], and case 7 shows the limit cycle it drives in the reactor temperature loop. Figure 8.5 sets the decentralized structure on such hardware, from each sensor to the valve it moves.

> **TECHNOLOGY Final Element: A Control Valve That Sticks**
>
> Friction between a valve stem and its packing is inevitable, and static friction is larger than the friction of a moving stem. A pneumatic actuator therefore holds the stem still until enough pressure builds to break it free. The stem then jumps until friction grabs it again, so a smooth signal turns into stair steps. This is stiction. To reverse direction the actuator must overcome the friction twice, which acts as a dead band in the loop. Over-tightened packing nuts are a common cause, and a test of small reversals of the actuator pressure measures the total friction. Many digital positioners record the actuator force against stem travel, and the gap between the opening and closing traces grows with the friction [18].

Where the Controller Runs

In a plant the seventeen loops of Table 8.3 would run in a distributed control system, whose supervisory level oversees localized controllers. Loops there are often nested or cascaded, with cycle times from milliseconds to minutes [19]. The application executes every loop at each 1 s step, and its predictive controller decides every 0.05 h above them. Supervisory control in the process industry is advanced regulatory control, a setpoint-based predictive controller on top of the PID layer, or a combination,

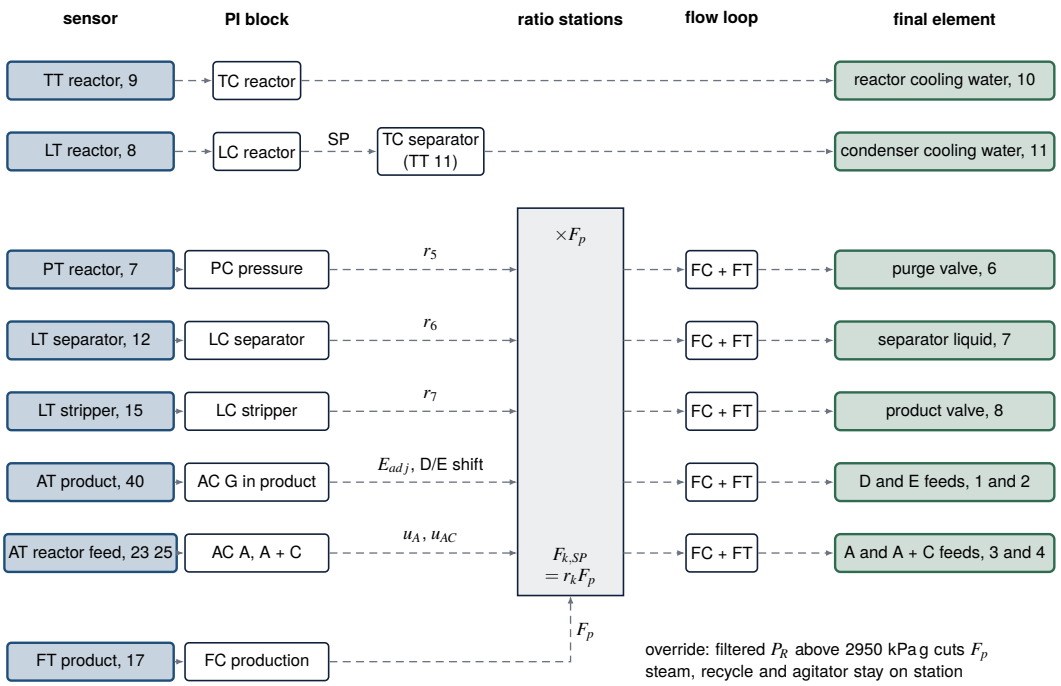

Numbers are XMEAS for sensors and XMV for valves. In a plant the blocks run in the distributed control system, and the application executes all of them at every 1 s step.

Figure 8.5 The decentralized structure on plant hardware, one row per loop of Table 8.3. The two temperature rows act on cooling water directly, and the reactor level sets the separator temperature. The five rows below send ratios to the ratio stations, which multiply each by the production index F_p set by the production loop at the bottom, and every product of a ratio and the index is a flow setpoint for a flow loop on its valve. Notice that one handle, F_p, moves every flow at once.

often implemented in the same distributed control system [15]. The production index with its ratio stations is such an advanced regulatory layer, and the chapter's predictive structure is the setpoint-based kind. An override is built with a select function that lets only one controller output reach the process at a time, and the de-selected controller must not wind up while it has no control [18]. The chapter's override meets that need by writing the index it applies back into the production loop's output.

> **IN PRACTICE** **Nameplate First, Optimization Later**
>
> From startup the primary objective of a new plant is to reach its nameplate capacity in a reliable and predictable way. Optimization often comes only after the plant has run for a few years, and a top-down strategy can then be added if the plant has a good regulatory control system. A strategy that drives the plant to the right economics in a counterintuitive way is hard to keep in service, because operator understanding suffers [3]. A dynamic model is usually not available at startup either, so a new plant first gets a simpler system of advanced regulatory elements, and predictive control is considered only if that system performs poorly [15].

What Is Modernizing It

Two developments bear on this plant. The first is the choice between advanced regulatory control and model predictive control in the supervisory layer. Constraints, often claimed as the special

advantage of predictive control, can in most cases also be handled by regulatory elements such as selectors, split-range control and anti-windup. Yet such solutions can grow complex, with many layers of cascades and selectors, and predictive control may then be the simpler one to implement and understand [15]. On the challenge the chapter's predictive structure passes at 188.0 $/h against 187.3 for the decentralized loops (Table 8.7). The second is valve diagnostics. Digital positioners that monitor the actuator's force against the stem's motion turn friction into a measured signature [18]. Case 7 shows why that matters: no gain of the loop removes the cycle of a sticking valve.

8.8 Limitations

The model is the public code, so it carries that code's simplifications: equilibrium vessels, a stripper of fixed separation factors, algebraic flows and a chemistry named by letters only, which maps onto no named process. Its limits and prices are the benchmark's teaching values. The port has been compared point by point with two translations, but the Fortran was checked by reading, not by compiling and running it, and the 1993 and 1996 papers were not available when the port was checked, so nothing printed in them is quoted or compared. The open-loop trip time depends on the rounding of the constants, by 0.59 h across the three readings measured. The noise realizations differ from the original generator's, so a trajectory published with the original code is matched in its statistics and not point by point.

The decentralized structure is the archived Simulink one, not the 1996 strategy, and it differs from Ricker's recordings in the stripper level and under the A feed loss. The override and the predictive controller are this application's designs, the predictive model is linear and identified at the base case, and the controllers meet no model error beyond what the plant itself presents. The individual case results use the application's seed, with the additional realizations stated beside each result. The challenge also separates 31 instrument-noise streams from 31 physical feed paths, as reported above. Half-step comparisons preserve the physical analyzer and disturbance clocks.

> **OPEN CHALLENGE Coordinating the Layers**
>
> The plantwide procedures written around plants like this one say what each layer of control must do, but many research issues remain in filling in the detail of each step, above all in the bottom-up part. More work is needed to understand how to decompose the layers of a control system and coordinate them [2]. A review of 2000 judged the field to be at a relatively early stage and called for a common ground on which case studies could be compared. It expected that, once alternative structures could be generated and analyzed automatically, the main problem would be building process models efficiently [4]. A predictive layer shares that need, because obtaining and maintaining its full dynamic model is costly [15].

8.9 Conclusion

The Tennessee Eastman plant runs here as a JavaScript port of the public Fortran model, translated through a BSD-3-Clause Python port, checked against the Fortran by reading, against the Python port to 1.76×10^{-11} and against Ricker's C translation bit for bit at the base case and to 4×10^{-10} over six hours, with identical shutdown steps. Its base case is a steady state that is not stable, and its open-loop trip at 3.47, 3.59 or 4.06 h identifies the rounding of the constants more than the plant. The inventory loops hold it until a feed changes, the archived decentralized structure rides the feed change and cuts the mean cost from 215.7 to 187.3 $/h when the pressure setpoint moves from 2705 to 2800 kPa g, and the same structure trips at 8.40 h on an A feed loss that Ricker's recorded strategy rode. Only an integrating override and a production cut keep it alive. The port omits the 2015 revision's added disturbances, and its noise matches the original only in distribution. A fuller port would add the 1996 formulation with its overrides, the revision's disturbances and measurements, and the remaining operating modes as challenges. The simulator puts the plantwide benchmark of three decades where a student can watch it shut down.

KEY POINTS

- Left with every valve fixed and no disturbance, the published steady state does not hold: the reactor pressure reaches its interlock at 3.59 h.
- The inventory loops make the plant stable, and only the composition loops above them make it survive a feed change.
- Moving the pressure setpoint from 2705 to 2800 kPa g cuts the challenge's mean cost from 215.7 to 187.3 $/h, and the price is margin.
- The archived structure trips at 8.40 h on an A feed loss that Ricker's recorded strategy rode. Only an integrating override and a production cut keep it alive.
- A trip time quoted for this plant identifies an implementation as much as a dynamic.

Where else this applies. An override that sits idle in normal operation also guards the compressor of Volume 4, Chapter 2, and an inert purged from a gas recycle under a pressure override returns in Volume 13, Chapter 2.

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] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[4] T. Larsson and S. Skogestad, "Plantwide control - A review and a new design procedure," *Modeling, Identification and Control: A Norwegian Research Bulletin*, vol. 21, no. 4, pp. 209–240, 2000.

[5] A. Bathelt, N. L. Ricker, and M. Jelali, "Revision of the Tennessee Eastman process model," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 309–314, 2015.

[6] S. Yin, S. X. Ding, A. Haghani, H. Hao, and P. Zhang, "A comparison study of basic data-driven fault diagnosis and process monitoring methods on the benchmark Tennessee Eastman process," *Journal of Process Control*, vol. 22, no. 9, pp. 1567–1581, 2012.

[7] J. J. Downs and E. F. Vogel, "A plant-wide industrial process control problem," *Computers & Chemical Engineering*, vol. 17, no. 3, pp. 245–255, 1993.

[8] N. L. Ricker, "Decentralized control of the Tennessee Eastman challenge process," *Journal of Process Control*, vol. 6, no. 4, pp. 205–221, 1996.

[9] T. J. McAvoy and N. Ye, "Base control for the Tennessee Eastman problem," *Computers & Chemical Engineering*, vol. 18, no. 5, pp. 383–413, 1994.

[10] P. R. Lyman and C. Georgakis, "Plant-wide control of the Tennessee Eastman problem," *Computers & Chemical Engineering*, vol. 19, no. 3, pp. 321–331, 1995.

[11] A. Banerjee and Y. Arkun, "Control configuration design applied to the Tennessee Eastman plant-wide control problem," *Computers & Chemical Engineering*, vol. 19, no. 4, pp. 453–480, 1995.

[12] N. L. Ricker, "Optimal steady-state operation of the Tennessee Eastman challenge process," *Computers & Chemical Engineering*, vol. 19, no. 9, pp. 949–959, 1995.

[13] N. L. Ricker and J. H. Lee, "Nonlinear model predictive control of the Tennessee Eastman challenge process," *Computers & Chemical Engineering*, vol. 19, no. 9, pp. 961–981, 1995.

[14] T. Larsson, K. Hestetun, E. Hovland, and S. Skogestad, "Self-optimizing control of a large-scale plant: The Tennessee Eastman process," *Industrial & Engineering Chemistry Research*, vol. 40, no. 22, pp. 4889–4901, 2001.

[15] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[16] 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.

[17] 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.

[18] 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.

[19] 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

Three translations of the same public code are run open loop, with every valve fixed and no disturbance, on the same 1 s Euler step. They trip on high reactor pressure at 3.47, 3.59 and 4.06 h. What does the spread show?

(a) The Euler step is too coarse, so each translation integrates the unstable mode with its own error

(b) One translation carries an error in its balances that the check of the base case did not catch

(c) The measurement noise differs between the translations, and the noise moves the time of the trip

(d) The unstable mode grows from the rounding residue of the decimal constants, which each reading leaves

8.2 Concept check · 1 min

In case 5 the feed composition wanders at random for a day, and raising the reactor pressure setpoint from 2705 to 2800 kPa g lowers the mean operating cost from 176.1 to 146.3 \$/h. Which mechanism does the chapter give?

(a) A higher pressure carries the same inerts out in less purge gas, so less A, C, D and E leave with it

(b) A higher pressure speeds up the reactions, so the plant makes more product from the same feeds

(c) A higher pressure lowers the compressor work, which is the largest term of the operating cost

(d) A higher pressure damps the random feed walk, so the composition loops move the feeds less

8.3 Concept check · 2 min

Case 3 ramps the production setpoint down by 15 %. The run is repeated with the stripper level gain raised to four times Ricker's value. What does the chapter measure?

(a) The stripper level swings less, and the production also stays closer to its setpoint in force

(b) The stripper level swings less, and the production strays further from its setpoint in force

(c) The stripper level swings as before, because its 200 min integral time sets the swing

(d) The stripper level begins to cycle, and the plant trips on the stripper level interlock

8.4 Theory and calculation · 6 min

The decentralized structure sets the production index from Equation (8.9), and its production setpoint passes a rate limit of 30 % of the base rate of 22.95 m^3/h per day before it is in force.

(a) With the production loop's trim at zero, find the index F_p for a setpoint of 22.95 m^3/h and for the challenge's raised setpoint of 25.25 m^3/h. *Checkpoint:* 100.26 % and 110.31 %.

(b) The challenge raises the setpoint from 22.95 to 25.25 m^3/h at 24 h. Find the rate limit and the time the setpoint in force needs to reach the new value. *Checkpoint:* 0.287 m^3/h per hour and 8.0 h.

(c) The challenge scores the mean true production over 48 h against the mean of the setpoint the reader enters, not the one in force. Find the score a plant would get if its production followed the setpoint in force exactly, and compare it with the 99.1 % of the suggested run. *Checkpoint:* 99.2 %.

8.5 Theory and calculation · 6 min

The pressure override of Equation (8.10) filters the measured reactor pressure over 0.02 h at the 1 s step, and then takes $K_{ovr} = 2$ points of F_p per kPa of filtered excess every second.

(a) Find the filter coefficient $1 - \mathrm{e}^{-\Delta t/0.02}$ per step and the filter's time constant in steps. *Checkpoint:* 0.0138 and 72 steps.

(b) The measurement noise is 0.298 kPa rms and independent from step to step. Driven by such noise, this filter's output has a variance $\alpha/(2-\alpha)$ times that of its input, where α is the coefficient of part (a). Find the rms noise left after the filter. *Checkpoint:* 0.025 kPa, twelve times less.

(c) With the production loop frozen at the override's first step in case 8, F_p fell from 83.8 to 42.1 % in 3 min. Find the mean filtered excess this needs, and compare it with the largest excess of 1.11 kPa. *Checkpoint:* 0.116 kPa.

8.6 Theory and calculation · 6 min

At the base case the port reads a compressor work $W = 341.43$ kW, a steam flow $S = 230.31$ kg/h and a purge $F_9 = 0.3371$ kscmh of 32.96, 13.82, 23.98, 1.257, 18.58, 2.263, 4.844 and 2.299 mol % A to H. The product holds 0.0179, 0.8357 and 0.0986 mol % D, E and F, with the stripper valve at its base opening of 46.534 %. Use Equation (8.8) and its prices.

(a) Find the power and steam terms. *Checkpoint:* 18.30 $/h and 7.32 $/h.

(b) Find the purge term and its share of the total of 170.62 $/h. *Checkpoint:* 114.73 $/h, 67 %.

(c) Find the product-loss term, and the same term with the composition printed in Ricker's table, 0.02, 0.84 and 0.10 mol %. *Checkpoint:* 30.27 $/h and 30.56 $/h, 0.28 $/h apart.

8.7 App challenge · 8 min

Use the default seed and change nothing that the parts do not name.

(a) Load scenario 1, Every valve fixed: the open-loop shutdown, and run it in Manual at the default speed of 0.25 h per second until it stops. Read the time and the cause on the status strip, and in the exported file the true reactor pressure, column `truth_xmeas07_reactor_press_kpag`, at 2 and 3 h. *Checkpoint:* 3.59 h, reactor pressure high. 2705.0 and 2689.6 kPa g, within 0.5 kPa.

(b) Load scenario 2, Inventory loops only, and run it as shipped at the same speed until it stops. Find in the exported file the first row with the true reactor pressure above 2900 kPa g, the shutdown time, and the lowest A share of A + C in the reactor feed analyzer, column `yA_pct`. *Checkpoint:* 4.85 h within 0.02 h, 5.84 h, 37.0 % within 0.3.

(c) Choose Decentralized with pressure override in Structure (in Auto), press Reset and run 8 h at 1 h per second. Read the highest true reactor pressure and the lowest A share in the new file. *Checkpoint:* 2751 kPa g within 2 kPa, 45.4 % within 0.3.

Explain the three results in terms of the layers of the structure.

8.8 App challenge · 8 min

Load scenario 8, A feed lost (IDV 6) and the pressure override. The A feed is lost at 1 h. Use the default seed and a speed of 4 h per second.

(a) Run the scenario as shipped until it stops, and read the time and the cause on the status strip. *Checkpoint:* 8.40 h, stripper level low.

(b) Press the reload button, set the Production setpoint to 17 m^3/h before Run, and run 24 h. In the exported file find the highest true reactor pressure and the mean true production, column `truth_xmeas17_strip_underflow_m3h`, over the rows after 20 h. *Checkpoint:* No shutdown. 2974 kPa g within 2 kPa, and 5.88 m^3/h within 0.02.

(c) Choose Decentralized without override in Structure (in Auto), press Reset and run again. *Checkpoint:* 13.34 h, reactor pressure high.

Say which part of the structure each result belongs to.

8.9 App challenge · 8 min

Load the Challenge, Two days of upsets at lower cost. Use the default seed and a speed of 4 h per second, and leave the model and the disturbances as given.

(a) Run it as shipped and read the time and the cause of the stop. *Checkpoint:* 6.68 h, reactor pressure high.

(b) Choose Decentralized with pressure override in Structure (in Auto), press Reset and run the 48 h. Read the Mean cost tile and the note at the end of the run. *Checkpoint:* 215.7 $/h, missed on cost (worst G 2.42 points, production 99.1 %).

(c) Set the Reactor pressure setpoint to 2800 kPa g, press Reset and run again. Read the same note, and find in the exported file the highest true reactor pressure and its time. *Checkpoint:* Target met, 187.3 $/h, worst G 2.33 points, production 99.1 %. 2944 kPa g near 5.53 h, within 2 kPa and 0.05 h.

Accept a mean cost within 0.1 $/h. Explain what the structure and the pressure setpoint each contribute, and what the second change costs.

8.10 Application area · 4 min

An operations engineer wants to move a plant like this one from a reactor pressure setpoint of 2705 to 2800 kPa g, on the evidence of case 5 and the challenge.

(a) Case 5 gives mean costs of 176.1 and 146.3 $/h at the two setpoints. Find the saving over a year of 8760 h. *Checkpoint:* About $261 000 per year.

(b) The peak pressures are 2756 and 2827 kPa g in case 5 and 2944 kPa g in the challenge at 2800 kPa g. Find each margin to the 3000 kPa g interlock, and name what must be established before the move on a real plant. *Checkpoint:* 244, 173 and 56 kPa.

8.11 Application area · 4 min

The two composition loops and the mix loop of the decentralized structure read sampled analyzers, as in Equation (8.7).

(a) Find the youngest, oldest and mean age of a reading of the product analyzer, period 0.25 h, and of the reactor feed analyzer, period 0.1 h. *Checkpoint:* 0.25 to 0.5 h, mean 0.375 h. 0.1 to 0.2 h, mean 0.15 h.

(b) Case 9 moves the G setpoint from 53.72 to 11.7 mol %. Find how long the setpoint ramp lasts at Ricker's 50 mol % per day and at 10 mol % per hour, and how far the setpoint moves during the age of the oldest product reading in each case. *Checkpoint:* 20.2 h and 4.2 h. 1.04 and 5 mol %.

8.12 Review the agent · 5 min

An AI assistant is asked why scenario 8, A feed lost (IDV 6) and the pressure override, ends at 8.40 h and replies as follows. "The loss of the A feed raises the reactor pressure, and the override, at 2 points of F_p per kPa per second, is too gentle to pull production down in time. Raise the Override gain to 5 and the decentralized structure will ride the loss at the full production setpoint." Find the flaw in the reply and say how it could be checked.

CHAPTER 9

GLP-1 Peptide Purification Where a Chromatography Cut Is Judged at the Filter

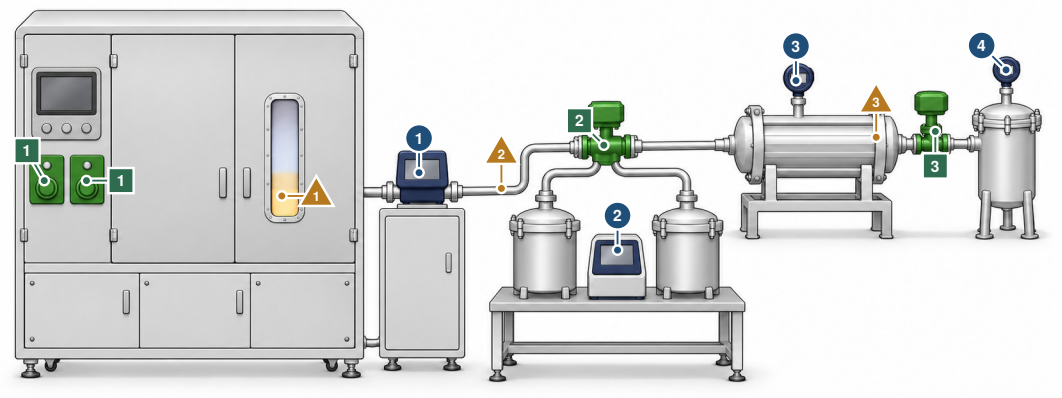

The controller moves 1 column pump and solvent blend 2 cut valve to pool, side cut or waste 3 membrane retentate valve

The controller reads 1 UV detector at the column outlet 2 laboratory assay of the pool 3 membrane pressure 4 filter pressure

What disturbs it 1 crude purity and its late variant 2 extra line before the cut valve 3 membrane passes or fouls more

The cut valve sits a line's length after the detector, and what it lets into the pool decides, hours later, how fast the cake on the filter drains. Illustration, not to scale.

IN PLAIN TERMS **One Cut Decides What the Filter Gets**

The main idea. This chapter follows a small protein drug through the steps that clean it. A tube of fine grains lets each part come out at its own time, and a valve keeps the middle part.

The challenge. A bit more of the flow can still pass the purity test, yet an impurity that comes out last makes the solid grains small. Small grains pack into a cake that drains slowly, and the batch runs out of filter time.

Why it matters. A choice made at the first step has to be judged by what the last step needs.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 9.2, the try-its and the two facts in Section 9.6, and the key points at the end of Section 9.9
Study	40 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question A chromatography cut is usually judged by the purity of its pool, but the filter two unit operations later answers to one late-eluting variant that the purity does not single out. Which variable should the cut hold, and what does a collection rule driven by a soft sensor add over a fixed window?

Learning objectives

- Explain why a pool that passes its purity assay can lose the batch on the filter clock, and name the variable that separates the cuts that pass from those that fail.
- Compute retention factors, selectivity, diafiltration losses and the split of a cake resistance from the chapter's model equations.
- Correct a cut set in minutes for an unmeasured holdup between the detector and the valve.
- Predict what a higher filter pressure buys on a compressible cake.
- Judge a soft-sensor collection rule against a fixed window on the same crude and on a crude outside its calibration.

Prerequisites Sections 1.4.2 (choice of controlled variable), 1.4.8 (disturbance propagation), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios, scorecard and export).

The app https://apmonitor.com/apps/glp1-manufacturing/

See also Volume 8, Chapter 11 for a finite-volume column switched on a capacity nobody measures, and Volume 8, Chapter 4 for a crystal population carried in its leading moments. Volume 5, Chapter 7 for material tracked in volume to a diversion valve, and Volume 5, Chapter 10 for a soft sensor whose bias a delayed laboratory result corrects.

9.1 Introduction

The purity assay of a chromatography pool answers one question about the material, how much of it is product. In the simulated campaign of this chapter, a cut from 23 to 31 min of a 60 min gradient pools 61.7 g of peptide at 98.92 % purity, and 56.2 g of it leaves the filter as acceptable product, 9.9 g/h of campaign time. Shutting the cut 0.75 min later adds 4.4 g to the pool, and the laboratory assays it at 98.35 %, still inside the 98.5 % limit of the isolated solid. The batch is nevertheless lost. The extra minute carried the pool's late-eluting variant from 0.358 to 0.939 %, that variant slows the growth of the precipitating solid, the mean particle size falls from 11.6 to 9.0 µm, and the cake on the filter needs 79.3 min against the 60 min its schedule allows.

The simulator is built around a belief that this case contradicts, that a chromatography cut is right when its pool meets the purity limit. The pool assay is a sufficient statistic only when the next unit operation responds to the total impurity and not to which impurity it is. A cooling crystallizer whose impurity rejection depends on the impurity, a filter whose cake depends on the particle habit, and a dryer whose endpoint depends on the solvent the cake occluded have the same shape. The filtration of an active pharmaceutical ingredient (API) can become the bottleneck of its whole process train, and its specific cake resistance is estimated from laboratory filtrations before scale-up [1]. Gradient purification of biomolecules needs a central cut between weakly and strongly adsorbing impurities, and a single batch step often cannot purify satisfactorily [2]. Kinetic models describe how an impurity in solution slows crystal growth [3], which is the link from the first problem to the second.

However, these effects are usually studied one unit at a time. Preparative chromatography models are fitted and optimized against the band and the pool [4, 5], and filtration studies estimate the cake a given slurry makes [1]. This chapter sets side by side, in a simulated campaign whose column, membrane and filter process the same material, local purity-maximizing cuts, recovery-maximizing cuts and a downstream-aware collection rule, and reports what each costs in membrane loss, solvent use, filtration time and isolated mass. In the terms of Section 1.4.2, the question is which variable the cut should hold. The finite-volume transport, the competitive isotherm, the population balance, the compressible-cake law and the velocity-form proportional-integral (PI) loops are established methods, and every comparison below is measured on a single declared teaching model. The model has 677 states, and its product ledger closes to 2×10^{-11} of the charge in every case measured. Controllers and collection rules read instruments only.

Section 9.2 follows the domain from peptide synthesis to isolation. Section 9.3 gives the route and the model, with the equations and defaults from which the simulator can be rebuilt, and Section 9.4 the recipe, the loops and the three collection rules. Section 9.5 separates code verification from agreement with sources. Section 9.6 reports the nine cases, the challenge and the comparison of cut policies, Section 9.7 describes the real detectors, pumps and controllers, Section 9.8 states the limitations, and Section 9.9 draws the conclusions. The cases and the challenge carry completion rules, and the exercises of Section 9.10 are built on them.

9.2 Peptides on Resin, Peptides on Silica

Peptide manufacture by synthesis begins with the solid-phase method of 1963, in which the chain is assembled step by step on an insoluble resin and the excess reagents of each coupling are washed away [6]. Every coupling that falls short leaves a deletion sequence and every side reaction a related variant, so a crude peptide of thirty residues is a family of molecules rather than one. GLP-1, the hormone behind this class, has its own history, told in the box below. Drug design then turned the native hormone into acylated analogues that bind serum albumin, and semaglutide was selected as a once-weekly candidate with two amino acid substitutions and a derivatized lysine at position 26 [7].

> ⧖ **HISTORY** **Early 1980s to 2012: From a Gene Sequence to a Weekly Medicine**
>
> In the early 1980s the proglucagon sequences of the anglerfish and then of the rat, hamster, cow and human showed that glucagon is cut from a larger precursor. The mammalian precursors hold glucagon and two related peptides, named GLP-1 and GLP-2 [8, 9]. The work of the following three decades established GLP-1 as an insulin-stimulating glucoregulatory hormone. Injected, it proved disappointingly ineffective: an enzyme cleaves it, and in people with diabetes its half-life in the circulation is only 1.5 to 2 minutes. Analogues that resist the enzyme last longer. The first receptor agonist approved for clinical use, a synthetic form of a peptide isolated from lizard venom, reached approval for twice-daily injection in 2005, and a once-weekly form followed in 2012 [9].

Purification grew up with high-performance liquid chromatography, whose gradient elution was given a theoretical treatment in 1979 [10]. Preparative work rests on a theory of nonlinear, nonideal chromatography whose models are used to optimize operating conditions, and by 2002 preparative chromatography was becoming an established purification method in the pharmaceutical industry [4]. For peptides the dependence on the modifier is steep. A heptadecapeptide's retention factor changed by more than 400 % between 16 and 19 % acetonitrile, and its overloaded gradient band was modeled with reasonable accuracy by a Langmuir isotherm whose parameters follow the linear solvent strength model [5]. Because the product elutes between weakly and strongly adsorbing impurities, a batch gradient run needs a central cut, and multicolumn countercurrent solvent gradient purification was introduced to recycle the impure side streams internally [2].

The public record of GLP-1 manufacture is mostly in patents, which disclose process options rather than validate a train. One describes multidimensional reversed-phase purification of a GLP-1 analogue, desalting by tangential-flow filtration (TFF) on a membrane of 1 kDa cutoff, and the propensity of glucagon-like peptides to aggregate at acidic pH [11]. A second purifies semaglutide in three reversed-phase passes from a crude of about 60 to 70 % purity and names lyophilization, spray drying and antisolvent addition among the isolation options [12]. A third, for a dual GLP-1R/GIPR agonist, exchanges acetonitrile for 2-propanol/water by diafiltration over 4 to 8 diavolumes and isolates the peptide by antisolvent precipitation followed by filtration and drying, by spray drying or by freeze drying [13]. The volumes involved are large. A perspective written for an industry roundtable team describes peptide manufacture as dependent on large amounts of hazardous reagents and solvents and names greener processes a critical unmet need [14].

The isolation end has older foundations. General equations for cake filtration were derived in 1935 [15], membrane processes became a reviewed technology of bioprocessing in their own

Table 9.1 Variables by Role, With Defaults and Ranges

Symbol	Role	Default	Range
Q	CV, column flow (FIC-201)	1.5 L/min	0.3–2.5
φ_{in}	CV, modifier at the inlet (AIC-202)	program	10–70 %
pH	CV, conditioning (AIC-401)	7.3	3.5–8.5
TMP	CV, membrane (PIC-501)	2 bar	0.5–6
ΔP	CV, filter (PIC-601)	1.5 bar	0.2–4
u_P, u_B	MV, feed pump, B valve	—	0–100 %
u_{pH}	MV, acid (−) or base (+)	—	±100 %
u_V, u_G	MV, retentate valve, gas regulator	—	0–100 %
t_1, t_2	decision, cut opens, shuts	23, 31 min	10–75
rule	decision, times, UV or downstream aware	times	—
$\hat{x}_{L,max}$	decision, shut limit of the soft sensor	0.80 %	0.1–3
m_{load}	decision, crude charged	120 g	20–260
N_D	decision, diavolumes	4	0–12
F_s	decision, antisolvent rate	0.06 L/min	0.01–0.5
x_c, x_ℓ	DV, crude purity, late share	70, 23 %	45–90, 5–60
ΔV_c	DV, extra collector holdup	0 L	0–3
$s_\times$, $f_\times$	DV, sieving, fouling multiples	1, 1	0.5–12, 0–6
c_i, q_i, φ	states, 60 cells	—	—
V, Z, C_T, M_k	states, vessel	—	—
M_c, V_{cake}	states, filter	—	—

right [16], and the population balance gave particle technology a way to carry a size distribution through a dynamic model [17]. What none of these supply is a peptide's cake resistance, and no source consulted for this chapter reports one. The route below is therefore a proposed teaching route assembled from these disclosed options, and every filtration coefficient in it is a teaching value.

9.3 The Route as Modeled

One campaign charges 120 g of crude at 70 % peptide, dissolved in 3 L, onto a column 0.20 m across and 0.25 m long (7.85 L of bed), as shown in Figure 9.1. A gradient from 22 to 45 % organic over 60 min at 1.5 L/min elutes the bands past a UV detector to the cut valve XV-203, which routes the stream to the pool, a side cut or waste. The pool is diluted with 4 L of buffer in the conditioning vessel T-401, set to pH 7.3, concentrated on a 1 m^2 membrane to 3 L, diafiltered with four diavolumes of 60 % 2-propanol in water, and concentrated again to 100 g/L. Antisolvent dosed at 0.06 L/min to a 60 % fraction precipitates the peptide in R-701, and an agitated Nutsche filter of 0.05 m^2 forms, washes and dries the cake at 1.5 bar. The train is a teaching assembly of the options in Section 9.2, not a reproduction of any manufacturer's process, and its quality scores describe simulated attributes, not batch release.

The variables are classified in Table 9.1. The model carries 677 states, 60 column cells of four mobile concentrations, four stationary loads and the modifier, two holdup lines of five and four tanks, the pump and mixer, three collection vessels, the vessel's dissolved species, solvents, buffer, solids, four moments, membrane adsorption and fouling, the permeate, the filter with its free liquor, suspended solids, cake and cake liquor, and seven ledgers. Physics is in SI units, and the display shows minutes, liters, grams, g/L and bar.

9.3.1 The Column and Its Binding Law

The column is one-dimensional transport with axial dispersion and a lumped kinetic exchange for each species i (P, E, L and the unretained smalls X), written on the mobile-volume basis, as shown in

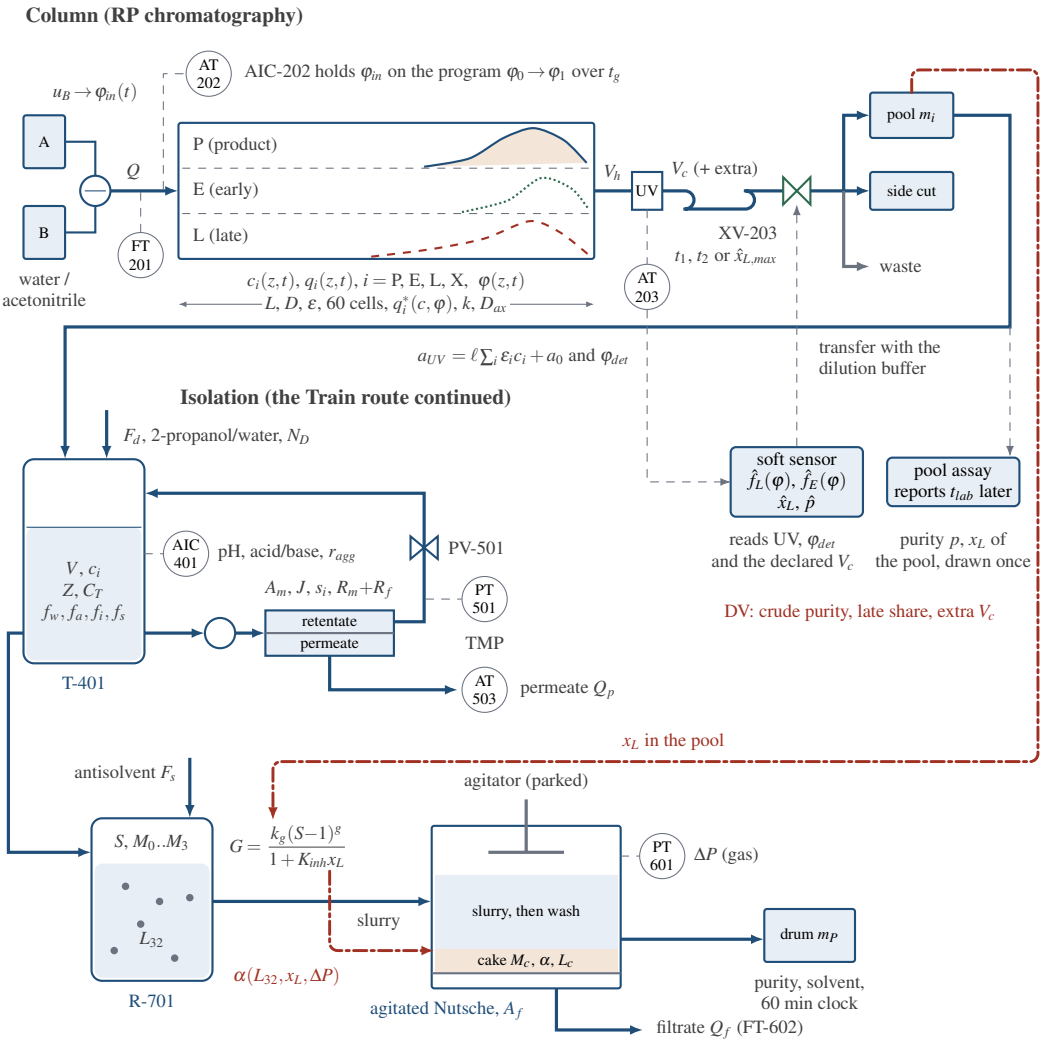

Figure 9.1 GLP-1 purification and isolation as modeled, the Column view's rig continued along the Train route. The gradient pump blends A and B to $\varphi_{in}(t)$ at flow Q (FT-201, AT-202), the bed of length L and diameter D carries the bands $c_i(z,t)$ and loads $q_i(z,t)$ of product P, early variant E and late variant L, and the stream passes the holdup V_h, the UV cell AT-203 and the collector holdup V_c to XV-203, which cuts on t_1 and t_2 or on the soft sensor's $\hat{x}_L$. The pool's assay reports after t_{lab}. T-401 holds V, c_i, the strong ion Z and buffer C_T, and recirculates through the membrane (A_m, J, s_i) under PV-501 and PT-501 with permeate UV AT-503 and diafiltration buffer F_d. R-701 precipitates with antisolvent F_s, and the Nutsche at ΔP (PT-601) holds the cake M_c of resistance α. The thick dash-dotted line (red) is the coupling this chapter studies, the late variant x_L slowing growth G, which sets the Sauter size L_{32} and through it α. In the bed the product band is solid, the early variant dotted and the late variant dashed. Process model only, not to scale.

Equation (9.1).

$$\begin{aligned}
&\frac{\partial c_i}{\partial t}+\frac{1-\varepsilon}{\varepsilon}\frac{\partial q_i}{\partial t}+v\frac{\partial c_i}{\partial z}=D_{ax}\frac{\partial^2 c_i}{\partial z^2}, \qquad v=\frac{Q}{\varepsilon A_c}\\
&\frac{\partial q_i}{\partial t}=k_i\big(q_i^{*}(c,\varphi)-q_i\big), \qquad \frac{\partial \varphi}{\partial t}+v\frac{\partial \varphi}{\partial z}=D_{ax}\frac{\partial^2 \varphi}{\partial z^2}\\
&v c_{i,in}=v c_i(0)-D_{ax}\frac{\partial c_i}{\partial z}(0), \qquad \frac{\partial c_i}{\partial z}(L)=0
\end{aligned} \tag{9.1}$$

where c_i is mass per mobile volume, q_i mass per stationary volume, $\varepsilon = 0.60$ the effective mobile fraction with pore liquid included, A_c the bed area and $k_i = k$ except $k_X = \max(k, 0.2\,\mathrm{s}^{-1})$. The modifier φ travels with the same transport and does not adsorb, so each cell sees the gradient when the liquid carrying it arrives. The isotherm is a competitive Langmuir with one capacity whose affinities fall exponentially with the modifier, as shown in Equation (9.2).

$$q_i^{*}=\frac{q_{max} b_i c_i}{1+\sum_j b_j c_j}, \qquad b_i(\varphi)=b_{i,ref}\,\mathrm{e}^{-S_i(\varphi-\varphi_{ref})} \tag{9.2}$$

At vanishing load the retention factor is $k_i' = q_{max} b_i (1-\varepsilon)/\varepsilon$, which is 10.0 for the product at the reference 30 % organic and 1215 at the 22 % the charge is loaded at. The bed saturates near 157 g of peptide. The ratio $b_L/b_P = 1.73$ at φ_{ref} is the selectivity that separates the product from its late variant. The charge enters at $c_{i,in} = m_{load} f_i / V_{feed}$ during the load and zero after it, where the crude fractions are $f_P = x_c$, $f_E = 0.27(1-x_c)$, $f_L = x_\ell(1-x_c)$ and f_X the rest. The column starts with $c = q = 0$ at φ_0.

The equations are discretized as 60 equal finite volumes with a van Leer limited advective flux, the Danckwerts flux at the inlet face and a zero gradient at the outlet. At the loading modifier the exchange relaxes in milliseconds while the band moves over minutes, so the exchange is solved in closed form per cell, with the competition frozen on a midpoint secant, and split Strang-wise around a fourth-order Runge-Kutta step of the transport (half an exchange, the transport, half an exchange). The step conserves each cell's inventory to round-off for any step length. The column outlet passes a holdup $V_h = 0.4\,\mathrm{L}$ of five well-mixed tanks to the detector and a holdup $V_c = 1.2\,\mathrm{L}$ of four tanks, plus any unmeasured extra ΔV_c, to the valve. Each destination integrates $\dot m_{i,d} = r_d\, Q\, c_{i,out}$ with $r_d \in \{0,1\}$ set by the rule, so every collected gram is a gram that left the column. The antibody capture train of Volume 8, Chapter 11 uses the same finite-volume column with a binding law of its own.

9.3.2 Conditioning, Aggregation and the Membrane

The transfer moves the pool and its phosphate into T-401 once, with a dilution buffer. The vessel integrates the strong-ion amount Z and the buffer total C_T, and the pH is the bracketed root of the charge balance in Equation (9.3), never an integrated state.

$$\begin{aligned}
&H+\frac{Z}{V_\ell}-\frac{K_w}{H}-\frac{C_T}{V_\ell}\frac{K_a}{K_a+H}=0\\
&\mathrm{pH}=-\log_{10}H+0.9\,f_{org}\\
&r_{agg}=k_{agg}\,\frac{V}{1+10^{(\mathrm{pH}-\mathrm{pH}_{agg})/0.5}}\left(\frac{c_P}{c_{ref}}\right)^2
\end{aligned} \tag{9.3}$$

where V_ℓ is the volume in liters, f_{org} the organic fraction and $0.9 f_{org}$ a declared apparent-pH offset. Acid and base enter as 1 M reagents at up to 0.05 L/min, and r_{agg} converts product into an insoluble aggregate G, half activated at $\mathrm{pH}_{agg} = 5.6$.

During concentration and diafiltration the vessel recirculates through the membrane, as shown in Equation (9.4).

$$\begin{aligned} J &= \left(\frac{1}{J_{\Delta P}} + \frac{1}{J_{lim}}\right)^{-1}, \qquad J_{\Delta P} = \frac{\mathrm{TMP}}{\mu (R_m + R_f)} \\ J_{lim} &= k_m \ln \frac{c_g}{c_{pep}}, \qquad \frac{\mathrm{d}R_f}{\mathrm{d}t} = k_f J c_{pep} \\ \frac{\mathrm{d}(V c_i)}{\mathrm{d}t} &= -A_m J s_i c_i - (\dot{a}_P + r_{agg})\, \delta_{iP} \\ \frac{\mathrm{d}V_k}{\mathrm{d}t} &= F_{d,k} - A_m J f_k \\ \frac{\mathrm{d}a_P}{\mathrm{d}t} &= \frac{1}{\tau_{ads}} \left(\frac{a_{max} A_m c_P}{K_{ads} + c_P} - a_P\right) \end{aligned} \tag{9.4}$$

where s_i is the observed sieving ($s_P = 0.010$, $s_{E,L} = 0.025$, $s_X = 0.95$, times the disturbance multiple), V_k the volume of solvent k (water, acetonitrile, 2-propanol, antisolvent) with fraction f_k, μ a logarithmic volume-fraction blend of the solvent viscosities, and a_P the peptide adsorbed on the membrane, kept as an inventory. The diafiltration buffer follows the permeate, $F_d = A_m J$, so the exchange runs at constant volume and N_D counts buffer volumes over the vessel volume. The soft minimum caps the flux at $J_{lim} = 8.25\,\mathrm{L/(m^2\,h)}$ at 100 g/L whatever the pressure.

9.3.3 Precipitation

Antisolvent lowers the solubility exponentially, residual acetonitrile raises it, and the solid is born and grown as described by extensive moments, as shown in Equation (9.5). The batch crystallizer of Volume 8, Chapter 4 carries its crystal population in the same leading moments.

$$\begin{aligned} c^* &= c_0^* \exp(-\kappa f_s + \kappa_a f_{a,rel}), \qquad S = c_P / c^* \\ B &= k_b (S-1)^b, \qquad G = \frac{k_g (S-1)^g}{1 + K_{inh} x_{L,liq}} \\ \frac{\mathrm{d}M_0}{\mathrm{d}t} &= VB, \qquad \frac{\mathrm{d}M_k}{\mathrm{d}t} = kGM_{k-1} + VBL_0^k, \quad k = 1,2,3 \\ \dot{m}_{solid} &= \rho_s k_v \frac{\mathrm{d}M_3}{\mathrm{d}t}, \qquad L_{32} = M_3 / M_2 \end{aligned} \tag{9.5}$$

where $f_{a,rel}$ is acetonitrile as a fraction of the non-antisolvent liquor, $x_{L,liq} = c_L/(c_P + c_E + c_L)$ is the late variant's share of the dissolved peptide and L_{32} is the Sauter mean size. The new solid is exactly $\dot{m}_{solid}$, split among the species in proportion to $c_P - c^*$, $1.35 c_L$, $0.60 c_E$ and $50 c_G$, so the variants enter by incorporation and every aggregate reports to the solid. The growth inhibition $K_{inh} = 700$ is the coupling of this chapter, a teaching value chosen so that the lesson appears at a realistic cut and not fitted to any data. Its form is a simple saturating inhibition in the spirit of impurity-adsorption models of growth [3].

9.3.4 The Nutsche Filter and the Discharge

The cake obeys Ruth's constant-pressure relation [15] with a specific resistance that follows the Kozeny scaling on particle size, a compressibility law, and a gel contribution of the late variant and the aggregate, as shown in Equation (9.6).

$$\begin{aligned} \alpha &= \alpha_{ref} \left(\frac{L_{ref}}{L_{32}}\right)^2 \left(\frac{\Delta P}{\Delta P_{ref}}\right)^s (1 + \gamma_L x_L + \gamma_G x_G) \\ Q_f &= \frac{A_f \Delta P}{\mu\,(R_{cl} + \alpha M_c / A_f)}, \qquad L_c = \frac{M_c}{\rho_s (1 - \varepsilon_c) A_f} \end{aligned} \tag{9.6}$$

Table 9.2 Parameters and Their Evidence. S Anchored in a Read Source, T Teaching Value, F Fitted Inside the App, N Numerics

Symbol	Meaning	Value	Unit	Ev.
L, D, ε	bed length, diameter, mobile fraction	0.25, 0.20, 0.60	m, —	T
q_{max}	saturation capacity	50	kg/m^3	T
b_P, b_E, b_L	affinities at $\varphi_{ref} = 0.30$	0.30, 0.16, 0.52	m^3/kg	T
S_P, S_E, S_L	modifier sensitivities	60, 50, 72	—	T
k, D_{ax}	exchange rate, dispersion	0.05, 3×10^{-7}	s^{-1}, m^2/s	T
V_h, V_c	holdups to detector, collector	0.4, 1.2	L	T
x_c	crude purity	70	%	S
C_T, pH	phosphate, mobile-phase pH	20, 7.5	mM, —	S
φ_0, φ_1, t_g	gradient	22, 45, 60	%, min	T
k_{agg}, pH$_{agg}$	aggregation at 100 g/L, half point	2.6×10^{-3}, 5.6	kg/(m^3 s)	T
A_m, R_m	membrane area, resistance	1.0, 1.5×10^{13}	m^2, m^{-1}	T
k_m, c_g	polarization coefficient, gel	2.5×10^{-6}, 250	m/s, g/L	T
s_P, $s_{E,L}$, s_X	observed sieving	0.010, 0.025, 0.95	—	T
a_{max}, τ_{ads}, K_{ads}	membrane adsorption	0.35, 300, 5	g/m^2, s, g/L	T
k_f	fouling per peptide carried	2×10^{14}	m/kg	T
TMP, N_D	pressure, diavolumes	2, 4	bar, —	S
c_0^*, κ, κ_a	solubility, antisolvent, acetonitrile	150, 9.5, 5.0	g/L, —	T
k_b, b	nucleation	2×10^{11}, 2.4	m^{-3} s^{-1}	T
k_g, g	growth	1.6×10^{-7}, 1.3	m/s	T
K_{inh}	growth inhibition by L	700	—	T
ρ_s, k_v, L_0	solid density, shape, birth size	1150, 0.52, 1	kg/m^3, —, µm	T
α_{ref}, s	cake resistance (1 bar, 20 µm), compressibility	2×10^{12}, 0.60	m/kg	T
γ_L, γ_G	gel weights of L and G	200, 55	—	T
A_f, ε_c, R_{cl}	filter area, porosity, cloth	0.05, 0.60, 10^{11}	m^2, —, m^{-1}	T
f_{occ}, τ_{dew}, k_{dis}	occlusion, drying, redissolution	0.015, 240, 7×10^{-5}	—, s, m/s	T
$\hat{f}_L$ constants	φ_{L50}, w_L, $f_{L,max}$	0.3303, 0.00329, 0.785	—	F
Δt, cells	integration step, finite volumes	1, 60	s, —	N
limits	purity, acetonitrile, solvent, clock	98.5, 0.010, 2.5, 60	%, min	T

where x_L and x_G are mass fractions in the cake, R_{cl} the cloth and L_c the cake depth, 2.46 mm for the 56.6 g cake of the reference case. In cake-dominated flow $Q_f \propto \Delta P^{1-s}$, so with $s = 0.60$ more pressure gives diminishing returns and never less flow. The filtrate carries the suspended solids into the cake, which takes up liquor at $\varepsilon_c/(1-\varepsilon_c)$ of its solid volume. The wash is one well-mixed stage that displaces the pore liquor and redissolves cake at $k_{dis} a_s (c^* - c)$ over the solid surface $a_s = 6M_c/(\rho_s L_{32})$. The drying relaxes the pore liquor toward 1 % of the pore volume at the rate $\lambda = \max(0.1, \Delta P/1\,\text{bar})/\tau_{dew}$. A share $f_{occ}(1 + 3(x_L + 2x_G))$ of the solid volume is occluded mother liquor that no wash reaches, so the residual acetonitrile of a washed cake is set by the diafiltration. Filtration, wash and drying share one 60 min clock. At the discharge the solid is acceptable only if its peptide purity is at least 98.5 %, its acetonitrile at most 0.010 %, its total solvent at most 2.5 % and the clock did not run out. These limits are teaching values. Real ones come from the product's specification, a pharmacopeia and the regulator's assessment, and passing them here certifies no batch. The scorecard leads with the acceptable mass and its rate per campaign hour, both read at the discharge.

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

9.3.5 Parameters, Evidence and Instruments

Every parameter appears in Table 9.2 with its evidence. Four are anchored in read sources. The crude purity of 70 % sits inside the 60 to 70 % of [12], and the transmembrane pressure and diavolumes inside the 2 to 9 bar and 4 to 8 diavolumes of [13]. Everything else is a teaching value, and the membrane, precipitation and cake values carry no data at all. No controller sees the plant except through an instrument. FT-201 reads the flow through a 3 s lag with 0.5 % colored noise, AT-202 the inlet modifier through 6 s, and AT-203 the absorbance $a_{UV} = \ell \sum_i \varepsilon_i c_i + a_0$ of the detector cell through 2 s with 0.004 AU of noise, a mixed signal in which product and variants absorb nearly alike

(1.00, 0.98 and 1.06 AU per g/L per cm). An inline refractive index gives the modifier at the detector. AT-401 reads the pH through 25 s with 0.01 of noise, PT-501 and PT-601 the pressures with 0.02 bar of noise, and AT-503 the permeate's absorbance. The pool assay draws the pool once, at the transfer, and reports its purity and late-variant fraction 15 min later with standard deviations of 0.06 and 0.03 points. The time step is 1 s, the controllers run every 5 s, one row is logged every 20 s, and the seed is 20261020.

9.3.6 Assumptions

The crude is a mass and a composition, with three impurity classes and one downstream aggregate standing in for a related-substance profile that in practice has dozens of members. The isotherm is a candidate effective law with fixed temperature, pH and ion pairing, and one preparative pass replaces the two or three of the patents, so the trade-offs appear at a lower purity than a multipass route reaches. Solvent volumes are additive, the apparent pH of organic-rich liquor is an offset, and the moments assume size-independent growth with no agglomeration, breakage or dissolution. The cake is washed as one mixed stage, which understates a real displacement wash, and its resistance is evaluated on its current composition, so the wash lowers the resistance of a formed cake by 16 % in the reference case. The scorecard therefore reports the value at the end of the filtration. The precipitation vessel and the conditioning vessel share one inventory in the model although the scene draws R-701 apart.

9.4 The Recipe, the Loops and the Three Cuts

A supervisor runs thirteen stages, load, column wash, gradient, strip, conditioning, concentration, diafiltration, final concentration, precipitation, filtration, cake wash, dewatering and discharge. Each ends on a measurement or a time, a loaded volume, 1.5 column volumes of wash, the gradient time, 6 min of strip at $\varphi_1 + 0.15$, the pH held within 0.12 of setpoint for 2 min after at least 3 min, a vessel volume of 3 L, the diavolume count, a peptide concentration from the soft sensor's inventory, the antisolvent fraction plus a 10 min hold, and the filter volumes. Each transition fires once, after the integration of the step in which its condition is met. In the reference case the stages take 4.7, 60.0, 6.0, 5.3, 44.8, 114.6, 37.1, 24.5, 15.3, 5.2 and 20.2 min from the wash to the drying, a campaign of 5.66 h.

The five loops are the velocity-form PID of the framework, Equations (1.11) to (1.14) of Section 1.6.2, with the derivative off, executed every $\Delta t_c = 5\,\mathrm{s}$. The bias u_b is the output the loop is initialized at, and the anti-windup is back-calculation on the integral with a tracking time $T_t = \tau_I$. FIC-201 drives the pump (3 L/min at full speed, 8 s lag) at $K_c = 18\,\%$ per L/min and $\tau_I = 25\,\mathrm{s}$. AIC-202 drives the B valve on the gradient program at 2.2 % per percent and 30 s, and idles with setpoint tracking when the column stops. AIC-401 drives a split range with a 0.04 deadband at 70 % per pH and 150 s, dosing only during conditioning. PIC-501 sets a retentate valve whose demand is $8\,\mathrm{bar}\,(u/100)^{1.2}$ through a 20 s lag at 22 % per bar and 70 s, and PIC-601 a gas regulator of 5 bar span through 15 s at 35 % per bar and 70 s. Every one has a positive gain from its output to its measurement, and each tracks its operator station in Manual, so transfers are bumpless (Section 1.6.2).

The cut is the learner's decision, and three rules make it. The time rule collects while $t_1 \le t_{el} \le t_2$ on the gradient clock. The UV rule collects while $a_{UV} \ge 0.60\,\mathrm{AU}$ and, once it has shut, stays shut. Its stop threshold of 0.35 AU has no effect as implemented, because the opening threshold is tested at every scan (a defect of the app as shipped). The downstream-aware rule reads a soft sensor that turns the UV into a peptide concentration with one average absorptivity and splits it by a calibration against the modifier at the detector, as shown in Equation (9.7).

$$\begin{aligned} \hat{f}_L &= \frac{f_{L,max}}{1+\mathrm{e}^{-(\varphi-\varphi_{L50})/w_L}}, \qquad \hat{f}_E = \frac{1}{1+\mathrm{e}^{(\varphi-\varphi_{E50})/w_E}} \\ \hat{x}_L &= 100\,\frac{\int Q\hat{c}\,\hat{f}_L\,\mathrm{d}t}{\int Q\hat{c}\,\mathrm{d}t}, \qquad \hat{c} = \frac{a_{UV}-a_0}{\ell\bar{\varepsilon}} \end{aligned} \tag{9.7}$$

Table 9.3 Code Verification and the Checks a Source Allows

Check	Result
Code verification	
product ledger, charge to cake, 45 campaigns	$\le 2 \times 10^{-11}$ relative
solid against $\rho_s k_v M_3$	$\le 3.7 \times 10^{-15}$
column with integrated fluxes, 30–480 cells	4×10^{-15}–10^{-14}
tracer mean against V_{liq}/Q	−0.000 %
band exits against linear solvent strength	0.033, 0.046, 0.034 min
linear pulse, mean and spread at 60 cells	0.12, 2.2 %
diafiltration $c/c_0 = \mathrm{e}^{-sN_D}$, $s_P = 0.010$	0.43 %
$s = 0$ retention	no product passes
Ruth filtration time, reference cake	14.2 against 15.3 min[a]
drying time by formula and engine	20.2, 20.2 min
grid, 60 against 480 cells at 0.25 s	0.12 %, 0.04 points
Δt 1 s against 0.25 s	0.12 %, 0.13 points
Δt halved, challenge (suggested)	59.91 and 59.89 g
playback at 1800 times real time	bit-identical
Against read sources (consistency only)	
91 % net yield over 10 diavolumes [11]	$s_P \le 0.0094$, model 0.010
crude purity 60–70 % [12]	model 70 %
TMP 2–9 bar, 4–8 diavolumes [13]	model 2 bar, 4
cake resistance of a peptide precipitate	no source

[a] The rest is PIC-601's pressure ramp, 1.22 bar after 1 min.

where both signals are carried forward by the declared V_c in volume, so the sensor describes what is at the valve. The supervisor of Volume 5, Chapter 7 follows material in volume in the same way. It opens when the valve stream is above 0.60 AU and its predicted purity $100(1 - \hat{f}_E - \hat{f}_L)$ reaches 80 %, and it shuts for good when $\hat{x}_L$ reaches $\hat{x}_{L,max}$ or either opening condition fails. The calibration constants were fitted inside the app to seven runs of this column with crudes of 65 to 75 %, with a weighted root-mean-square error of 6.23 points of mass percent on those runs. No loop or rule reads a true composition. Four stage ends read a model volume standing in for a level, a totalizer, a sight glass and a dew point, and the application keeps three models apart: the plant of Section 9.3, the purity and recovery curve that trends let a learner draw, and the soft sensor, which knows only the calibration and the declared holdup.

9.5 Code Verification, and What the Sources Can Check

Code verification, agreement with sources and educational evaluation are kept apart. Whether the program solves the equations it states is the subject of Table 9.3, and every entry passes. The column's inventory, with its inlet and outlet fluxes integrated in the same stages, closes to 10^{-14} at 30 to 480 cells. A tracer leaves after $V_{liq}/Q = 188.5$ s to round-off, and at vanishing load each band leaves the gradient within 0.05 min of linear-solvent-strength theory [10]. With the step held at 0.25 s, 60 cells put the outlet peak within 0.12 % and the pooled recovery within 0.04 points of a 480-cell reference, while first-order upwinding is 1.58 % and 0.91 points out. The shipped 1 s step adds 0.12 % and 0.13 points. Halving it over a whole campaign moves the isolated mass of case 2 by 0.033 %, and no verdict among 31 lesson variants changes at 120 cells, at half the step or at 30 cells.

Agreement with sources is thin, and the table says how thin. A net yield of 91 % over ten water diavolumes in [11] would correspond to an observed product sieving of $-\ln(0.91)/10 = 0.0094$ if every loss were sieving, which is consistent with the declared 0.010 without fixing it. The crude purity and the membrane operating window lie inside the disclosed ranges. No source consulted for this chapter measures the retention of these peptides on this column, a membrane's sieving for this solvent, or the filterability of a peptide precipitate, so the isotherm, the growth inhibition and the cake are teaching closures. Validation runs of the load, the gradient and the filtration against withheld data cannot be made without such data. Educational evaluation is not attempted, and no learning gain is claimed.

Table 9.4 The Nine Cases and the Challenge

#	Case and departure	Teaching point
1	crude at 60 %	a cut belongs to its batch
2	reference, cut 23–31 min	purity against recovery
3	200 g charged	capacity and displacement
4	collector holdup +1.2 L, unmeasured	volume is delay
5	conditioning at pH 4.5	aggregate the filter feels
6	membrane sieving ten times specified	each diavolume costs
7	antisolvent at 0.40 L/min, cut to 31.25 min	fines
8	filter pressure and wash	compressible cake
9	cut shut at 31.75 min, UV rule infeasible	the pool that passes
C	cut to 32, 2 diavolumes, 0.25 L/min, 3 bar	the whole campaign

Table 9.5 Measured Outcomes of the Cases (Default Seed, at the Discharge)

#	Measured
1	23–31: pool 96.27 %, solid 96.88 %, 53.0 g rejected. Re-cut 24–30.5: pool 98.47 %, 45.3 g acceptable. Downstream aware: pool 94.43 %, rejected.
2	24–30: 47.7 g pooled, 99.48 %, 29.7 min on the filter, 43.4 g, 9.1 g/h. 23–31: 61.7 g, 98.92 %, 11.6 µm, $\alpha = 12.6 \times 10^{12}$ m/kg, 40.7 min, 56.2 g, 9.9 g/h. 23–32: 98.11 %, 1.235 % late, 8.3 µm, out of time filtering.
3	200 g: recovery 42.7 % (73.4 at 120 g), 2.79 % late, $\alpha = 1.6 \times 10^{14}$ m/kg, out of time filtering.
4	pool 97.84 %, solid 98.26 %, rejected. Both cut times +0.8 min: 55.6 g acceptable.
5	pH 4.5: 2.21 g aggregate, solid 95.41 %, 66.4 min. pH 6.0: 0.25 g, acceptable.
6	33.9 g through the membrane, 27.0 g isolated. 2 diavolumes: 0.0127 % acetonitrile, rejected. 8: 17.8 g.
7	0.40 L/min: 8.1 µm, 68.3 min, rejected. 0.15, 0.06, 0.02: 9.3, 10.7, 12.5 µm in 56.2, 47.8, 41.5 min.
8	3.5 against 1.5 bar: 24.8 against 40.7 min. 0.6 bar: drying unfinished, 3.11 % solvent.
9	pool 98.39 % (assay 98.35), 0.939 % late, 9.0 µm, $\alpha = 3.4 \times 10^{13}$ m/kg, out of time drying, solid 98.83 %. Rescued by $t_2 = 31$ or 3 bar. UV rule: 0 of 24 isolation settings rescue it.
C	shipped: 53.5 g at 98.22 %, 0.0273 % acetonitrile, out of time. Suggested: 59.9 g, 10.0 g/h.

9.6 Nine Cases, a Challenge and Three Cut Policies

The cases are listed in Table 9.4. The figures quoted here come from the app's own engine, run without a display at seed 20261020 with the browser's step sequence and with each setting typed in before the clock starts, and every case loads identically when opened first, chosen after another case or rerun from Reset. The outcomes of the cases and their main variants are listed in Table 9.5. Credit for a case requires both a change the learner made and the plant state it produced, and all 11 progress checks pass the automated walk-through.

9.6.1 The Column's Own Trade

The product band is above 0.1 g/L at the valve from 20.0 to 36.0 min of the gradient, with a flat top near 5.5 g/L from 25.9 to 28.5 min, the early variant peaking at 20.4 min on its front and the late variant at 33.9 min on its tail. That ordering is the overloaded displacement of Equation (9.2), and it is why the cut is a decision. In case 3 a 200 g charge exceeds the bed's 157 g and breaks through during the load, the recovery falls from 73.4 to 42.7 % and the pool carries 2.79 % of the late variant. In case 4 an unmeasured 1.2 L extension of the collector line delays the valve by 48 s at 1.5 L/min, and the same cut collects the front shoulder, so the pool falls from 98.92 to 97.84 %.

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

Table 9.6 Cut Policies on the Challenge Case With Equal Downstream Settings (4 Diavolumes, 0.06 L/min, 1.5 bar, 0.30 L Wash), Default Seed

Policy	Window (min)	Rec. (%)	Pool (%)	Late (%)	Memb. (g)	Size (μm)	α (10^{12} m/kg)	Filter (min)	Isol. (g)	Acc. (g)	Rate (g/h)	Solv. (L/g)
purity-maximizing	24.05–30.03	56.8	99.48	0.086	3.70	14.4	5.7	30.7	43.4	43.4	9.02	2.91
reference	23.05–31.03	73.4	98.92	0.358	4.67	11.6	12.6	41.8	56.2	56.2	9.89	2.25
recovery-maximizing at the pool limit	23.05–31.53	77.0	98.62	0.698	4.88	9.8	23.9	62.0[a]	59.0	0	0	—
UV threshold	20.30–32.78	95.0	91.64	2.069	5.96	7.3	91.2	80.2[b]	38.9	0	0	—
downstream aware, $\hat{x}_{L,max} = 0.80\,\%$	22.55–31.20	78.2	98.51	0.430	4.95	11.2	14.4	46.8	59.9	59.9	10.03	2.11
fixed at that window	22.63–31.28	78.3	98.53	0.485	4.95	10.9	16.3	49.4	60.0	60.0	9.96	2.11
best fixed window in the sweep	22.55–31.03	77.0	98.58	0.340	4.88	11.9	11.8	42.4	58.9	58.9	10.08	2.14

Rec.: recovery into the pool. Isol.: isolated. Acc.: acceptable. Window: the collection times realized on the gradient clock (the rule scans every 5 s). Memb.: product through the membrane plus product adsorbed on it. Solv.: all water and organic solvent used per gram of acceptable peptide, 80.9 L of water and 45.5 L of organic solvent per campaign (37.0 L of it acetonitrile in the mobile phase) for every policy here. [a]Out of filter time while drying. [b]Out of filter time while filtering.

> **TRY IT IN THE APP** **4 · A Longer Line to the Collector**
>
> Run the scenario and compare the UV trace with the shaded cut window on the Column view. The pool falls from 98.92 to 97.84 % and the solid is rejected. Press Reset, move both Cut opens and Cut shuts 0.8 min later on the XV-203 card, and run again: 55.6 g is now acceptable.

9.6.2 Purity, Recovery and the Filter

The comparison of cut policies is run on the challenge case with every policy given the same downstream recipe, four diavolumes, 0.06 L/min of antisolvent, 1.5 bar and a 0.30 L wash, as listed in Table 9.6. A purity-maximizing cut (24 to 30 min) pools 99.48 % and isolates every gram it pools, but only 56.8 % of the product reaches the pool, and the rate is 9.02 g/h. A recovery-maximizing cut that holds the pool at the solid's 98.5 % limit shuts at 31.5 min, pools 98.62 % at 77.0 % recovery and is lost on the filter clock at 62.0 min, still drying. The UV rule recovers 95.0 % into a 91.64 % pool and never finishes filtering. The downstream-aware rule at a 0.80 % limit opens at 22.55 min, shuts at 31.20 min and isolates 59.9 g at 10.03 g/h. Figure 9.2 draws the recovery and the acceptable mass of every policy side by side.

The comparison over the whole cut grid is shown in Figure 9.3. Of 44 windows from $t_1 = 22.5$ to 24 min and $t_2 = 29$ to 33 min, the 20 that failed all ran out of filter time, and the dividing line is the pool's late variant, at most 0.568 % in the 24 that passed and at least 0.665 % in those that failed. A failing pool reached a purity of 98.86 % while a passing one held 98.46 %, so the boundary in the purity and recovery plane is not a purity contour. The last acceptable shutting time is 31.25 min for every opening time tried, and the acceptable mass rises with recovery up to that edge and drops to zero past it, as shown in Figure 9.3b. Membrane loss rises with the pooled mass, from 3.70 to 5.96 g, because sieving takes a fixed share of what the membrane sees. Solvent use is nearly the same for every policy, about 126 L per campaign, because the column program and the diafiltration volume do not depend on the cut, so solvent per gram follows the acceptable mass, from 2.91 to 2.11 L/g.

Pause. Try Exercise 9.1 before reading on.

> **TRY IT IN THE APP** **9 · The Wide Pool That Passes Its Assay**
>
> Load the scenario, press Run, set the speed to 1800× and follow the campaign to the Filter view. The pool assays at 98.35 %, inside the limit. Watch the filter clock: the cake needs 79.3 min of the 60 it is allowed. Press Reset, set Cut shuts to 31 min and run again, and the batch is saved.

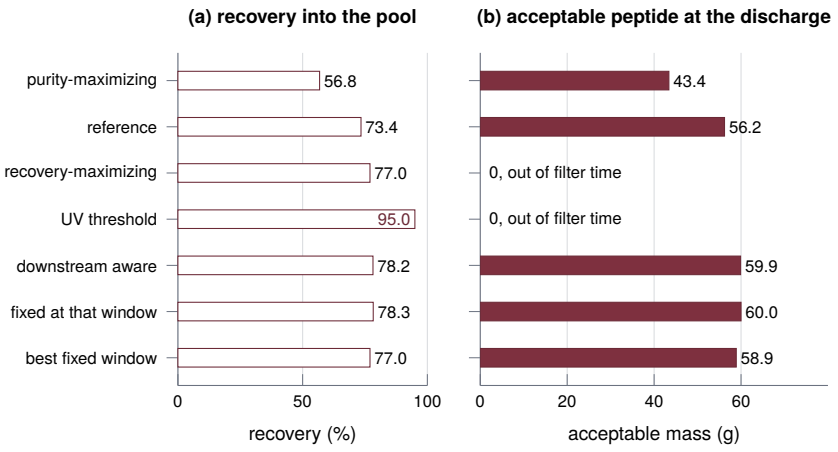

Figure 9.2 The cut policies of Table 9.6 on the challenge case with equal downstream settings: (a) recovery into the pool, (b) acceptable peptide at the discharge. Notice that the UV threshold rule recovers the most and isolates nothing acceptable, like the cut that holds the pool at its purity limit, and that the two fixed windows match the downstream-aware rule.

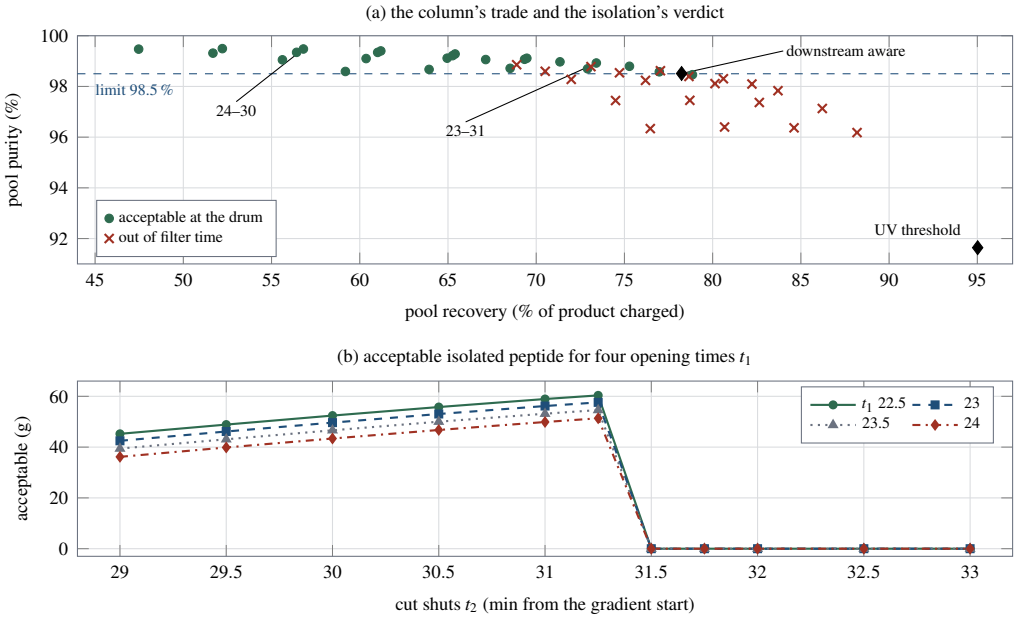

Figure 9.3 Cut sweep on the challenge case with the suggested downstream settings. (a) Pool purity against recovery for 44 windows, marked by the isolation's verdict, with the UV and downstream-aware rules. (b) Acceptable isolated peptide against the shutting time for four opening times, from the earliest to the latest solid with dots, dashed with squares, dotted with triangles and dash-dotted with diamonds. Every window that fails runs out of filter time, and none shutting after 31.25 min survives.

9.6.3 What the Soft Sensor Contributes

The downstream-aware rule's margin over the reference cut, 59.9 g against 56.2 g, comes from its window rather than from what it knows about the pool. A fixed cut set to the window it realized, 22.55

Table 9.7 The Soft Sensor Against the Crude (Reference Case, Cut 23–31 min), and the Downstream-Aware Rule on the Same Crude

Crude (%)	True pool purity	True pool late	Predicted purity	Predicted late	Ratio late	Rule, acceptable (g)
60	96.27	0.484	97.26	0.786	1.62	0, pool 94.43 %
65	98.11	0.416	97.21	0.778	1.87	0, pool 97.14 %
70	98.92	0.358	97.19	0.776	2.17	59.9
75	99.28	0.304	97.18	0.775	2.55	60.9
80	99.47	0.251	97.18	0.776	3.10	61.6

to 31.20 min, isolates 60.0 g at 9.96 g/h, and the best fixed window of the sweep, 22.5 to 31 min, reaches 10.08 g/h, slightly above the rule's 10.03 g/h. The reason is the calibration. The soft sensor's prediction is a function of the modifier at the valve and of the UV, and it hardly moves with the crude, as listed in Table 9.7. On the 23 to 31 min cut it predicts 0.78 % of the late variant for every crude from 60 to 80 %, while the truth falls from 0.484 to 0.251 %, and the rule's window stays at 22.55 to 31.20 min for crudes from 65 to 80 %. Over six paired seeds the suggested configuration isolates 59.9 to 64.3 g and passes the challenge every time, the fixed 23 to 31 min cut 55.9 to 56.7 g and never, and a fixed 23 to 31.5 min cut filtered at 3 bar 58.6 to 59.5 g and every time. Over the 31 noise seeds of the final validation the suggested configuration passes 18 (19 at half the step). At six seeds the rule's cut shuts within half a minute of opening and is not reopened (the stop-threshold defect of Section 9.4), at five the filtration, wash and drying need 62.7 to 68.9 min, and at two the solid is 98.06 to 98.39 % pure. The rule's advantage disappears on a weaker crude. At 65 % it pools 97.14 % and is rejected while the fixed cut isolates 55.0 g.

> ✦ DID YOU KNOW? **A Soft Sensor That Does Not See the Crude**
>
> From the simulator (Table 9.7). On the reference cut the soft sensor predicts between 0.775 and 0.786 % of the late variant for every crude from 60 to 80 %, while the true pool holds from 0.484 down to 0.251 %. Its estimate follows the modifier at the valve, not the crude.

9.6.4 Sensitivity to Selectivity, Sieving and Cake Resistance

The plant was then perturbed while the soft sensor kept its calibration, as listed in Table 9.8. Selectivity matters most. With the late variant's affinity 10 % lower (b_L/b_P from 1.73 to 1.56) its band moves into the product, the reference cut pools 0.96 % of it and every policy but the purity-maximizing cut runs out of filter time, while the narrow cut still isolates 43.5 g. With 10 % more affinity every policy gains filter margin and the ranking of the declared plant holds. Sieving costs mass in proportion. Five times the specified sieving takes the downstream-aware rule from 59.9 to 44.0 g with 21.1 g lost to the membrane, and the ranking is unchanged. A higher cake resistance costs time and not mass until the clock binds. At 1.5 times α_{ref} the rule's cake takes 59.0 min of the 60, the closest of the four to failing, because its pool carries the most late variant.

> ✦ DID YOU KNOW? **A Weaker Selectivity Reverses the Ranking**
>
> From the simulator (Table 9.8). On the declared plant the narrow 24 to 30 min cut isolates the least acceptable peptide, 43.4 g. With the late variant's affinity 10 % lower, the three wider policies all run out of filter time, and that narrow cut, at 43.5 g, is the only one left.

9.6.5 The Isolation Cases

Cases 5 to 8 break the isolation one lever at a time. At pH 4.5 the loop overshoots to pH 3.0 on the way down, 2.21 g of aggregate forms, mostly during the diafiltration, and the solid falls to 95.41 %

Table 9.8 Acceptable Peptide (g) With the Plant Perturbed and the Soft Sensor's Calibration Kept. Filter: Out of Filter Time

Plant	24–30	23–31	22.5–31	Downstream
declared	43.4	56.2	58.9	59.9
b_L −10 %	43.5	filter	filter	filter
b_L +10 %	43.3	56.0	58.7	59.7
sieving ×2	40.2	52.1	54.7	55.6
sieving ×5	31.8	41.2	43.3	44.0
α_{ref} ×0.8	43.4	56.2	58.9	59.9
α_{ref} ×1.25	43.4	56.2	58.9	59.9
α_{ref} ×1.5	43.4	56.2	58.9	59.9[a]

[a]59.0 min on the filter, against 50.9 for 22.5–31.

because the aggregate is insoluble. A membrane sieving ten times its specification passes 33.9 g, and the permeate's own UV cell shows it within minutes. Dosing the antisolvent at 0.40 L/min instead of 0.06 makes 8.1 µm particles from the same 57.6 g, and the cake needs 68.3 min. On the filter, 2.33 times the pressure is 1.64 times faster, because the cake's resistance rises from 1.3×10^{13} to 2.1×10^{13} m/kg as it compresses. The challenge as shipped fails twice, on the clock and on 0.0273 % acetonitrile left by two diavolumes in the occluded liquor. A headless campaign of 5.66 h takes 0.82 s on the build host, 24 719 steps per second, and in a browser a campaign plays in 11 to 15 s at 1800 times real time. The shoulder experiment of Figure 9.4, the reference cut and then the cut at 32 min, took 27.6 s of wall time at 1800 times real time in a usability review.

9.7 Technology in Practice

What Measures It

The cut is decided on one optical signal. AT-203 reads the absorbance of the stream in a UV cell at the column outlet. In the model it carries a 2 s lag and 0.004 AU of noise, and it sees the product and its variants nearly alike. Real purification meets the same limit, because one optical method alone offers limited selectivity between a product and contaminants of closely similar structure [18]. The variable of each loop comes from a common transmitter: a flowmeter, an inline analyzer of the modifier, a pH electrode and two pressure transmitters, each with the lag and noise of Section 9.3. Composition beyond the UV signal comes only from the pool assay, 15 min after the transfer, the role that analytical chromatography of typically 30 min or more plays in practice [18].

> **TECHNOLOGY Sensor: The UV Absorbance Detector**
>
> The most popular detectors for liquid chromatography use the absorption of ultraviolet or visible light by the analytes. The stream passes a flow cell, which in an analytical instrument holds 1 to 10 µL over a path of 0.2 to 1 cm, and the chromatogram is absorbance against elution time. The mobile phase must not absorb at the chosen wavelength, and the choice of wavelength changes each analyte's signal. Gas released as the pressure falls across the column can form bubbles that interfere with the detector's response [19]. One optical method alone also has limited selectivity. In protein purification, contaminants with subtle structural differences from the product are found by analytical chromatography that typically takes 30 min or more [18].

What Moves It

The pump of FIC-201 delivers 3 L/min at full speed through an 8 s lag, and AIC-202 sets the share of solvent B with a valve, the arrangement of a proportioning valve ahead of a single pump [19]. The flow measurement carries 0.5 % colored noise and no stroke pulses. The cut valve XV-203 switches at once in the model, but the line between the detector and the valve holds 1.2 L, and case 4 shows that 1.2 L more delays the cut by 48 s at 1.5 L/min. Figure 9.5 sets the loops on such hardware.



Figure 9.4 The shoulder collected, case 2 with the cut shut at 32 min, on the Column view at $t = 2443$ s (40.7 min), 34.0 min into the gradient. The figure is a composite of the app's header and its Column view, with the control panel, the trends and the Learn panel cut. Callouts: (1) the soft sensor's badges, (2) the UV cell AT-203, (3) the cut valve XV-203, (4) the pool with the soft sensor's estimate, (5) the shaded collection. The shaded collection from 23.0 to 32.0 min is complete and shut, 2.3 g remain on the bed, and the late-variant band peaks at the valve, about 1.4 g/L against under 1 g/L of product on the app's trend panel. The pool holds 13.5 L and 68.9 g by the soft sensor, which predicts 96.4 % purity and 1.78 % late variant against a true 1.24 % that the trend panel shows, and 80.1 % of the charged product is in the pool. The status reads ELUTING TO WASTE at 0.40 AU and 33.8 % organic. This pool will assay at 98.07 % and run out of filter time.

TECHNOLOGY **Final Element: The Gradient Pump**

A liquid chromatograph draws its mobile phase from one or more reservoirs. In one design two piston pumps feed a mixing chamber: their combined speed sets the flow and their relative speed the composition. In another, a solvent proportioning valve ahead of a single pump sets how much of each solvent is drawn in each pump cycle. The back-and-forth stroke of a reciprocating pump gives a pulsed flow that adds noise to the chromatogram. A second cylinder that delivers while the first refills, or a pulse damper with a compressible fluid behind a flexible diaphragm, smooths the flow [19].

Where the Controller Runs

In a plant, loops like these would run in a distributed control system, whose supervisory level oversees localized controllers. Programmable logic controllers often serve as those local controllers, and the loops run with cycle times from milliseconds to minutes [20]. The app executes the five PI loops every 5 s on a 1 s integration step. Above them sits the plantwide structure: the supervisor of thirteen stages, each ended on a measurement or a time, and the cut rule, the one decision that couples the column to the filter. ICH Q11 lists in-process controls and controls implicit in the design of the process, such as the sequence of purification steps, among the parts of a control strategy [21], the two parts that the loops and the stage sequence play in the model.

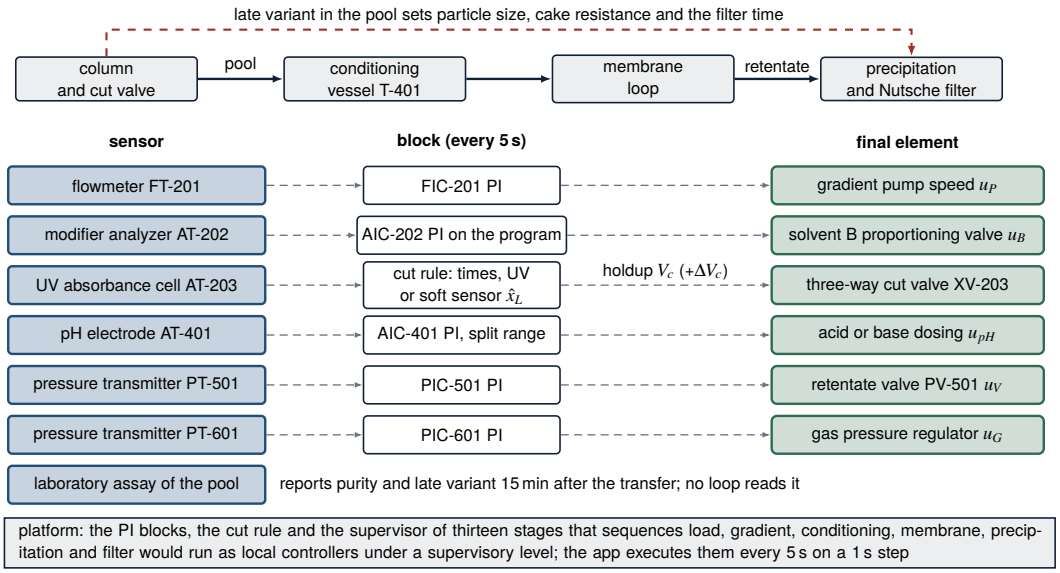

Figure 9.5 The loops of Section 9.4 on real hardware, sensor to block to final element, below the four blocks of the train. Notice that no loop closes from the filter back to the cut: the coupling drawn as the dashed line across the top is left to the cut rule and its holdup, and the laboratory assay arrives after the cut is over.

> **IN PRACTICE** **One Design Space Across Several Unit Operations**
>
> The international guideline on the development and manufacture of drug substances, ICH Q11, asks that the formation, fate and purge of impurities be understood, and that controls be set for impurities as they progress through multiple process operations. A design space may be set for one unit operation or for a combination of selected unit operations. These are chosen for their impact on the critical quality attributes and need not be sequential, and the linkages between steps are evaluated so that the cumulative generation and removal of impurities is controlled. A design space that spans multiple unit operations can give more operational flexibility [21].

What Is Modernizing It

The first development is in-line spectroscopy with multivariate calibration in place of a late laboratory result. In the purification of protein drugs, spectroscopic methods detect rapidly and automatically with no sample preparation, conditioning or destruction, and several sensors combined measure better than one. Where sensors sit at different points of a stream, the dispersion between them must be accounted for [18], which is the lesson of case 4 in another form. The second is model predictive control. Once a fast dynamic model of the critical quality attributes exists, it can be used for control, and a predictive controller respects constraints and can tailor production toward real-time release [18]. In this chapter nothing optimizes the cut, the diavolumes or the dosing, and a predictive or batch-to-batch policy is not implemented (Section 9.8).

9.8 Limitations

The route is a teaching assembly of disclosed options, and patents establish that the options exist, not that the train works. The isotherm, the growth inhibition, the membrane and every cake coefficient are teaching values, and the coupling of this chapter rests on K_{inh}, a number chosen so the lesson appears at a realistic cut. The soft sensor is a modifier calibration biased high by a factor of 1.6 to 3.1 on the late variant and blind to the crude, so the downstream-aware rule acts as a fixed window in

disguise, and the best fixed window found with hindsight matches it. The pool assay arrives after the cut is over, so it cannot correct the cut it describes, unlike the laboratory result that corrects the soft sensor of Volume 5, Chapter 10. The weaker crude of case 1 arrives 10 points below the calibration in one step, which is not how a synthesis drifts, but it isolates the calibration's blind spot. Nothing optimizes the cut, the diavolumes or the dosing, and a model predictive or batch-to-batch policy is not implemented. The precipitation has no agglomeration, the wash is one mixed stage, and the drying is a relaxation. Every limit and the 60 min clock are teaching values.

OPEN CHALLENGE Telling the Product From Its Variants on Line

Ideally a process measurement would tell the product, process-related and product-related contaminants apart in real time. For protein drugs this is hard: the contaminants are chemically close to the product, and one optical method alone has limited selectivity. Several sensors combined by multivariate data analysis do better, but a calibration model must be tested on an external data set before productive use, because internal validation is not sufficient. In-line soft sensors have reached the needed detection limits for aggregates and fragments at laboratory scale, and future work may close the loop on the steps of interest [18]. The soft sensor of this chapter, fitted inside the app to seven runs, shows the gap: it is blind to the crude.

9.9 Conclusion

This chapter studied a GLP-1 purification and isolation train as a simulator in which the chromatography cut, the membrane exchange and the filter cake share one material, with a product ledger that closes to 2×10^{-11} of the charge. A cut that adds 4.4 g to the pool and keeps it inside its purity limit loses the batch, because the late variant it collects shrinks the particles from 11.6 to 9.0 µm and the cake misses its hour. On the challenge, a purity-maximizing cut isolates 43.4 g, cuts that maximize recovery fail on the filter, and the downstream-aware rule isolates 59.9 g, but a fixed cut at the window it realizes does as well, because its soft sensor predicts 0.78 % late variant whatever the crude. Lower selectivity leaves only the narrow cut standing. The coupling, the cake and the isotherm are teaching closures without data. A fuller model would add measured retention, sieving and cake resistance for one peptide, a sensor that sees the crude, and an optimized schedule of cut and isolation. The simulator shows that a fraction cut is a decision about the filter as much as about the pool.

KEY POINTS

- A pool can pass its purity assay and still lose the batch: the late variant it carries shrinks the particles, and the cake misses its hour.
- Over 44 cut windows the dividing line is the pool's late variant, not its purity: at most 0.568 % where the batch passed, at least 0.665 % where it failed.
- An unmeasured 1.2 L between the detector and the valve delays the cut by 48 s at 1.5 L/min and drops the pool from 98.92 to 97.84 %.
- The downstream-aware rule isolates 59.9 g, but a fixed cut at its window does as well: its soft sensor predicts 0.78 % late variant whatever the crude.

Where else this applies. A product stream cut on an analyzer's running average returns on the batch column of Volume 8, Chapter 3.

References

[1] S. Murugesan, D. M. Hallow, J. P. Vernille, J. W. Tom, and J. E. Tabora, "Lean filtration: Approaches for the estimation of cake properties," *Organic Process Research & Development*, vol. 16, no. 1, pp. 42–48, 2012.

[2] L. Aumann and M. Morbidelli, "A continuous multicolumn countercurrent solvent gradient purification (MCSGP) process," *Biotechnology and Bioengineering*, vol. 98, no. 5, pp. 1043–1055, 2007.

[3] N. Kubota and J. W. Mullin, "A kinetic model for crystal growth from aqueous solution in the presence of impurity," *Journal of Crystal Growth*, vol. 152, no. 3, pp. 203–208, 1995.

[4] G. Guiochon, “Preparative liquid chromatography,” *Journal of Chromatography A*, vol. 965, no. 1–2, pp. 129–161, 2002.

[5] N. Marchetti, F. Dondi, A. Felinger, R. Guerrini, S. Salvadori, and A. Cavazzini, “Modeling of overloaded gradient elution of nociceptin/orphanin FQ in reversed-phase liquid chromatography,” *Journal of Chromatography A*, vol. 1079, no. 1–2, pp. 162–172, 2005.

[6] R. B. Merrifield, “Solid phase peptide synthesis. I. the synthesis of a tetrapeptide,” *Journal of the American Chemical Society*, vol. 85, no. 14, pp. 2149–2154, 1963.

[7] J. Lau, P. Bloch, L. Schäffer, I. Pettersson, J. Spetzler, J. Kofoed, K. Madsen, L. B. Knudsen, J. McGuire, D. B. Steensgaard, H. M. Strauss, D. X. Gram, S. M. Knudsen, F. S. Nielsen, P. Thygesen, S. Reedtz-Runge, and T. Kruse, “Discovery of the once-weekly glucagon-like peptide-1 (GLP-1) analogue semaglutide,” *Journal of Medicinal Chemistry*, vol. 58, no. 18, pp. 7370–7380, 2015.

[8] G. I. Bell, R. F. Santerre, and G. T. Mullenbach, “Hamster preproglucagon contains the sequence of glucagon and two related peptides,” *Nature*, vol. 302, no. 5910, pp. 716–718, 1983.

[9] D. J. Drucker, J. F. Habener, and J. J. Holst, “Discovery, characterization, and clinical development of the glucagon-like peptides,” *Journal of Clinical Investigation*, vol. 127, no. 12, pp. 4217–4227, 2017.

[10] L. R. Snyder, J. W. Dolan, and J. R. Gant, “Gradient elution in high-performance liquid chromatography,” *Journal of Chromatography A*, vol. 165, no. 1, pp. 3–30, 1979.

[11] A. Stadelmaier, R. O. Schönleber, D. Samson, and F. Dettner, “Purification of glucagon-like peptide 1 analogs,” European Patent 3 433 268 B1, Jul. 28, 2021.

[12] C. T. Dewar *et al.*, “Improved purification process of semaglutide,” European Patent 4 181 946 A1, May 24, 2023.

[13] S. Krishnan *et al.*, “Synthesis of an intermediate and its use for the preparation of GLP-1R/GIPR agonists,” U.S. Patent Application Patent 2026/0 116 848 A1, Apr. 30, 2026.

[14] A. Isidro-Llobet, M. N. Kenworthy, S. Mukherjee, M. E. Kopach, K. Wegner, F. Gallou, A. G. Smith, and F. Roschangar, “Sustainability challenges in peptide synthesis and purification: From R&D to production,” *The Journal of Organic Chemistry*, vol. 84, no. 8, pp. 4615–4628, 2019.

[15] B. F. Ruth, “Studies in filtration III. derivation of general filtration equations,” *Industrial & Engineering Chemistry*, vol. 27, no. 6, pp. 708–723, 1935.

[16] R. van Reis and A. Zydney, “Bioprocess membrane technology,” *Journal of Membrane Science*, vol. 297, no. 1–2, pp. 16–50, 2007.

[17] H. M. Hulburt and S. Katz, “Some problems in particle technology,” *Chemical Engineering Science*, vol. 19, no. 8, pp. 555–574, 1964.

[18] L. Rolinger, M. Rüdt, and J. Hubbuch, “A critical review of recent trends, and a future perspective of optical spectroscopy as PAT in biopharmaceutical downstream processing,” *Analytical and Bioanalytical Chemistry*, vol. 412, no. 9, pp. 2047–2064, 2020.

[19] D. Harvey, “High-performance liquid chromatography,” Section 12.5 of *Analytical Chemistry 2.1*, Chemistry LibreTexts open textbook, Creative Commons BY-NC-SA 4.0 license, 2021.

[20] 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.

[21] International Conference on Harmonisation, “Development and manufacture of drug substances (chemical entities and biotechnological/biological entities) Q11,” ICH harmonised tripartite guideline, Step 4 version, https://database.ich.org, May 2012.

9.10 Exercises

9.1 ○ Concept check · 2 min

In case 9 the cut shuts at 31.75 min instead of 31 min. The pool assays at 98.35 %, and the isolated solid is 98.83 % pure with 0.0019 % acetonitrile. Why is the batch lost?

(a) The isolated solid fails the 98.5 % purity limit once the assay error is allowed for

(b) The longer collection leaves too much acetonitrile in the liquor occluded in the cake

(c) The late variant slows particle growth, the cake resists more, and the filter misses its hour

(d) The larger pool passes more peptide through the membrane, and the cake is too small to wash

9.2 Concept check · 2 min

In case 4 the line from the detector to the cut valve is 1.2 L longer than the recipe allows for, and nobody has measured it. The cut is still set at 23 to 31 min of the gradient clock at 1.5 L/min. What does the pool collect, and how should the cut be corrected?

(a) Material that eluted later than planned, so both cut times should move earlier by about 0.8 min

(b) The same material as before, because the detector trace is unchanged and the cut follows the clock

(c) The same material as before, because the soft sensor carries the signal forward to the valve

(d) Material that eluted about 48 s earlier than planned, so both cut times should move later

9.3 Concept check · 2 min

Case 8 raises the Nutsche filter pressure from 1.5 to 3.5 bar, a factor of 2.33. The cake is compressible with $s = 0.60$. What does Equation (9.6) predict for the filtrate rate once the cake dominates the resistance?

(a) It rises by about 1.4 times, because the specific resistance grows as $\Delta P^{0.6}$

(b) It rises by about 2.3 times, because the flow through a cake is proportional to the pressure

(c) It stays the same, because the specific resistance rises in proportion to the pressure

(d) It falls, because the compressed cake closes up faster than the pressure pushes liquor through

9.4 Theory and calculation · 6 min

Use Equation (9.2) and Table 9.2: $q_{max} = 50\,\mathrm{kg/m^3}$, $\varepsilon = 0.60$, $b_P = 0.30$ and $b_L = 0.52\,\mathrm{m^3/kg}$ at $\varphi_{ref} = 0.30$, and $S_P = 60$, $S_L = 72$.

(a) Find the product's retention factor at vanishing load at 30 and at 22 % organic. *Checkpoint:* 10.0 and 1215.

(b) Find the selectivity b_L/b_P at 22, 30 and 33 % organic. *Checkpoint:* 4.53, 1.73 and 1.21.

(c) The modifier analyzer AT-202 reads 0.5 points high, and AIC-202 holds its measurement on the program. By what factor do the retention factors of the product and of the late variant change? *Checkpoint:* 1.35 and 1.43, both higher.

9.5 Theory and calculation · 6 min

Diafiltration at constant volume keeps a fraction e^{-sN_D} of a species whose observed sieving is s. Acetonitrile passes freely ($s = 1$). In case 6 the membrane sieves ten times its specification ($s_P = 0.10$), four diavolumes isolate 27.0 g, and two diavolumes leave 0.0127 % acetonitrile in the solid against the 0.010 % limit.

(a) Find the product lost during four diavolumes at the specified $s_P = 0.010$ and at $s_P = 0.10$. *Checkpoint:* 3.92 % and 33.0 %.

(b) Scale the 27.0 g of case 6 to two and to eight diavolumes. *Checkpoint:* 33.0 g and 18.1 g.

(c) Find the smallest setting on the app's 0.5 diavolume grid that meets the acetonitrile limit in case 6, and the mass it isolates. *Checkpoint:* 2.5 diavolumes, 0.0077 % and 31.4 g.

9.6 Theory and calculation · 6 min

Use Equation (9.6) with $\rho_s = 1150\,\mathrm{kg/m^3}$, $\varepsilon_c = 0.60$, $A_f = 0.05\,\mathrm{m^2}$ and $s = 0.60$, and the scorecard values of Table 9.5.

(a) Find the depth of the 56.6 g cake of the reference case. *Checkpoint:* 2.46 mm.

(b) Case 2 gives 11.6 µm and 12.6×10^{12} m/kg, and case 9 gives 9.0 µm and 3.4×10^{13} m/kg, both at 1.5 bar. Split the rise in α into the factor from particle size and the rest. *Checkpoint:* 2.70 in all: 1.66 from size and 1.62 from the rest.

(c) Predict α for case 9 filtered at 3 bar, and the factor on the cake-dominated filtrate rate. *Checkpoint:* 5.15×10^{13} m/kg and 1.32.

9.7 App challenge · 8 min

Load scenario 4, A longer line to the collector. The collector line holds an unmeasured 1.2 L beyond the recipe. Use the default seed and run each case at 1800 × to the discharge.

(a) Find the extra delay at 1.5 L/min. Run the scenario as shipped and read the pool purity and the Isolated purity tile. *Checkpoint:* 48 s; 97.84 % and 98.26 %, nothing acceptable.

(b) Press Reset, set Cut opens to 23.5 and Cut shuts to 31.5 min, and run again. Read the same tiles and the acceptable mass. *Checkpoint:* 98.40 % and 98.81 %, 55.4 g.

(c) Press the reload button, set the Collection rule to Downstream aware, and run again. *Checkpoint:* 96.90 % and 97.40 %, nothing acceptable.

Accept purities within 0.01 points and masses within 0.1 g. Explain why the soft sensor does not correct the delay.

9.8 App challenge · 8 min

Load scenario 9, The wide pool that passes its assay, with the cut shut at 31.75 min. Use the default seed and run each case at 1800 × to the discharge.

(a) Run the scenario as shipped. Read Late variant in the pool (true), Particle size · cake resistance and Time on the filter. *Checkpoint:* 0.939 %; 9.0 µm and 34.0×10^{12} m/kg; out of time while drying, off at 79.3 min.

(b) Press Reset, set the Filter pressure to 3.0 bar, and run again. Read the same tiles and the acceptable mass. *Checkpoint:* 51.5×10^{12} m/kg, 54.9 min, 60.2 g at 9.8 g/h.

(c) Press the reload button, set the Collection rule to UV threshold, and run again. Read the pool tile and Time on the filter. *Checkpoint:* 95.0 % and 91.64 %; out of time while filtering.

Accept resistances within 0.2×10^{12} m/kg and times within 0.1 min. Compare (b) with the prediction of $\alpha \propto \Delta P^{0.6}$.

9.9 App challenge · 7 min

Load the challenge, the most acceptable peptide per hour. Use the default seed and run each case at 1800 × to the discharge.

(a) Run the challenge as shipped. Read the Challenge target tile, Residual solvent and Particle size · cake resistance. *Checkpoint:* cake not off the filter in time; 0.0273 % acetonitrile; 6.7 µm.

(b) Press the reload button, press Reveal suggested tuning, and run again. Read the acceptable mass, the pool tiles and the Challenge target tile. *Checkpoint:* 59.9 g at 10.0 g/h; 78.2 and 98.51 %, 0.430 %; met.

(c) Press Reset, set Crude purity from the synthesis to 65 %, and run again. This is an experiment, not an attempt. *Checkpoint:* pool 97.14 %, isolated purity 97.70 %, target missed on purity.

Accept masses within 0.1 g and purities within 0.01 points.

9.10 Application area · 4 min

Every policy of Table 9.6 uses 80.9 L of water and 45.5 L of organic solvent per campaign, 37.0 L of it acetonitrile.

(a) Find the total solvent, the acetonitrile share, and the solvent per gram of acceptable peptide for the reference cut (56.2 g) and the downstream-aware rule (59.9 g). *Checkpoint:* 126.4 L, 29.3 %, 2.25 and 2.11 L/g.

(b) Explain why the solvent per gram follows the acceptable mass, and say what the recovery-maximizing cut costs in solvent. *Checkpoint:* 6.2 % less per gram; all 126.4 L for no acceptable product.

9.11 Application area · 4 min

The pool assay reports purity and late-variant fraction 15 min after the transfer, with standard deviations of 0.06 and 0.03 points. In the reference case the stages from conditioning to precipitation take 5.3, 44.8, 114.6, 37.1 and 24.5 min.

(a) Find how long before the filtration starts the assay result is available. *Checkpoint:* 211 min.

(b) In the cut sweep of Figure 9.3, passing pools carry at most 0.568 % of the late variant and failing pools at least 0.665 %. Express the gap in assay standard deviations, and propose a use of the result. *Checkpoint:* 3.2 standard deviations.

9.12 Review the agent · 5 min

An AI assistant reviews the challenge and writes: “The downstream-aware rule isolates 59.9 g against 56.2 g for the fixed 23 to 31 min cut, because its soft sensor predicts the late variant that the filter answers to. The 3.7 g is the value of that prediction. Use the rule on every campaign. When a weaker crude arrives, the sensor will see the extra late variant and shut the cut earlier.” Find the flaw and say how you would check it.

CHAPTER 10

Sealed Orbiting Cabin Simulator Where Added Oxygen Shuts the Nitrogen Valve

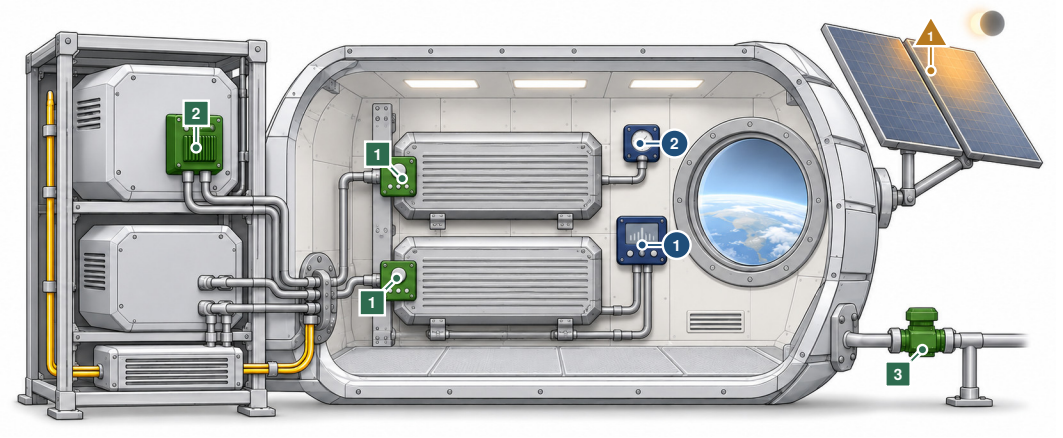

The controller moves 1 bed swaps and heater of the two carbon dioxide beds
2 electrolysis power, which makes the oxygen 3 nitrogen makeup valve

The controller reads 1 gas analyzer: oxygen, carbon dioxide, vapor and nitrogen 2 total pressure

What disturbs it 1 sunlight and eclipse on the array that charges the battery

Notice that the oxygen maker, the beds and the makeup valve all act on one sealed volume of air and draw on one array and battery, so each loop meets the others in the cabin's total pressure or on the power bus. Illustration, not to scale.

IN PLAIN TERMS A Sealed Room That Can Only Add Air

The main idea. Six people in a closed room in orbit breathe in oxygen and breathe out carbon dioxide and water. Machines add oxygen and nitrogen, and two beds take turns pulling out the carbon dioxide.

The challenge. Both gases push on one shared pressure, so the two controllers that add them answer each other. The beds and the oxygen maker also share one battery, which runs low each time the station passes into Earth's shadow.

Why it matters. A week aboard stays safe only when the controllers, the bed timing and a running count of the nitrogen, which finds a slow leak, work as one plan.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 10.2, the fact in Section 10.3, the two try-its and the fact in Sections 10.4 and 10.6, and the key points at the end of Section 10.9
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the two open challenges

The control question A sealed cabin can only add gas, so its oxygen and pressure loops answer each other through one total, and its carbon dioxide beds share a battery with the oxygen generator through every eclipse. How must the loops, the bed schedule and a nitrogen balance work together so that a week aboard holds every limit and a leak is found before the tank runs low?

Learning objectives

- Estimate the hours to each atmospheric limit with the machines stopped from the crew's rates and the cabin's moles per kilopascal.
- Explain why a higher oxygen setpoint shuts the nitrogen valve, and classify the coupling by its gain matrix and relative gain.
- Compute the eclipse energy budget, and say which load a power-aware schedule must move out of the dark.
- Describe how an hourly nitrogen balance with a cumulative-sum test finds a leak that the pressure loop feeds and hides.
- Judge which layers a week aboard needs by running it with each one removed.

Prerequisites Sections 1.4.3, 1.4.5 and 1.4.7 (inventories, triangular plants, conditioning), 1.6.2 and 1.6.4. Chapter 8 treats an inert that leaves by one exit.

The app https://apmonitor.com/apps/space-station-eclss/

See also Volume 8, Chapter 12 for on-off loads on one shared supply scheduled by a hybrid predictive controller, and Volume 8, Chapter 14 for an integer actuator under a PI loop. Volume 9, Chapter 11 for an integer pump schedule by branch and bound and a night-flow leakage estimate, and Volume 9, Chapter 9 for cooling shifted in time by the price of energy.

10.1 Introduction

Six people in a sealed room of 400 m^3 at 22 °C give off 29.4 mol of water vapor and 5.4 mol of carbon dioxide an hour and take 6.0 mol of oxygen. With every machine of the simulated module switched off, the relative humidity passes 60 % at 1.46 h, the carbon dioxide partial pressure passes 0.53 kPa at 4.84 h and the oxygen partial pressure falls below 19.5 kPa only at 39.19 h. The total pressure meanwhile drifts from 101.30 to 102.13 kPa over two days, because the crew trades oxygen for carbon dioxide and water almost mole for mole. A cabin's priorities follow from that order, and so do its control problems. As for the spent-fuel pool of Volume 1, Chapter 9, the operative quantity with the machines stopped is a time to a limit rather than an error.

The environmental control and life support systems of crewed spacecraft have been documented for designers in a reference publication of the 1990s [1], a textbook of their subsystems [2], a baseline of design loads and assumptions [3] and a standard for the habitable atmosphere [4]. Carbon dioxide is removed in flight by alternating molecular sieve beds that are tested as whole assemblies [5] and modeled in detail as packed sorption columns [6], oxygen is recovered from that carbon dioxide by a Sabatier reactor fed with electrolysis hydrogen [7], and the crew's health sets how much carbon dioxide is tolerated [8]. The control problem belongs to a broader class. Any closed inventory whose actuators can only add material (an inerted vessel, a glove box, a submarine's atmosphere), any cyclic adsorber such as a pressure swing unit [9], and any plant whose discrete operations are scheduled against a periodic energy supply share its shape, and the integration of scheduling with process control has been reviewed as a field of its own [10]. The discrete half of such a problem is a mixed-integer program, for which branch and bound was written [11], and the diagnostic half is the reconciliation of flows and inventories against a mass balance [12] with a sequential test for a change in its mean [13].

One belief the simulator sets out to correct is that a pressure loop protects the inventory it controls. In this model a nitrogen valve that holds the total pressure inside its band is also the actuator that feeds a new leak without anyone noticing. In case 7 the valve opens from 6.0 to 37.0 % after the leak steps up at 24 h while the pressure stays between 101.00 and 101.40 kPa, well inside its 97 kPa to 103 kPa band. The same sum of partial pressures couples the two gas loops. Raising the oxygen

setpoint by 1 kPa shuts the nitrogen valve two and a half hours later, which a reader who has met each loop alone may not expect.

The material a student meets treats these pieces apart. Design references tabulate loads, limits and assembly performance [1, 3, 5], bed studies resolve one column's breakthrough [6], and scheduling studies fix the plant they schedule [10]. This chapter works with a browser simulator of a sealed orbiting cabin in which the gas inventories, the alternating beds, the oxygen and Sabatier chain, the battery and the leak share one run, and it measures on the running engine what fixed, adaptive and power-aware bed schedules, a crew-activity feedforward and a nitrogen balance achieve on the same week. The ideal-gas balances, the bed law, the cumulative-sum test, the proportional-integral loops and branch and bound over a linear relaxation are established methods. Behind the measurements are a forty-state model whose element balances close to the integrator's precision, a three-orbit power-aware program on the branch and bound that the apps share, which reports its proof or its gap to the reader, and a wall between the instruments and the hidden state that no controller, predictor or balance crosses. Ten cases, each graded by a completion rule (Section 1.6.4), range from a cabin with nothing running to eight crew on a faulted bed that no schedule can hold, and a graded week aboard follows them.

Section 10.2 follows the domain from expendable canisters to regenerable beds. Section 10.3 sets out the cabin, its beds, its humidity and power and its instruments, and Section 10.4 measures the three behaviors that shape its control: one total pressure for two gas loops, the saw-tooth of the beds and an eclipse budget that does not close. Section 10.5 describes the three loops, the fixed, adaptive and power-aware bed schedules and the nitrogen balance. Section 10.6 measures the cases and the week, with the program's proofs, the code verification and the agreement with published values, Section 10.7 describes the real analyzer, removal assembly and controllers, Section 10.8 states the limitations and Section 10.9 concludes. The Exercises at the end of the chapter run the cases and the week.

10.2 From Expendable Canisters to Regenerable Beds

The domain's early story, from pure oxygen to the first cabin with two gases, is told in the box below. A later long-duration station replaced superoxide oxygen cartridges with water electrolysis and its expendable canisters with a four-bed molecular sieve vented to space [1]. The same reference states the principle the simulator is built on, that the oxygen or nitrogen added is decided by monitoring the total pressure and the oxygen partial pressure, and that a tighter habitat loses leakage and venting as ways to remove an excess constituent [1].

> ⧖ HISTORY **1967 to 1974: From Pure Oxygen to a Two-Gas Cabin**
>
> The early crewed spacecraft of the United States breathed pure oxygen at 34.5 kPa. After a fire during a practice session on the launch pad in 1967, prelaunch testing used a mixture of 40 percent nitrogen and 60 percent oxygen, because nitrogen inhibits the flammability of oxygen. Carbon dioxide was absorbed by lithium hydroxide, which reacts irreversibly, so for missions longer than about two weeks the stored canisters became prohibitive. A ground test that began in 1972 kept a crew of three for 56 days in a sealed chamber, with two-gas control of oxygen and nitrogen among its subsystems. The three-person station it prepared flew three missions of 28, 59 and 84 days in 1973 and 1974, with a mixed oxygen and nitrogen atmosphere and a two-canister molecular sieve, regenerated by exposure to space vacuum, in place of lithium hydroxide [1].

The removal assembly of the current station era has four beds, two desiccant and two sorbent, which swap each half-cycle while the regenerating sorbent bed is heated under vacuum. Ground testing in 1999 reported flows averaged over a full cycle of 288 min, a half-cycle of 144 min, and contrasted the water-saving desiccant beds with the earlier two-bed design that vented its water to space [5]. Reduction of the removed carbon dioxide followed. A Sabatier assembly that combines carbon dioxide with electrolysis hydrogen into methane and water, under development for more than

Table 10.1 Classification of the Cabin's Variables and Their Starting Values

Variable	Role	Start	Limit or range
p_{CO_2}	CV, analyzer	0.37 kPa	$\leq$ 0.53 kPa (teaching)
p_{O_2}	CV, analyzer	SP 21.0 kPa	19.5–23.1 kPa
$x_{O_2} = p_{O_2}/P$	constraint	20.73 %	$\leq$ 23.5 %
$P = \sum p_i$	CV, transducer	SP 101.3 kPa	97–103 kPa
RH	CV, from analyzer	SP 50 %	40–60 %
E/E_{bat}	constraint, meter	90 %	floor 60 %
P_e	MV, electrolysis	1.2 kW manual	0–2 kW
u_{N_2}	MV, makeup valve	6 % manual	0–100 %
u_c	MV, condenser flow	70 % manual	0–100 %
swap, x_k	discrete MV, beds	bed A adsorbs	half-cycle 75–180 min[a]
n_c, a	DV, crew, activity	6, 1	2–10, 1 or 2.5
ΔC_{leak}, f	DV, leak step, bed fault	0, 1	0.05, 0.5
Q_A, Q_B	states, hidden	3.70, 0.30 mol	—
T_A, T_B	states, thermocouples	22 °C	to 227 °C
n_{N_2}, n_{O_2}	states, cabin	12 813, 3423 mol	—
L	leak, hidden	0.93 mol/h N_2	—

[a] Adaptive and power-aware schedules. The fixed half-cycle is 144 min as shipped.

twenty years, joined an oxygen generator launched in 2006, and its installation and activation were reported in 2011 [7]. Design studies take their loads from a baseline document that gives 0.84 kg of oxygen consumed and 1.01 kg of carbon dioxide and 2.27 kg of water vapor produced per crew member and day [3].

The tolerance for carbon dioxide has tightened as the record grew. An analysis of in-flight headache reports found that the odds of a headache doubled for each 1 mmHg of carbon dioxide and that a seven-day average below 2.5 mmHg would keep the risk under one percent, against an observed range of 1 to 9 mmHg [8]. The current standard limits the one-hour average partial pressure to 3 mmHg (0.40 kPa) [4]. Bed models in the literature are axially dispersed, nonisothermal, linear-driving-force descriptions of molecular sieve 5A columns, and their comparison with data called for two dimensions to capture wall effects [6].

10.3 The Cabin as Modeled

The simulated module is a generic teaching design that names no flight assembly, as shown in Figure 10.1. A well-mixed cabin of free volume $V = 400\,\mathrm{m}^3$ at a fixed $T = 22\,^\circ\mathrm{C}$ holds nitrogen, oxygen, carbon dioxide and water vapor. Six crew live in it, two visitors may come aboard and two crew may leave through a $12\,\mathrm{m}^3$ airlock for an excursion. A rack holds two carbon dioxide beds on outlet valves to a vacuum line, a Sabatier reactor on that line, an electrolysis unit, a condenser, a nitrogen tank on a makeup valve, an oxygen tank on a backup valve and a water tank. The orbit lasts $T_{orb} = 92\,\mathrm{min}$, of which $T_{ecl} = 35\,\mathrm{min}$ are in Earth's shadow, and every case starts at sunrise. Time is in hours, amounts in moles, pressures in kilopascals, power in kilowatts and energy in kilowatt-hours.

The variables are classified in Table 10.1. Forty states are integrated. Fourteen are physical, the cabin's four inventories n_i, the loadings Q_A, Q_B and temperatures T_A, T_B of the two beds, the battery energy E and the nitrogen, oxygen and water tanks. The other twenty-six are ledgers that integrate every source and sink, so that each species can be balanced at every step.

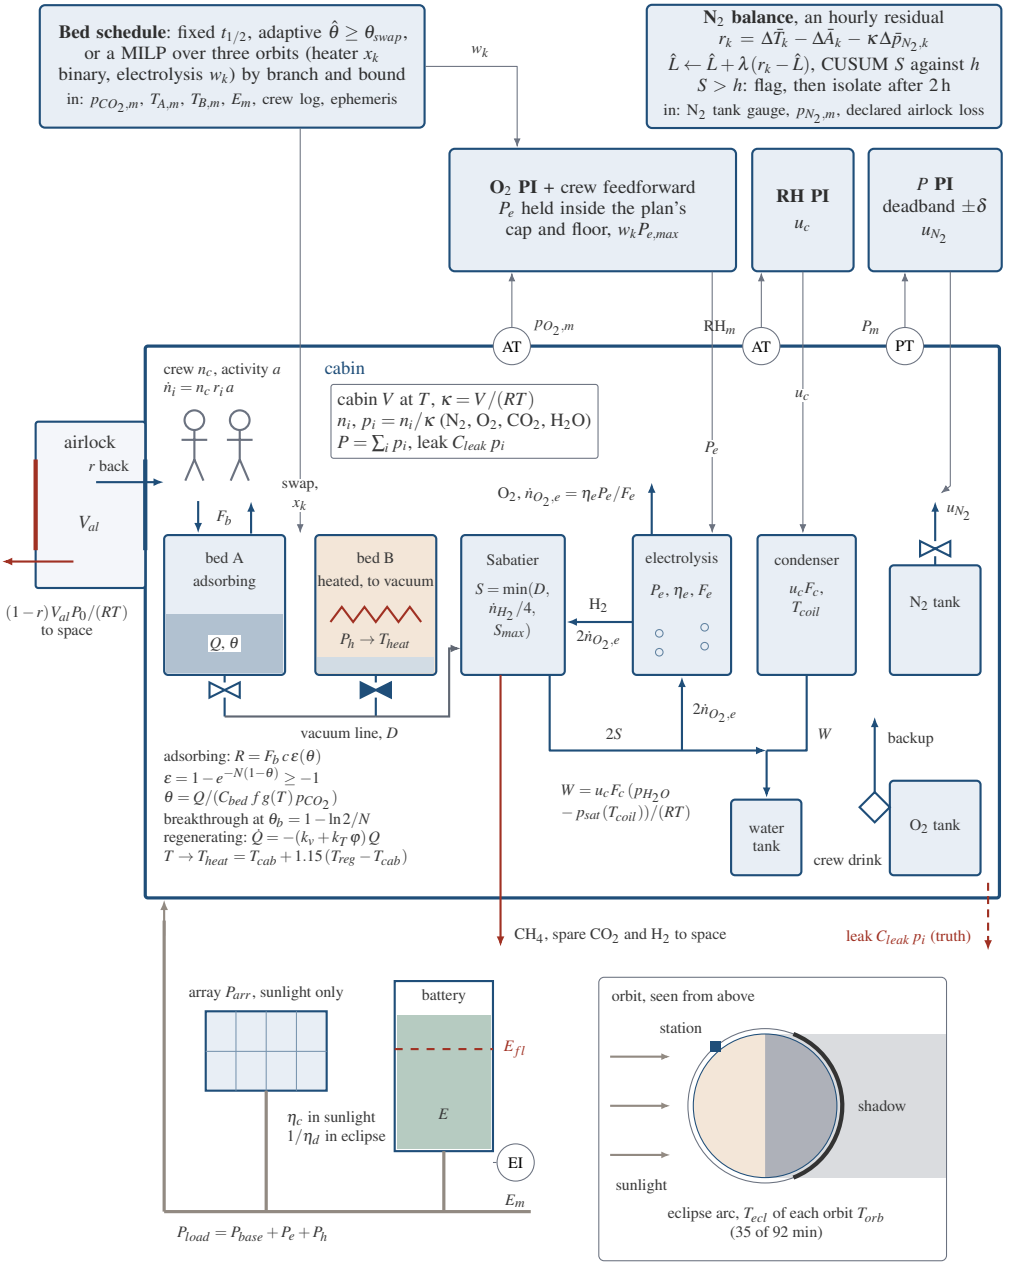

Figure 10.1 Sealed cabin, rack, power bus and orbit as modeled, drawn after the Module view. Bed A adsorbs cabin air at F_b with its outlet valve shut while bed B is heated toward T_{heat} with its valve open to the vacuum line that feeds the Sabatier reactor. The electrolysis unit sends oxygen to the cabin and hydrogen to the reactor, which vents methane and the spare carbon dioxide and hydrogen to space. The condenser, the reactor and the electrolysis share the water line, from which the crew also drink. The leak (dashed) is teaching truth that no instrument reads. The three loops read the analyzer (AT) and the pressure transducer (PT), and the schedule and the balance list their inputs. The inset draws the orbit at the radius whose shadow crossing lasts T_{ecl}. Process model only, not a flight layout. Not to scale.

10.3.1 A Sealed Inventory

Each gas changes only by what the crew, the machines, the leak and the airlock move, and its partial pressure is its inventory over the cabin's moles per kilopascal, as shown in Equation (10.1).

$$\begin{aligned}
\frac{\mathrm{d}n_{\mathrm{N_2}}}{\mathrm{d}t} &= u_{N_2} Q_{N_2} - (C+v)\, p_{\mathrm{N_2}} \\
\frac{\mathrm{d}n_{\mathrm{O_2}}}{\mathrm{d}t} &= \dot{n}_{O_2,e} + b\, Q_{O_2,b} - n_c r_{O_2} a - (C+v)\, p_{\mathrm{O_2}} \\
\frac{\mathrm{d}n_{\mathrm{CO_2}}}{\mathrm{d}t} &= n_c r_{CO_2} a - R - (C+v)\, p_{\mathrm{CO_2}} \\
\frac{\mathrm{d}n_{\mathrm{H_2O}}}{\mathrm{d}t} &= n_c r_{H_2O} a - W - W_{wall} - (C+v)\, p_{\mathrm{H_2O}} \\
p_i &= n_i/\kappa, \qquad \kappa = V/(RT), \qquad P = \textstyle\sum_i p_i
\end{aligned} \tag{10.1}$$

where n_c counts the crew and visitors in the cabin, a is 1 at rest and $a_{ex} = 2.5$ while everyone exercises, u_{N_2} is the makeup valve's opening on a full flow of $Q_{N_2} = 15\,\mathrm{mol/h}$, b is the backup oxygen valve (open or shut, $Q_{O_2,b} = 10\,\mathrm{mol/h}$), $C = C_{leak} + \Delta C_{leak}$ is the leak coefficient until a step is isolated, and v is the airlock's vent coefficient. At $R = 8.314\,462\,618 \times 10^{-3}\,\mathrm{kPa\,m^3/(mol\,K)}$, $\kappa = 163.00\,\mathrm{mol/kPa}$, so one mole of any gas raises the total by 0.006 135 kPa. A leak to vacuum takes the cabin's mixture, so the hole passes each species in proportion to its partial pressure. The airlock is depressurized over $t_{dep} = 0.5\,\mathrm{h}$, and the vent coefficient is

$$v = \frac{(1-r)\, V_{al}\, P_0}{R\, T\, t_{dep}\, P} \tag{10.2}$$

during that stage and zero otherwise, where $r = 0.90$ is the reclaimed fraction and P_0 the cabin pressure when the stage starts. The vent therefore loses $(1-r)V_{al}P_0/(RT)$, 49.5 mol at 101.3 kPa, and the crew outside for $t_{eva} = 6.5\,\mathrm{h}$ leave the cabin's count.

10.3.2 Crew, Electrolysis and Sabatier Reduction

The resting rates per person are $r_{O_2} = 1.0$, $r_{CO_2} = 0.9$ and $r_{H_2O} = 4.9\,\mathrm{mol/h}$. A day with one exercise hour then consumes 0.816 kg of oxygen and produces 1.010 kg of carbon dioxide and 2.251 kg of vapor per person. Electrolysis makes oxygen from the tank's water at the Gibbs energy of the reaction, and its hydrogen feeds the Sabatier reactor, which reduces the carbon dioxide arriving on the vacuum line, as shown in Equation (10.3).

$$\begin{aligned}
\dot{n}_{O_2,e} &= \eta_e P_e / F_e, \qquad \dot{n}_{H_2} = 2\dot{n}_{O_2,e} \\
S &= \max\left(0,\ \min(D,\ \dot{n}_{H_2}/4,\ S_{max})\right) \\
\frac{\mathrm{d}W_t}{\mathrm{d}t} &= W + 2S - 2\dot{n}_{O_2,e} - n_c r_{H_2O} a
\end{aligned} \tag{10.3}$$

where $F_e = 0.131\,74\,\mathrm{kWh/mol}$ is twice the Gibbs energy of formation of liquid water per mole of oxygen, $\eta_e = 0.70$, so the full $P_{e,max} = 2\,\mathrm{kW}$ makes 10.63 mol/h, D is the desorbing bed's flow and $S_{max} = 1.5\,\mathrm{mol/h}$. The reaction $\mathrm{CO_2 + 4\,H_2 \rightarrow CH_4 + 2\,H_2O}$ is the one tabulated for this service [1, 7]. Methane, the carbon dioxide the reactor cannot take and the spare hydrogen are vented. The crew drink exactly the water they exhale, a declared closure that makes the water tank a ledger rather than a water recovery system.

> ✦ DID YOU KNOW? **How Much Water Comes Back**
>
> On the current station, distillation recovers about 87 percent of the water in urine, a limit set to prevent unwanted precipitation. Missions beyond low Earth orbit are to show more than 98 percent [14]. The model of this chapter skips that loop: its crew drink exactly the water they exhale.

10.3.3 Two Beds That Fill and Empty

The adsorbing bed takes carbon dioxide from the air drawn through it with the efficiency of N transfer units at its loading fraction θ, and the regenerating bed gives its carbon dioxide up to vacuum at a rate that heat raises, as shown in Equation (10.4).

$$\begin{aligned} R &= F_b\, c\, p_{\mathrm{CO_2}}\, \varepsilon(\theta), \quad \varepsilon = \max\left(-1,\ 1-\mathrm{e}^{-N(1-\theta)}\right) \\ \theta &= \frac{Q_j}{C_{bed} f_j\, g(T_j)\, p_{\mathrm{CO_2}}}, \quad g(T) = \mathrm{e}^{-\max(0,\,T-T_{cab})/\Delta T_c} \\ \frac{\mathrm{d}Q_j}{\mathrm{d}t} &= R, \quad \frac{\mathrm{d}Q_k}{\mathrm{d}t} = -D = -\left(k_v + k_T\, \varphi(T_k)\right) Q_k \\ \varphi &= \min\left(1,\ \max(0,\ \Theta)\right), \qquad \Theta = \frac{T - T_{cab}}{T_{reg} - T_{cab}} \end{aligned} \tag{10.4}$$

where bed j adsorbs and bed k regenerates, $c = 1/(RT) = 0.407\,50\,\mathrm{mol/(m^3\,kPa)}$, $F_b = 55\,\mathrm{m^3/h}$, $N = 2.5$, $C_{bed} = 40\,\mathrm{mol/kPa}$, f_j is 1 or one half on a faulted bed, $\Delta T_c = 40\,\mathrm{K}$, $k_v = 0.15\,\mathrm{h^{-1}}$, $k_T = 3.5\,\mathrm{h^{-1}}$ and $T_{reg} = 200\,^\circ\mathrm{C}$. The outlet reaches half its inlet, which is breakthrough, at $\theta_b = 1 - \ln 2/N = 0.7227$. The floor on ε lets a bed swapped in while still hot give carbon dioxide back to the cabin. The bed temperatures follow first-order laws, $\tau_h \dot{T} = T_{heat} - T$ with the heater on, $\tau_v \dot{T} = T_{cab} - T$ for a bed under vacuum without heat and $\tau_a \dot{T} = T_{cab} - T$ for the bed in the air flow, with $\tau_h = 0.3\,\mathrm{h}$, $\tau_v = 0.8\,\mathrm{h}$, $\tau_a = 0.25\,\mathrm{h}$ and $T_{heat} = T_{cab} + 1.15\,(T_{reg} - T_{cab}) = 226.7\,^\circ\mathrm{C}$. With the beds stopped, neither takes air and both cool at τ_v. A regenerating bed's θ rises far above one because its hot capacity is small, so the scene draws it on its cool capacity $Q/(C_{bed} f\, p_{\mathrm{CO_2}})$ and the export logs both.

10.3.4 Humidity and Power

The condenser passes a fraction u_c of $F_c = 400\,\mathrm{m^3/h}$ of cabin air over a coil at $T_{coil} = 7\,^\circ\mathrm{C}$ and keeps the vapor above the coil's saturation, while vapor above saturation in the cabin itself condenses on the walls, as shown in Equation (10.5).

$$\begin{aligned} W &= u_c F_c\, c\, \max\left(0,\ p_{\mathrm{H_2O}} - p_{sat}(T_{coil})\right) \\ W_{wall} &= K_{wall}\, \kappa\, \max\left(0,\ p_{\mathrm{H_2O}} - p_{sat}(T)\right) \\ p_{sat}(T) &= 0.61094\, \exp\left(17.625\, T/(T + 243.04)\right) \end{aligned} \tag{10.5}$$

where $K_{wall} = 5\,\mathrm{h^{-1}}$ and p_{sat} is the Magnus form in kPa with T in °C, 2.638 kPa at the cabin temperature and 1.0007 kPa at the coil. The wall term is a teaching closure that limits the humidity growth but does not enforce a saturation cap: case 1 settles near 101.37 %. The bus carries a base load, the electrolysis and the regenerating bed's heater, as shown in Equation (10.6).

$$\begin{aligned} P_{load} &= P_{base} + P_e + x P_h, \qquad P_{sun} = P_{arr}\, s(t) \\ \frac{\mathrm{d}E}{\mathrm{d}t} &= \begin{cases} \eta_c (P_{sun} - P_{load}) & P_{sun} \geq P_{load},\ E < E_{bat} \\ (P_{sun} - P_{load})/\eta_d & P_{sun} < P_{load} \end{cases} \end{aligned} \tag{10.6}$$

where $s(t)$ is 1 in sunlight and 0 for the last 35 min of each orbit, $P_{arr} = 11\,\mathrm{kW}$, $P_{base} = 4\,\mathrm{kW}$, $P_h = 1\,\mathrm{kW}$, $E_{bat} = 8\,\mathrm{kWh}$ and $\eta_c = \eta_d = 0.95$. The battery is clipped to $[0, E_{bat}]$ after each step, and a protection sheds the electrolysis and the heater below 35 % and restores them above 50 %. The base load alone takes $4 \times 35/60/0.95 = 2.456\,\mathrm{kWh}$ in an eclipse, of the 3.2 kWh that a full battery holds above its 60 % floor.

Table 10.2 Model Parameters. S Is Read From a Source, T a Teaching Value, D Derived

Symbol	Meaning	Value	Unit	
V, T	cabin volume, temperature	400, 22	m^3, °C	T
κ	moles per kilopascal	163.00	mol/kPa	D
r_{O_2}, r_{CO_2}, r_{H_2O}	rates per person	1.0, 0.9, 4.9	mol/h	T[a]
a_{ex}	activity, exercising	2.5	—	T
C_{bed}	bed capacity, cool	40	mol/kPa	T
F_b, N	bed air flow, transfer units	55, 2.5	m^3/h, —	T
ΔT_c	capacity e-fold	40	K	T
k_v, k_T	desorption rates	0.15, 3.5	h^{-1}	T
T_{reg}, P_h	regeneration, heater	200, 1.0	°C, kW	T
τ_h, τ_v, τ_a	bed thermal lags	0.3, 0.8, 0.25	h	T
f	faulted bed's capacity	0.5	—	T
F_e	energy per mole O_2	0.13174	kWh/mol	S
η_e, $P_{e,max}$	electrolysis	0.70, 2.0	—, kW	T
S_{max}	Sabatier capacity	1.5	mol/h	T
Q_{N_2}, $Q_{O_2,b}$	makeup, backup flows	15, 10	mol/h	T
tanks	N_2, O_2, water	900, 600, 3000	mol	T
reserves	N_2, O_2	575, 300	mol	T
F_c, T_{coil}	condenser flow, coil	400, 7	m^3/h, °C	T
C_{leak}, ΔC_{leak}	leak, leak step	0.0118, 0.05	mol/(h kPa)	T[b]
V_{al}, r	airlock, reclaimed	12, 0.90	m^3, —	T, S[c]
t_{dep}, t_{eva}, t_{rep}	airlock stages	0.5, 6.5, 0.25	h	T
P_{arr}, P_{base}	array, base load	11, 4.0	kW	T
E_{bat}	battery	8	kWh	T
floor, shed	battery limits	0.60, 0.35	—	T
η_c, η_d	battery efficiencies	0.95, 0.95	—	T
T_{orb}, T_{ecl}	orbit, eclipse	92, 35	min	T

[a] A day with one exercise hour is within 3 % of the baseline's 0.84, 1.01 and 2.27 kg [3]. [b] About 0.83 kg of air a day, roughly ten times the baseline's upper module leakage of 0.09 kg/day [3]. [c] The baseline gives about 10 % of an airlock's gas lost per cycle [3].

10.3.5 Measurements

A gas analyzer reports the four partial pressures once a minute with white noise of standard deviation 0.03, 0.03, 0.0015 and 0.01 kPa for nitrogen, oxygen, carbon dioxide and vapor, and holds each report. The humidity is computed from its vapor reading. A pressure transducer (0.005 kPa), the three tank gauges (0.5 mol), a battery energy meter (0.02 kWh) and the two bed thermocouples (0.5 °C) are scanned with it, each with colored noise. The crew's log, the electrolysis trip contact, the bed fault annunciation and the airlock stage are status inputs. The bed loadings and the leak are never measured. Each instrument draws from its own stream seeded by the case, so a run replays bit for bit (Section 1.6.1). The exported file logs each measurement next to the truth it reads, with the actuators, the schedule and the program's status, nodes and gap, the nitrogen ledger and the balance's residual, estimate and flag.

10.3.6 Parameters, Start and Numerics

Table 10.2 lists the parameters. All appear in the Parameters panel and can be changed by the reader. Each case starts with the oxygen at its setpoint, the carbon dioxide at 0.37 kPa, the vapor at the humidity setpoint and nitrogen making up the total, bed A a quarter of the way to its equilibrium loading and bed B freshly regenerated, both at cabin temperature, the battery at 90 % and the tanks full. The engine of Section 1.6.1 integrates in hours with the classical fourth-order Runge-Kutta method at a step of $\Delta t = 10$ s, holds every event for the step it falls in, and applies discrete changes (an event, a bed swap, an airlock stage, the clipping of a tank or the battery) once at a step boundary. Halving the step moves the challenge week's nitrogen losses by 0.0014 % as shipped and 0.016 % with the suggested settings.

10.3.7 Assumptions

The cabin is one well-mixed volume at a fixed temperature, so the pocket of exhaled carbon dioxide around a sleeping crew member, which the standard's rationale names, and the thermal control loop are absent. Each bed is one loading and one temperature with a linear isotherm whose capacity falls as the bed warms, which reproduces breakthrough and the cost of swapping in a hot bed but not the axial and radial profiles of a real column. The real assembly's desiccant beds are left out, and the condenser carries all of the humidity. The crew's rates are constant per person, and exercise multiplies them for the whole crew at once. The airlock's lost gas leaves at a constant rate over the depressurization. The power model has one array, one battery with constant efficiencies and one fixed orbit. Every limit, reserve and efficiency other than those marked S in Table 10.2 is a teaching value, and a run that holds them is not evidence about any real vehicle.

10.4 The Sum, the Saw-Tooth and the Shadow

10.4.1 One Total Pressure for Two Loops

Adding oxygen alone raises both the total and the oxygen fraction. One mole raises P by $1/\kappa = 0.006\,135$ kPa and the fraction by $(P - p_{O_2})/(\kappa P^2)$, 0.004 801 points at the start. In case 2 the oxygen setpoint is raised from 21.0 to 22.0 kPa at 1 h. The electrolysis runs at its full 2 kW, the backup tank opens because the oxygen is more than 0.8 kPa under its new setpoint, and the extra oxygen pushes the total past the pressure loop's 0.2 kPa deadband. The nitrogen valve, open 6.0 % to feed the everyday leak, shuts at 3.52 h, 2.52 h after the step. Nothing can remove the excess except the crew and the leak, so at 24 h the total is 102.16 kPa, the oxygen 21.97 kPa and its fraction 21.51 % against 20.73 % before the step, and the backup tank has given 58 mol. The two loops are coupled through the sum, and only one of them can see it. The oxygen added moves the total, while the nitrogen added leaves the oxygen partial pressure alone, so the gain matrix from the two gases to the oxygen partial pressure and the total is triangular. Its relative gain is one however strong the coupling (Section 1.4.5), and the coupling runs one way. Nothing in the cabin removes the excess, which is the closed inventory of Section 1.4.3 with actuators that can only add.

Pause. Try Exercise 10.1 before reading on.

> **TRY IT IN THE APP 2 · The Oxygen Loop and the Pressure It Shares**
>
> Run the scenario at 0.25 h/s. After the first half hour raise the O_2 setpoint from 21.0 to 22.0 kPa, then follow the Nitrogen valve strip on the Atmosphere view. Open about 6 % to feed the everyday leak, it shuts about two and a half hours after the step, and the total pressure and the oxygen fraction stay high, because only the crew and the leak remove gas.

10.4.2 A Saw-Tooth Set by the Half-Cycle

A fresh bed takes most of the carbon dioxide in the air through it, and its removal falls as it loads, so the cabin's carbon dioxide rises through each half-cycle and drops at each swap. With six crew at rest and the fixed 144 min half-cycle of case 3, the beds produce one tooth every 144.0 min, and between 12 and 24 h the partial pressure stays between 0.3713 and 0.3831 kPa with a mean of 0.3754 kPa. A longer half-cycle of 180 min raises the band to 0.4043 kPa to 0.4328 kPa (mean 0.4156), and a shorter one of 96 min lowers the mean to 0.3324 kPa at the price of 15.00 kWh of heater energy a day against 10.00. The tooth itself is small because the cabin holds 163 mol per kilopascal against roughly 13 mol produced in a half-cycle. A separate periodic calculation, two beds integrated by Euler's method at a frozen mean partial pressure, gives teeth of 0.0104, 0.0196 and 0.0315 kPa at 144, 180 and 216 min against the engine's 0.0102, 0.0190 and 0.0301 kPa on day 3 of case 3. At 180 min the teeth grow with the crew as well, to 0.0129, 0.0190 and 0.0255 kPa for four, six and eight people against 0.0131, 0.0196 and 0.0262 kPa by the same calculation.

10.4.3 Breakthrough and a Bed Swapped Too Soon

On a 240 min half-cycle a cool adsorbing bed passes half its inlet 102 min into its half-cycle, at a true loading of 0.7226 to 0.7234 across the step, against the analytic $\theta_b = 0.7227$. Case 4 puts the crew's exercise hour at 11:00, at the end of a half-cycle that swaps at 12:00. The carbon dioxide then rides to 0.4289 kPa, against 0.3865 kPa for the same day without exercise. Swapping at 10:54 meets the hour with a regenerated bed and lowers the peak to 0.4177 kPa, while a swap at 09:45, nine minutes after the last one, brings on a bed still hot and loaded and raises it to 0.4770 kPa. The floor on ε is what lets that hot bed give carbon dioxide back.

10.4.4 An Eclipse Budget That Does Not Close

A full battery holds 3.2 kWh above its floor, and the base load takes 2.456 kWh of it in each eclipse. A heater run through a whole eclipse adds 0.614 kWh, which still fits and leaves 61.6 % of a full battery. The oxygen loop's steady 1.176 kW for six crew, run through the eclipse, would take the battery to 60.3 % on its own and to 52.6 % with the heater. Full electrolysis in sunlight makes 10.10 mol of oxygen per orbit against the 9.58 mol that six crew and the leak consume, so the oxygen can wait for the sun, and once it does the heater can run in the dark. Sunlight with the heater and full electrolysis returns 3.61 kWh per orbit, more than an eclipse may take. In case 5 the adaptive schedule heats wherever its swaps fall and the oxygen loop runs the electrolysis at all hours, and the battery is below its floor at 1.478 h in the first eclipse, spends 169.7 min of the day below it and reaches 51.55 %.

10.5 The Control Layers as Implemented

The controller executes every minute ($\Delta t_c = 1$ min), on the analyzer's report period, in the order of work of Figure 1.10. The plant-side bookkeeping of each step covers events, swaps, airlock stages, clipping and protection. Every decision reads the measurement record of the previous subsection, the controls, the crew's declared timeline and the declared model's parameters, and an automated check in the application's suite scans the controller, the scheduler, the balance and the crew procedures for any read of the simulator's truth.

10.5.1 Three Proportional-Integral Loops

Each of the three loops runs the velocity-form PID of Section 1.6.2, Equations (1.11) to (1.14), with no derivative term, the limit on the emitted output only and the integral back-calculated against it with a tracking time T_t equal to τ_I. An integral time of zero turns the integral off. The oxygen loop sets P_e from the analyzer's oxygen with $K_c = 8\,\mathrm{kW/kPa}$ and $\tau_I = 4\,\mathrm{h}$ on 0 to 2 kW. The optional crew-activity feedforward is its output bias u_b,

$$u_{ff} = \frac{n_c\, r_{O_2}\, a(t) + C_{leak}\, p_{O_2,SP}}{\eta_e / F_e}, \tag{10.7}$$

taken from the crew's log and the exercise timetable. Toggling it shifts the integral by the feedforward's value, which keeps the output continuous. The pressure loop sets u_{N_2} with $K_c = 200\,\%/\mathrm{kPa}$ and $\tau_I = 6\,\mathrm{h}$ on an error that is zero inside $\pm\delta = \pm 0.2\,\mathrm{kPa}$ of the setpoint and is reduced by δ outside it, so the valve holds still inside the band. The humidity loop sets u_c with $K_c = 10\,\%$ per point of relative humidity and $\tau_I = 1\,\mathrm{h}$, direct acting, with the action $a = -1$ of Section 1.6.2. Every loop tracks its manual station in Manual, so a transfer is bumpless. The backup oxygen valve opens when the analyzer's oxygen falls 0.8 kPa under its setpoint and shuts at 0.4 kPa under it, in both modes.

10.5.2 Fixed and Adaptive Bed Schedules

The fixed schedule swaps the beds every $t_{1/2} = 144\,\mathrm{min}$ and heats the regenerating bed for the first $t_h = 60\,\mathrm{min}$ of each half-cycle, wherever the orbit is. Because the loadings are not measured, the adaptive schedule carries both predicted loadings forward each minute by the declared bed law of Equation (10.4), from the declared starting loadings, on the measured carbon dioxide and bed

temperatures, and swaps when the predicted $\hat{\theta}$ of the adsorbing bed reaches $\theta_{swap} = 0.60$, no sooner than 75 min and no later than 180 min after the last swap. A faulted bed is annunciated, so the predictor uses its halved capacity. A visiting-crew mode shortens the fixed half-cycle to 96 min while visitors are aboard.

10.5.3 A Power-Aware Program Over Three Orbits

The power-aware schedule keeps the adaptive swaps and places the heater runs and the electrolysis by a mixed-integer linear program over the next $H = 3T_{orb} = 4.6$ h. A forecast with the declared bed law and the crew aboard now predicts the swaps over the horizon, and the horizon is cut at every sunrise, sunset and predicted swap and then into slots k of at most 20 min, so each slot lies in one light and one window j between swaps. Per slot of length Δ_k the heater $x_k \in \{0,1\}$ is binary and the electrolysis fraction $w_k \in [0,1]$, the spilled array energy g_k, the floor shortfall σ_k, the tank oxygen o_k and the start indicator ζ_k are continuous, and each window carries a heating shortfall h_j. The program is Equation (10.8).

$$
\begin{aligned}
\min\ & \textstyle\sum_k \left(100\,\sigma_k + 2\,o_k + 0.05\,\zeta_k\right) + 50\sum_j h_j - E_K \\
& \quad + \textstyle\sum_k 0.01\,\Delta_k\,(t_k - t_{j(k)})\,x_k \\
\text{s.t. } & E_k = E_0 + \textstyle\sum_{i\le k}(\alpha_i - \beta_i x_i - \gamma_i w_i - g_i) \\
& E_k + \sigma_k \ge E_{fl} + 0.2, \qquad E_k \le E_{bat} \\
& n_k = n_0 + \textstyle\sum_{i\le k}(q_{max}\Delta_i w_i + o_i - \Delta_i \bar{D}_i) \\
& |n_k - n_{SP}| \le 0.25\,\kappa \\
& \textstyle\sum_{k\in j} \Delta_k x_k + h_j \ge t_{h,j}, \qquad \zeta_k \ge x_k - x_{k-1} \\
& \textstyle\sum_{k\in j} \zeta_k \le 2, \qquad x_k = 0 \text{ near a swap}
\end{aligned}
\tag{10.8}
$$

where $\alpha_k = \eta_c \Delta_k (P_{arr} - P_{base})$, $\beta_k = \eta_c \Delta_k P_h$ and $\gamma_k = \eta_c \Delta_k P_{e,max}$ in sunlight, $\alpha_k = -\Delta_k P_{base}/\eta_d$, $\beta_k = \Delta_k P_h/\eta_d$ and $\gamma_k = \Delta_k P_{e,max}/\eta_d$ in eclipse, $g_k = 0$ in eclipse, $E_{fl} = 4.8$ kWh is the floor, $q_{max} = 10.63$ mol/h, $\bar{D}_k$ is the crew timeline's oxygen use plus the declared leak at the slot's middle, $t_{h,j}$ is what window j still needs of its hour of heat, the heater is off for $t_{cool} = 12$ min before each swap, $t_{j(k)}$ is the start of slot k's window, and the oxygen at the horizon's end may sit no more than 0.05 kPa under the setpoint. The penalized slacks keep every program feasible, so a result reports how far the floor, the heating or the oxygen band is missed rather than failing. The battery model is linear because the sign of each slot's battery flow is known from the light. The same mix of on-off and continuous decisions on one shared supply recurs in the refrigeration pack of Volume 8, Chapter 12, whose hybrid predictive controller schedules the compressor starts.

Over the suggested week the program is solved 236 times, once at the start, 121 times at a swap, 108 times at sunrise and 6 times on a change of status, and every solve branches, with a median of 31 nodes. It is solved by depth-first branch and bound on the linear relaxation with the bounded simplex that the apps share, in the module that the hourly pump schedule of Volume 9, Chapter 11 first used, branching on the earliest fractional heater slot and exploring first the child nearer its fractional value. Any node whose relaxation is infeasible, or no better than the incumbent, is discarded. At a limit of 200 nodes the search returns its incumbent with the best open bound, and the scene says whether the plan was proved optimal or by how much it lies above that bound. It is re-solved at the start, at every swap, at every sunrise, when the crew aboard, a trip or a fault changes, and before the horizon runs out. In the eclipse the electrolysis runs at the plan's rate $w_k P_{e,max}$, which is both a cap and a floor, and in sunlight the plan's rate is a floor that the oxygen loop may exceed. The loop's integral is back-calculated against what the plan let through, and the backup valve delivers the plan's tank oxygen at the start of its slot.

10.5.4 The Nitrogen Balance and the Crew's Procedures

Nitrogen does not react, so it enters only from the tank and leaves only by the leak and the airlock, much as the inert of Chapter 8 leaves only by the purge (Section 1.4.3). Each hour the balance

Table 10.3 The Ten Cases and the Week

#	Case and events	Teaching point
1	Manual, every machine off, 48 h	the hours to each limit
2	Oxygen setpoint 21 to 22 kPa after the first half hour	loops coupled by the sum
3	Fixed half-cycle 144, then 180 or 96 min	the saw-tooth
4	Exercise hour at 11:00, a manual swap	timing the swap
5	Adaptive, then power-aware, 24 h	the eclipse budget
6	Excursion at 2 h, 12 h	the airlock's cost
7	Balance on, leak step at 24 h, 48 h	the balance finds the leak
8	Bed A faulted, two visitors at 2 h, adaptive	rationing capacity
9	Challenge, 168 h: excursions at 30 and 54 h, visitors 72 to 120 h, leak step at 126 h	the week aboard
10	Eight crew, bed A faulted, exercise, CO_2 from 0.45 kPa	infeasible at any schedule

averages the tank gauge T, the declared airlock loss A and the analyzer's inventory $\kappa p_{\mathrm{N}_2,m}$ over the hour and forms the residual and the tests of Equation (10.9).

$$
\begin{aligned}
r_k &= \frac{(\bar{T}_{k-1} - \bar{T}_k) - (\bar{A}_k - \bar{A}_{k-1}) - \kappa(\bar{p}_{N_2,k} - \bar{p}_{N_2,k-1})}{1\,\mathrm{h}} \\
\hat{L}_k &= \hat{L}_{k-1} + \lambda\,(r_k - \hat{L}_{k-1}) \\
S_k &= \max\left(0,\; S_{k-1} + (r_k - L_{ref,k-1})/\sigma_r - \tfrac{1}{2}\right)
\end{aligned}
\tag{10.9}
$$

where $\lambda = 0.15$, the declared $\sigma_r = 1.0\,\mathrm{mol/h}$ comes from the instruments' noise, and the reference L_{ref} follows the residual with the slow weight 0.05. When S passes $h = 5$ the test flags a step, estimates its size as the mean residual since S left zero less the reference, resets S and moves the reference by that step. The declared airlock loss is $(1-r)V_{al}Px_{N_2}/(RT)$ at the pressure and nitrogen fraction measured when the depressurization starts, spread over it. With "Isolate on a flag" on, the crew closes off the leaking volume two hours after a flag, a single search that removes a leak step if one is there. An excursion is a four-stage procedure, depressurize, outside, repressurize and idle, run on the clock.

10.5.5 Plant, Controller Model and Identified Model

The plant is everything in Section 10.3. The controller's model is the declared bed law and power arithmetic with the plant's own parameters, run on measurements, so the cases meet model error only through the hidden loadings, the noise and visitors the forecast cannot know about. A reader identifies the hours to each limit in case 1 and the tooth against the half-cycle in case 3.

10.6 Ten Cases and a Week, Measured

The cases are listed in Table 10.3. All share the starting state of Section 10.3, the loop tunings above and the fixed schedule unless a row says otherwise. Every number here was measured on the engine running headless and stepped as the browser steps it, on the case's own seed unless several are named, with any setting a sentence names made before Run. Limits are judged on the truth at every step.

10.6.1 The Airlock and the Leak

In case 6 the airlock vents 49.534 mol, of which 38.437 mol is nitrogen and 10.269 mol oxygen, equal to $(1-r)V_{al}P_0/(RT)$ at $P_0 = 101.298\,\mathrm{kPa}$. The pressure falls to 101.005 kPa at 2.733 h, just past the deadband, so the valve answers a small error slowly, and by 12 h the total is back only to 101.129 kPa while the tank has given 21.97 mol. Case 7 opens a leak step of $\Delta C_{leak} = 0.05$ at 24 h, which adds about 3.9 mol/h of nitrogen to the 0.93 mol/h the everyday leak takes. The pressure never leaves 101.001 kPa to 101.403 kPa over the two days, because the valve opens from 6.00 to 37.01 % at 48 h (39.89 % at its widest). The cumulative sum flags the step at 27.00 h, three hours after it opened, and at 48 h the estimate reads 4.689 mol/h against a true nitrogen leak of 4.841 mol/h. Over

Table 10.4 Measured Results of the Teaching Cases

Case	Configuration	CO_2 peak (kPa)	Battery low (%)	Below floor (min)
4, 24 h	no extra swap	0.4289	—	—
	swap at 10:54	0.4177	—	—
	swap at 09:45	0.4770	—	—
5, 24 h	adaptive	—	51.55	169.7
	power-aware	—	61.58	0.0
8, 24 h	fixed 144 min	0.5926	48.68	150.5
	adaptive	0.4738	49.18	225.3
	power-aware	0.4740	61.64	0.0
10, 48 h	best fixed (72 min)	0.565	—	—
	adaptive	0.5565	—	—
	power-aware	0.5585	—	—

the case's seed and ten others the flag comes 2.0 to 3.0 h after the step with no flag before it, the estimate is within 0.27 mol/h of the truth at 48 h, and the hourly residual scatters about the true leak with a standard deviation of 0.71 to 1.05 mol/h on the first day, against the declared 1.0 mol/h.

Pause. Try Exercise 10.2 before reading on.

> **TRY IT IN THE APP** **7 · A Leak Step and the Nitrogen Balance**
>
> Load the scenario, press Run at 1 h/s on the Module view and watch the total pressure and the Nitrogen valve strips after the leak opens at 24 h. The pressure never leaves 101.00 to 101.40 kPa, while the valve opens from 6.0 to 37.0 % to feed the leak. Only the nitrogen balance panel shows the step, with its flag at 27.00 h.

10.6.2 Schedules Against the Eclipse

The adaptive and power-aware schedules of case 5 make the same 17 swaps in the day. What the program moves is the electrolysis. Over the day it puts 0.86 kWh of its 28.98 kWh of electrolysis in the eclipses, against 11.23 kWh of 30.23 kWh on the adaptive schedule, while its heater still runs for 380 min of the 528 min of eclipse (402 min on the adaptive schedule), as the budget above allows. The battery's lowest point rises from 51.55 to 61.58 % with no time below the floor, as shown in Figure 10.2. Over the suggested week the same pattern holds, with 14.62 kWh of electrolysis in the eclipses against 88.65 kWh as shipped. Case 8 adds a faulted bed and two visitors, and the three schedules each fail differently, as listed in Table 10.4. The fixed half-cycle is too long for eight people on one and a half beds, and its carbon dioxide peaks at 0.5926 kPa, over the limit for 519.5 min. The adaptive schedule shortens its half-cycles and holds the carbon dioxide at 0.4738 kPa, but its 18 regenerations heat in the eclipses and the battery spends 225.3 min below its floor. The power-aware program keeps the same 18 swaps and 0.4740 kPa and holds the battery at 61.64 % or more. Figure 10.3 draws the peaks and the battery lows of Table 10.4 side by side.

10.6.3 An Honest Infeasibility

Case 10 puts eight crew on one healthy and one faulted bed with the daily exercise hour, from a steady 0.45 kPa. Every fixed half-cycle from 36 to 240 min in 12 min steps breaks the 0.53 kPa limit, the best at 72 min with a peak of 0.565 kPa, because short half-cycles bring the beds back hot and long ones let them saturate. The adaptive schedule does best and still passes the limit in the first exercise hour at 10.84 h and peaks at 0.5565 kPa, and the power-aware program peaks at 0.5585 kPa. Eight people exercising produce 18.0 mol/h of carbon dioxide, while a fresh cool bed at 0.5 kPa takes at most 10.29 mol/h, so the answer lies in the demand (staggered exercise, a repaired bed, a postponed visit) and not in the controller.

Pause. Try Exercise 10.3 before reading on.

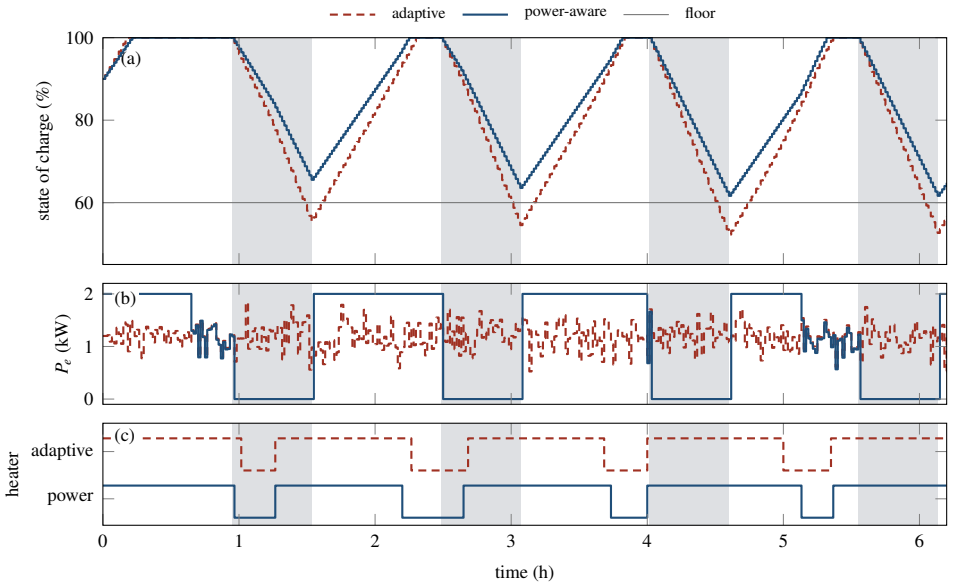

Figure 10.2 The first four orbits of case 5 on the adaptive (dashed) and the power-aware (solid) bed schedules, from the engine, eclipses shaded. (a) True state of charge against the 60 % floor (thin gray line). (b) Electrolysis power. (c) Heater on or off, the adaptive run on the upper row. The adaptive run breaches the floor at 1.478 h, while the program keeps the electrolysis out of the dark and lets the heater run there.

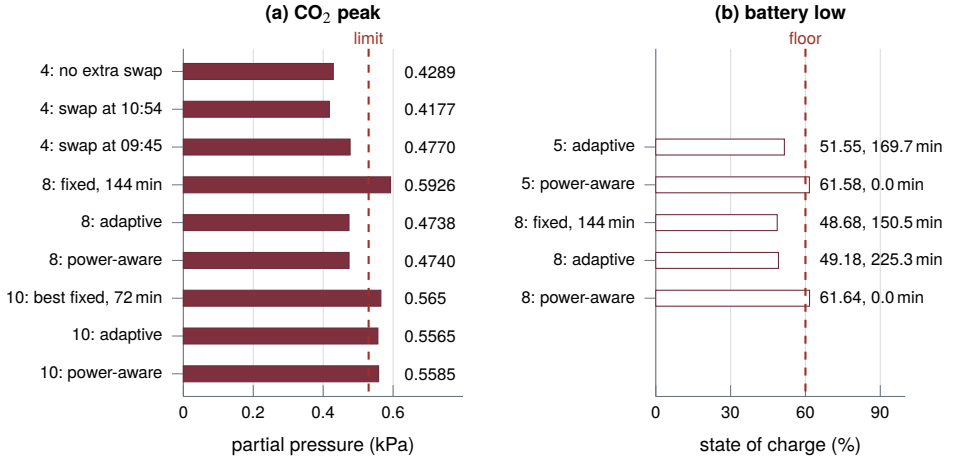

Figure 10.3 The teaching cases of Table 10.4: (a) the carbon dioxide peak against the 0.53 kPa teaching limit (dashed), (b) the battery's lowest state of charge against its 60 % floor (dashed), with the minutes spent below the floor. Notice that in case 8 the power-aware program keeps the adaptive schedule's carbon dioxide and alone holds the battery, and that in case 10 no schedule stays under the limit.

10.6.4 The Week Aboard

The challenge runs 168 h with six crew and their exercise hour at 10:00, excursions at 30 and 54 h, two visitors from 72 to 120 h and a leak step at 126 h. Its target is every atmospheric limit held on the truth, the battery never below its floor, both tanks above their reserves at the end (575 mol

Table 10.5 The Week Over the Case's Seed and Ten Others, Judged on the Truth

Configuration	Pass	CO_2 (kPa)	Battery (%)	N_2 left (mol)	Flag (h)
shipped	0/11	0.564	46.51–46.87	520–522	—
suggested	11/11	0.493–0.496	61.23–61.29	646–666	2.0–3.0
no feedforward	11/11	0.493–0.496	61.23–61.29	646–668	2.0–3.0
no balance	0/11	0.493–0.496	61.23	484–503	—
isolation off	0/11	0.493–0.496	61.23	484–503	2.0–3.0
adaptive	0/11	0.489–0.491	46.29–47.68	666–674	2.0–3.0
fixed, visit mode	0/11	0.497	47.40–48.05	683–690	2.0–3.0

of nitrogen and 300 mol of oxygen) and the leak step flagged within 12 h. Table 10.5 gives seven configurations over the case's seed and ten others, judged from the logged truth. As shipped (fixed schedule, no feedforward, no balance) the week fails four ways, with the carbon dioxide over its limit for 494 min while the visitors are aboard (worst 0.564 kPa), the battery below its floor for 901 min, the leak never flagged and 520 mol of nitrogen left. The suggested settings (power-aware, feedforward, balance) pass every seed with at least 0.034 kPa, 1.23 points of charge, 71 mol of nitrogen and 9 h of margin. The schedule and the balance are each necessary, while the feedforward is not. Without the balance the same schedule never flags the leak and ends with 484 to 503 mol. Under the program the electrolysis runs at the plan's rate, so the week's largest oxygen dip in an exercise hour is 0.1204 kPa without the feedforward and 0.1194 kPa with it. Where the loop owns the electrolysis, as in case 4, the feedforward roughly halves the dip, from 0.047 to 0.027 kPa.

The reserve of 575 mol sits at the midpoint of the measured gap between weeks whose leak is isolated (646 to 666 mol) and weeks whose leak is left open (484 to 503 mol), so the verdict turns on the response by at least 71 mol at every seed. The carbon dioxide limit of 0.53 kPa is a teaching value. Under the standard's one-hour average of 0.40 kPa [4] the suggested week would fail, since its highest one-hour average is 0.490 kPa, and even the six-person design point of case 3 averages 0.375 kPa. The crews of this simulator would not meet the real limit through a week. The week at 129.5 h, half an hour after the balance flagged the leak, is shown in Figure 10.4.

✦ DID YOU KNOW? A Leak Found but Not Stopped

From the simulator (Table 10.5). With Isolate on a flag off, the balance still flags the week's leak 2.0 to 3.0 h after it opens, yet the week fails on all eleven seeds. It ends with 484 to 503 mol of nitrogen, exactly as with no balance at all, against 646 to 666 mol when the crew isolate the leak.

10.6.5 The Program's Proofs and Gaps

Over the suggested week the program is solved 236 times, and 230 solves are proved optimal within the 200-node limit. The other six return an incumbent up to 20.07 above the best open bound in the objective's units, which the Orbit view prints as "not proved optimal". Searched to the end, each of the six incumbents is the optimum, reached after 205 to 311 nodes, so the large gaps measure a weak relaxation rather than a poor plan. At 500 and 5,000 nodes every solve is proved, and the eleven seeds give the same verdicts and ranges as at 200. At the earlier limit of 60 nodes, 67 of 237 solves stopped short.

10.6.6 Code Verification

Verification is kept apart from agreement with published data. Counting atoms through the electrolysis and the Sabatier stoichiometry, the nitrogen, carbon, hydrogen and oxygen held in the cabin, the beds, the tanks and the ledgers close to within 2.2×10^{-9} mol at every step of the shipped and suggested weeks and of cases 1, 6, 8 and 10. The total pressure equals $\sum_i n_i / \kappa$ to 4.3×10^{-14} kPa. Each species balances to its sources and sinks within 8.8×10^{-10} mol over the first 72 h of the suggested week.

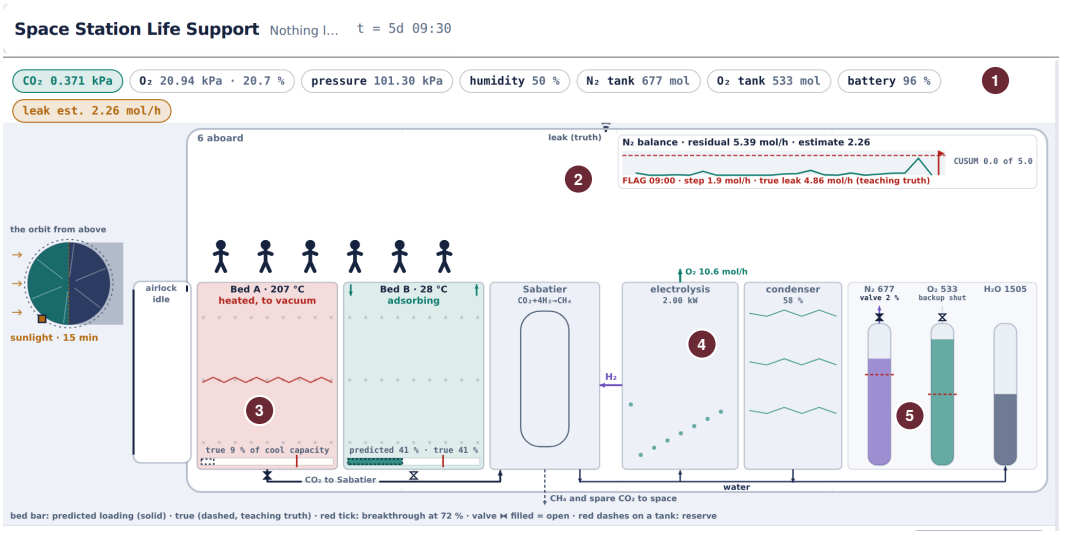

Figure 10.4 The week aboard with the suggested settings at $t = 129.5$ h (day 6, 09:30) on the Module view, in sunlight with 15 min to sunset, a composite of the app's header and its Module view. The leak step opened at 126 h, and the balance panel (2) shows the flag raised at 09:00 with an estimated step of 1.9 mol/h, a last hourly residual of 5.39 mol/h, an estimate of 2.26 mol/h against a true nitrogen leak of 4.86 mol/h (teaching truth) and the cumulative sum reset to 0.0 of 5.0. The crew isolate the leak at 131.0 h. Six are aboard, bed A (3) is heated to 207 °C under vacuum at 9 % of its cool capacity, and bed B adsorbs at 28 °C with a predicted and a true loading of 41 %. The badges (1) read 0.371 kPa of carbon dioxide, 20.94 kPa of oxygen (20.7 %), 101.30 kPa, 50 % humidity, 677 mol of nitrogen, 533 mol of oxygen in the tanks and a battery at 96 %. The electrolysis (4) and the nitrogen tank with its makeup valve (5) are the final elements of the oxygen and pressure loops. On the app's carbon dioxide trend, cut from this picture, the visitors' stay shows over the 48-hour window, and their exercise hour on day 5 sets the week's highest value, 0.494 kPa at 107.0 h.

The airlock loss matches its declaration, breakthrough matches its analytic loading and the saw-tooth matches the independent periodic calculation within 5 %, as set out in Section 10.4. The application's own suite checks the branch and bound against exhaustive enumeration on small instances. Halving the step moves the week's nitrogen losses (leak and airlock) by 0.0014 % as shipped and 0.016 % suggested, and the shipped, adaptive, visiting-mode and no-feedforward weeks rerun at half the step keep their verdicts. The tank's draw moves 0.72 % with the suggested settings, which fails a literal reading of a 0.5 % tolerance on nitrogen use. The difference, 1.7 mol, is the cabin's nitrogen left at a different point inside the pressure deadband of ± 0.2 kPa (± 33 mol), since the draw less the losses less the change in the cabin is zero at both steps.

10.6.7 Agreement With Published Values

Agreement is limited to declared values. The crew's day with one exercise hour is within 3 % of the baseline's 0.84, 1.01 and 2.27 kg [3], the half-cycle of the fixed schedule is the 144 min of the tested assembly [5], and F_e agrees within 0.004 % with the 0.131 745 kWh/mol computed from the reference enthalpy of formation and entropies of liquid water, hydrogen and oxygen [15]. The leak is ten times the baseline's upper module leakage on purpose, so that the balance has something to find within a day, and the bed, battery and electrolysis figures are teaching values. No measured trajectory of a flight system was compared.

10.6.8 Runtime

A week takes 60 480 steps. In a browser engine on one host shared with other work, three suggested weeks at each of two window sizes put the median program solve at 25.6 to 26.4 ms. The slowest solve of a week was 149 to 196 ms in four weeks and 274 and 462 ms in the other two. Both outliers are the solve at 19.28 h, which took 129 to 149 ms in three of the other weeks, so they come from the host rather than the problem. A week took 42.0 to 42.1 s shipped and 43.2 to 43.9 s suggested at the top speed of 4 h/s, bounds on one machine rather than benchmarks. A slow solve stalls one frame and changes no result.

10.7 Technology in Practice

What Measures It

Every loop of this chapter reads one gas analyzer and one pressure transducer. On the current station the analyzer is a mass spectrometer, the primary resource for keeping oxygen and carbon dioxide at safe levels, and its partial pressures are normalized to the total of the compartment sampled [16]. The sum that couples the two gas loops of Section 10.4 is thus built into the instrument as well. The model's analyzer reports four partial pressures once a minute with white noise of 0.03, 0.03, 0.0015 and 0.01 kPa and holds each report, and it neither drifts nor fails. A real one is recalibrated against a verification gas and is replaced when its ion pump wears out [16], so separate analyzers for single gases are kept beside it for redundancy [1].

> **TECHNOLOGY Sensor: The Magnetic-Sector Mass Spectrometer**
>
> On the current station a mass-spectrometer-based analyzer monitors nitrogen, oxygen, carbon dioxide, methane, hydrogen and water vapor. Sampled gas passes an inlet into an ion source, a single-focusing magnetic sector spreads the ions, and six electrometers in a spatial array collect them while an ion pump holds the vacuum. The nitrogen reading is used to watch the cabin for air leakage. The partial pressures are normalized to sum to the total, so an error in one gas appears as compensating errors in the others. A verification gas recalibrates the analyzer in orbit. As the ion pump ages its current rises, and at a set current the analyzer protects its filaments and shuts down [16]. A mass spectrometer cannot tell apart two compounds of the same molecular weight, such as nitrogen and carbon monoxide [1].

What Moves It

The oxygen loop moves the electrolysis power, the pressure loop the nitrogen makeup valve, the humidity loop the condenser's air flow, and the bed schedule the swaps and the heater, as Figure 10.5 draws them. In the model a swap takes effect at a step boundary and the beds warm and cool by first-order lags. On the real assembly the selector valves, the bed heaters, the air-save pump and the blower have their own controllers [5], and its air flow, half-cycle and heater set point trade removal against power [17], the trade that Section 10.4 measures as heater energy against the tooth.

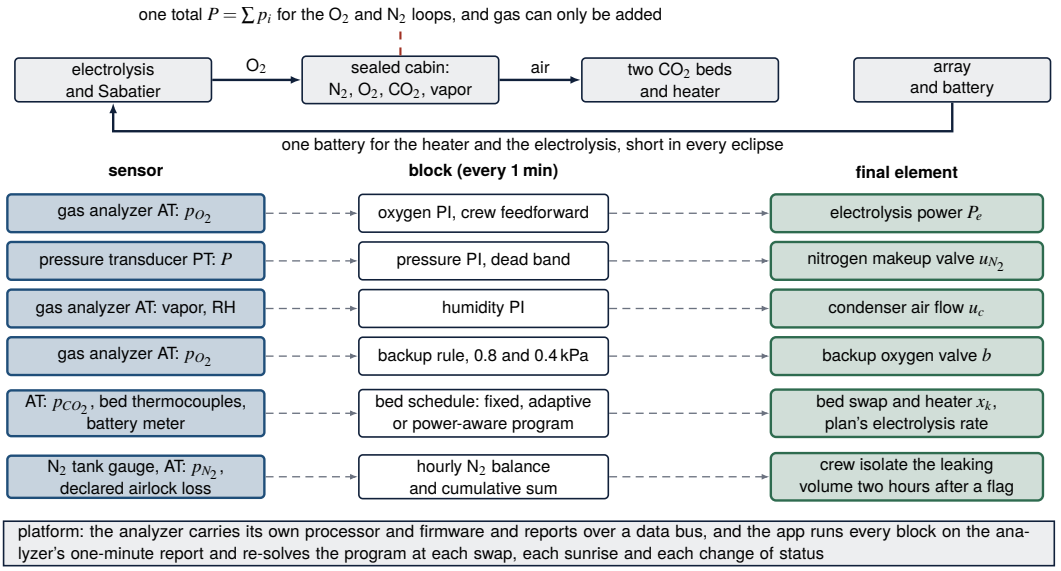

Figure 10.5 The loops of Section 10.5 on real hardware, sensor to block to final element, below the shared quantities that couple them. Notice that one analyzer serves five of the six rows, and that the only path that sees the leak runs through the nitrogen balance, not through the pressure loop that feeds it.

> **TECHNOLOGY** **Final Element: The Four-Bed Removal Assembly**
>
> The removal assembly of the current station holds two desiccant and two sorbent beds. One of each adsorbs while the other two desorb, and selector valves switch them at each half-cycle. Its air comes from the outlet of a condensing heat exchanger, cool and dry, which favors removal. Integral heaters warm the desorbing sorbent bed under space vacuum, and for the first 15 min of each half-cycle a pump returns its residual air to the cabin. The vent valves to vacuum close on a power failure [5]. Cutting the nominal cycle time of 144 min to 90 min is near the limit that still lets the beds heat to their 204 °C set point. The current revision came from an investigation of an anomaly in orbit and brought redesigned heaters, screens serviceable in orbit and new sorbents [17].

Where the Controller Runs

The station's analyzer carries its own processor, memory and firmware, reports to the station's command and data handling system over a 1553B data bus, and a built-in test in its firmware shuts it down when its internal transducers read wrong [16]. A design reference asks for a system controller that adjusts the operation automatically when more people come aboard, and for monitoring that alerts the crew to a failure [1]. The app executes its three PI loops, the bed schedule and the nitrogen balance every minute, on the analyzer's report period, over a 10 s integration step, and re-solves the power-aware program at each swap, each sunrise and each change of status. The plantwide structure sits above the loops: a schedule that shares the battery between the heater and the electrolysis, and a balance that watches what the pressure loop feeds.

What Is Modernizing It

Three developments bear on this chapter. New carbon dioxide scrubbers are tested in orbit: a thermal amine scrubber launched in 2019 has removed four crew members' carbon dioxide at slightly above the target cabin level, and it is to join a four-bed scrubber in a selection for long-duration testing [14]. A virtual test bed that solves the coupled heat transfer, mass transfer and low-pressure flow of an adsorption assembly tries many combinations of flow, cycle time, heater ramp rate and set point

before ground hardware confirms the best [17], the role the declared bed law plays for the adaptive and power-aware schedules here. Model-based reasoning compares sensor data with the predictions of a computer model of the system to locate a problem and alert the crew [1], which the nitrogen balance does for one species.

> **IN PRACTICE** **Atmosphere Requirements Taken as a Whole**
>
> The human-system standard for spaceflight, NASA-STD-3001 Volume 2, sets pressure, humidity, temperature, ventilation, oxygen and carbon dioxide one by one, then states that they must be considered as a whole because they depend on each other. It asks for local and remote control of pressure, humidity, temperature, ventilation and oxygen [V2 6017], for environmental control that meets the added oxygen use and carbon dioxide of exercise [V2 7041], for automatic, continuous recording of pressure, humidity, temperature and both partial pressures in every compartment that a hatch can isolate [V2 6020], and for alerts to the crew, locally and remotely, when any of them leaves its safe limits [V2 6022] [4].

10.8 Limitations

The cabin is well mixed and isothermal, so ventilation dead zones, trace contaminants and thermal control are outside the model. The beds are lumped and the desiccant half of a real assembly is absent, so the tooth and the cost of a hot swap are plausible for a teaching bed rather than predicted for a flight assembly. The leak is ten times a module's published upper value, the crews here would not meet the standard's carbon dioxide limit, and the battery, array, electrolysis and Sabatier figures are teaching values.

> **OPEN CHALLENGE** **Lower Carbon Dioxide Without More Power**
>
> Crew members on the current station have had headaches at its carbon dioxide levels, and a future removal system must hold 2 Torr with four crew. A ground test met that level with more air flow and shorter half-cycles, but the heater power alone rose by 200 W on average, and power is at a premium, the more so on missions that cannot draw on the sun [17]. A key performance parameter now asks for 2600 ppm of carbon dioxide with four crew [14]. Progress would be the flow, cycle time and heater set point with the best removal at the least power, found on a virtual test bed and confirmed on ground hardware [17].

The program's battery model is linear, its heater runs whole slots and its swap forecast does not know about visitors before they arrive. Its node limit leaves six of 236 solves unproved in the suggested week. The controller's model is the plant's own law with the plant's parameters, so the cases meet model error only through hidden states and noise. The balance's residual scatters by about 1 mol/h, so a leak step much smaller than that is found late or not at all. A simulated pass is a statement about this teaching model and not a certification of any design.

> **OPEN CHALLENGE** **Recovering More of the Oxygen**
>
> For long transit missions the target is to recover more than 75 percent of the oxygen in exhaled carbon dioxide, and more than 90 percent on a planetary surface. The paths under development improve the Sabatier reactor, recover the hydrogen that leaves in its methane by pyrolysis and return it to make more water for oxygen, or replace the reactor with more efficient Bosch technologies. On a transit to Mars every spare must launch with the vehicle, so reliability becomes paramount [14]. In this chapter's model the methane, the spare hydrogen and the carbon dioxide the reactor cannot take are all vented.

10.9 Conclusion

A sealed cabin is an inventory that machines can only add to. The simulator of this chapter integrates its four gases, two cyclic beds, oxygen and Sabatier chain, airlock, leak and battery with element balances that close to 2.2×10^{-9} mol. With nothing running, the humidity and the carbon dioxide reach their limits in hours and the oxygen in more than a day. Raising the oxygen setpoint by 1 kPa

shuts the nitrogen valve 2.52 h later because both gases share one total. The beds' tooth grows with the half-cycle, and the eclipse takes the battery below its floor at 1.478 h unless a three-orbit program keeps the electrolysis out of the dark, which holds it at 61.58 %. A pressure loop feeds a new leak inside its band while a nitrogen balance flags it in three hours. Over a week aboard, the power-aware schedule and the balance are each necessary and together pass eleven seeds out of eleven, while the feedforward changes little. The bed model is lumped and the limits are teaching values. A fuller model would add a desiccant pair, a distributed bed, a thermal loop and a stronger relaxation for the program. In a sealed system the schedule, the loops and the ledger are one control problem.

KEY POINTS

- A sealed cabin is an inventory that machines can only add to: nothing but the crew and the leak removes an excess of gas.
- Raising the oxygen setpoint by 1 kPa shuts the nitrogen valve 2.52 h later, because both gases share one total pressure.
- A pressure loop feeds a new leak inside its band, while an hourly nitrogen balance flags it in three hours.
- The eclipse takes the battery below its floor unless a three-orbit program keeps the electrolysis out of the dark, which holds it at 61.58 %.
- Over a week aboard, the power-aware schedule and the balance are each necessary, while the feedforward changes little.

Where else this applies. The gas transmission line of Volume 10, Chapter 12 finds its leaks with a balance of the same kind, metered inflow less metered outflow less the change in a line pack computed from three pressure transmitters.

References

[1] P. O. Wieland, "Designing for human presence in space: An introduction to environmental control and life support systems," National Aeronautics and Space Administration, George C. Marshall Space Flight Center, Marshall Space Flight Center, AL, NASA Reference Publication RP-1324, 1994. [Online]. Available: https://ntrs.nasa.gov/citations/19940022934

[2] P. Eckart, *Spaceflight Life Support and Biospherics*. Springer Netherlands, 1996.

[3] M. K. Ewert, T. T. Chen, and C. D. Powell, Eds., *Life Support Baseline Values and Assumptions Document*, ser. NASA Technical Publication. Houston, TX: National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, Feb. 2022, no. NASA/TP-2015-218570/REV2. [Online]. Available: https://ntrs.nasa.gov/citations/20210024855

[4] National Aeronautics and Space Administration, "NASA spaceflight human-system standard, volume 2: Human factors, habitability, and environmental health," National Aeronautics and Space Administration, Office of the Chief Health and Medical Officer, Technical Standard NASA-STD-3001, Volume 2, Revision E, 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250004555

[5] J. C. Knox, "International space station carbon dioxide removal assembly testing," in *SAE Technical Paper Series*, no. 2000-01-2345. SAE International, 2000.

[6] H. Mohamadinejad, J. C. Knox, and J. E. Smith, "Experimental and numerical investigation of adsorption/desorption in packed sorption beds under ideal and nonideal flows," *Separation Science and Technology*, vol. 35, no. 1, pp. 1–22, 2000.

[7] D. Samplatsky, K. Grohs, M. Edeen, J. Crusan, and R. Burkey, "Development and integration of the flight Sabatier assembly on the ISS," in *41st International Conference on Environmental Systems*. American Institute of Aeronautics and Astronautics, 2011, paper AIAA 2011-5151.

[8] J. Law, M. Van Baalen, M. Foy, S. S. Mason, C. Mendez, M. L. Wear, V. E. Meyers, and D. Alexander, "Relationship between carbon dioxide levels and reported headaches on the International Space Station," *Journal of Occupational & Environmental Medicine*, vol. 56, no. 5, pp. 477–483, 2014.

[9] S. Sircar, "Pressure swing adsorption," *Industrial & Engineering Chemistry Research*, vol. 41, no. 6, pp. 1389–1392, 2002.

[10] M. Baldea and I. Harjunkoski, "Integrated production scheduling and process control: A systematic review," *Computers & Chemical Engineering*, vol. 71, pp. 377–390, 2014.

[11] A. H. Land and A. G. Doig, "An automatic method of solving discrete programming problems," *Econometrica*, vol. 28, no. 3, p. 497, 1960.

[12] R. S. Mah, G. M. Stanley, and D. M. Downing, "Reconcillation and rectification of process flow and inventory data," *Industrial & Engineering Chemistry Process Design and Development*, vol. 15, no. 1, pp. 175–183, 1976.

[13] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1-2, pp. 100–115, 1954.

[14] J. L. Broyan, Jr., L. Shaw, M. McKinley, C. Meyer, M. K. Ewert, W. F. Schneider, M. Meyer, G. A. Ruff, A. C. Owens, and R. L. Gatens, "NASA environmental control and life support technology development for exploration: 2020 to 2021 overview," in *50th International Conference on Environmental Systems*, no. ICES-2021-384, 2021. [Online]. Available: https://ntrs.nasa.gov/citations/20210010866

[15] P. Linstrom, "NIST chemistry WebBook, NIST standard reference database 69," National Institute of Standards and Technology, 1997.

[16] B. D. Gardner, P. M. Erwin, T. Cougar, and B. Ulrich, "International space station major constituent analyzer on-orbit performance," in *47th International Conference on Environmental Systems*, no. ICES-2017-17, Charleston, SC, 2017. [Online]. Available: https://ntrs.nasa.gov/citations/20170005402

[17] J. C. Knox and C. M. Stanley, "Optimization of the carbon dioxide removal assembly (CDRA-4EU) in support of the international space system and advanced exploration systems," in *45th International Conference on Environmental Systems*, no. ICES-2015-165, Bellevue, WA, 2015. [Online]. Available: https://ntrs.nasa.gov/citations/20150016500

10.10 Exercises

10.1 Concept check · 2 min

In case 2 the oxygen setpoint is raised from 21.0 to 22.0 kPa at 1 h, and the nitrogen valve, open 6 % until then, shuts 2.52 h later. Why does it shut?

(a) The pressure loop is cascaded under the oxygen loop and takes its setpoint from the oxygen error

(b) The added oxygen lifts the total pressure past the deadband, and the pressure loop answers the total

(c) The backup oxygen tank opens and shares a supply line with the makeup valve, which it closes

(d) The analyzer reads nitrogen low when oxygen rises, so the valve answers a false reading

10.2 Concept check · 1 min

Case 7 opens a leak step at 24 h that adds about 3.9 mol/h of nitrogen loss, yet the total pressure stays between 101.001 and 101.403 kPa for the two days. What does this show?

(a) The leak is too small to move the total by more than the transducer's noise

(b) The deadband absorbs the loss, so the valve never has to move

(c) The makeup valve replaces what the leak takes, so the pressure loop feeds the leak and hides it

(d) The leak takes oxygen first, which the electrolysis replaces before the total can move

10.3 Concept check · 2 min

In case 10 eight crew live on one healthy and one faulted bed with a daily exercise hour. Every fixed half-cycle from 36 to 240 min breaks the 0.53 kPa carbon dioxide limit, and so do the adaptive and the power-aware schedules. What is the right conclusion?

(a) A half-cycle between two of the swept values, finer than the 12 min step, would hold the limit

(b) The power-aware program needs a larger node limit to find a schedule that holds the limit

(c) The crew-activity feedforward should be on, so the exercise hour is met before the peak

(d) The exercising crew make more carbon dioxide than a fresh bed can take, so the demand must change

10.4 Theory and calculation · 6 min

The cabin of Equation (10.1) holds 400 $\mathrm{m^3}$ at 22 °C. At the start the oxygen is 21.0 kPa and the total 101.3 kPa. Take the gas constant as $R = 8.314\,462\,618 \times 10^{-3}\ \mathrm{kPa\,m^3/(mol\,K)}$.

(a) Find $\kappa = V/(RT)$ and the rise in total pressure for one mole of any gas added. *Checkpoint:* 163.00 mol/kPa and 0.006 135 kPa.

(b) The oxygen fraction is $x = p_{O_2}/P$. Show that one mole of oxygen added raises it by $(P - p_{O_2})/(\kappa P^2)$, and evaluate this in percentage points. *Checkpoint:* 0.004 801 points per mole.

(c) Write the steady gain matrix from the moles of oxygen and nitrogen added, $(\Delta n_{O_2}, \Delta n_{N_2})$, to $(\Delta p_{O_2}, \Delta P)$. Find the relative gain λ_{11} of the pairing oxygen to p_{O_2} and nitrogen to P, and the condition number of the matrix. *Checkpoint:* $\kappa^{-1}\left[\begin{smallmatrix}1 & 0\\ 1 & 1\end{smallmatrix}\right]$, $\lambda_{11} = 1$, $\gamma = 2.62$.

10.5 Theory and calculation · 6 min

Use the power model of Equation (10.6): the eclipse lasts 35 min of a 92 min orbit, the base load is 4 kW, the heater 1 kW, the array 11 kW, the battery 8 kWh with a 60 % floor, and $\eta_c = \eta_d = 0.95$.

(a) Starting full, find the energy the base load draws in one eclipse and the state of charge after an eclipse with the heater on throughout. *Checkpoint:* 2.456 kWh of the 3.2 kWh above the floor; 61.6 %.

(b) The oxygen loop's steady load replaces six resting crew (1.0 mol/h each) and the oxygen in the leak, $C_{leak}\, p_{O_2,SP}$ with $C_{leak} = 0.0118$ mol/(h kPa). With $F_e = 0.131\,74$ kWh/mol and $\eta_e = 0.70$, find that load and the state of charge after a full eclipse of it, without and with the heater. *Checkpoint:* 1.176 kW; 60.3 % and 52.6 %.

(c) Find the oxygen that full electrolysis (2 kW) makes in the sunlit part of one orbit, and compare it with what the crew and the leak use in a whole orbit. *Checkpoint:* 10.10 mol against 9.58 mol.

10.6 Theory and calculation · 6 min

Use the bed law of Equation (10.4) with $N = 2.5$, $F_b = 55\,\mathrm{m^3/h}$ and $c = 0.407\,50\,\mathrm{mol/(m^3\,kPa)}$, and the balance of Equation (10.9) with $\sigma_r = 1.0$ mol/h and $h = 5$.

(a) Find the breakthrough loading θ_b and the removal efficiency ε of a fresh bed ($\theta = 0$). *Checkpoint:* $\theta_b = 0.7227$, $\varepsilon = 0.918$.

(b) Find the largest removal of a fresh cool bed at $p_{CO_2} = 0.5$ kPa and compare it with eight crew exercising at $a = 2.5$, each giving 0.9 mol/h of carbon dioxide at rest. *Checkpoint:* 10.29 mol/h against 18.0 mol/h.

(c) A leak step raises every hourly residual by 3.9 mol/h above the reference, with no noise. How many full hours of the step does the cumulative sum need to pass h? *Checkpoint:* Two hours ($S = 3.4$, then 6.8).

10.7 App challenge · 8 min

Load scenario 1, Nothing runs: the hours to each limit, and change no setting. Use the default seed and 1 h/s. The run pauses when a limit has been out for an hour. Press Run each time until 48 h.

(a) Before running, estimate from the crew's rates and $\kappa = 163.00$ mol/kPa the hours until the humidity passes 60 % (from 50 % of 2.638 kPa), the carbon dioxide passes 0.53 kPa (from 0.37) and the oxygen falls below 19.5 kPa (from 21.0). Ignore the leak. *Checkpoint:* 1.46, 4.83 and 40.75 h.

(b) Export the file and find the first row past each limit in the columns `truth_rh_pct`, `truth_co2_kPa` and `truth_o2_kPa`. Explain the oxygen's difference from (a). *Checkpoint:* 1.47, 4.85 and 39.20 h, within 0.02 h.

(c) Read `truth_p_kPa`, `truth_co2_kPa`, `truth_o2_kPa` and `truth_n2_kPa` at 48 h, and account for the change in the total since the start. *Checkpoint:* 102.13 kPa: carbon dioxide +1.59, oxygen −1.84, nitrogen −0.27, water vapor about +1.36.

10.8 App challenge · 8 min

Load scenario 2, The oxygen loop and the pressure it shares, at the default seed and the default 0.25 h/s. Pause between 0.95 and 1.05 h, set the O_2 setpoint to 22.0 kPa, and run to 24 h.

(a) In the exported file, read `n2_valve_pct` just before the step, and find how long after the step the valve first reads zero. *Checkpoint:* 6 %; 2.53 h, within 0.05 h.

(b) At 24 h read `truth_p_kPa` and `truth_frac_pct`, and the oxygen tank used on the Consumables tile. *Checkpoint:* 102.16 kPa and 21.51 %, within 0.02; 58 mol.

(c) With the valve shut, the total falls only as the everyday leak removes nitrogen, about 0.93 mol/h. Estimate from the total in (b) when the pressure will reach the lower edge of the deadband, 101.1 kPa, so that the valve can open again. *Checkpoint:* About 186 h after 24 h, near 210 h.

10.9 App challenge · 9 min

Load the challenge, A week aboard (case 9 of Table 10.3), at the default seed. Run each week to 168 h at 4 h/s, about 42 s, and read the Challenge target tile at the end.

(a) Run the week as shipped. List the requirements it fails. *Checkpoint:* Carbon dioxide over the limit 494 min (worst 0.564 kPa), the battery below its floor 901 min, 520 mol of nitrogen, under the reserve, and the leak not flagged in time.

(b) Set the Bed schedule to Power-aware, the Crew-activity feedforward On and the Nitrogen balance test On, press Reset (which keeps the settings) and run the week again. *Checkpoint:* Met: carbon dioxide worst 0.494 kPa, battery lowest 61 %, 664 mol of nitrogen, the leak flagged after 3.0 h.

(c) From the settings of (b), turn the Nitrogen balance test Off, Reset and run. Then turn it On, turn the feedforward Off, Reset and run. Which of the two layers is necessary? *Checkpoint:* Balance off: 495 mol, under the reserve, not met. Feedforward off: met with 658 mol.

10.10 Application area · 3 min

The challenge limits the carbon dioxide to 0.53 kPa at every instant, a teaching value. The standard limits the one-hour average to 3 mmHg, and an analysis of flight reports found that the odds of a headache double for each 1 mmHg. Take $1\,\text{mmHg} = 0.133\,32\,\text{kPa}$.

(a) Convert the standard's limit to kilopascals, and the teaching limit and the suggested week's highest one-hour average, 0.490 kPa, to mmHg. *Checkpoint:* 0.400 kPa; 3.98 and 3.68 mmHg.

(b) By the headache relation, how much higher are the odds at the suggested week's highest hour than at the standard's limit? Case 3's six crew at rest average 0.375 kPa. What does that say about where a remedy must come from? *Checkpoint:* About 1.6 times; 2.81 mmHg, above a seven-day 2.5 mmHg.

10.11 Application area · 4 min

The everyday leak is $C_{leak}P$ with $C_{leak} = 0.0118\,\text{mol}/(\text{h}\,\text{kPa})$ at 101.3 kPa, and it takes the cabin's mixture. The nitrogen tank holds 900 mol at the start with a reserve of 575 mol.

(a) Find the air lost per day in moles and in kilograms, with a molar mass of the cabin air of 28.77 g/mol, and compare it with a module leakage of 0.09 kg/day. *Checkpoint:* 28.7 mol, 0.83 kg, about nine times.

(b) With the nitrogen part of the leak at 0.93 mol/h, how many days of everyday leak does the tank hold above its reserve, and how many hours with the case 7 leak step open (4.84 mol/h in all)? Each excursion vents 38.4 mol of nitrogen. *Checkpoint:* About 14.6 days; about 67 h.

10.12 Review the agent · 5 min

An AI assistant summarizes case 5 for a design review: “The adaptive schedule breaks the battery floor because its bed heater runs land in the eclipse. The power-aware program fixes this by moving the heater runs into sunlight, while the oxygen loop keeps the electrolysis as it was. On a vehicle with a smaller battery, the first measure is therefore an interlock that allows the bed heaters only in sunlight.” Find the flaw, and say how you would check the claim.

CHAPTER 11

Supply-Line Feedback and the Bullwhip Effect in a Perishable Dairy Chain



The controller moves 1 the plant's bottling plan 2 the depot's order to the plant 3 the supermarket's order to the depot

The controller reads 1 the stock record at the plant and the depot 2 the tills' sales record

What disturbs it 1 shoppers' demand, which a promotion doubles

Notice that each stage sees only its own stock and the orders or sales that reach it, while every order waits days for its truck. Illustration, not to scale.

IN PLAIN TERMS **Milk Orders That Grow on the Way Up**

The main idea. A supermarket, a depot and a dairy each order milk from the one above, and every order takes days to arrive. A manager who forgets the milk already on its way orders the same milk twice.

The challenge. A one-week sale at the store then reaches the dairy as a swing many times larger. Milk goes past its date on the shelf, and the dairy pours away fresh milk from the farms when the orders stop.

Why it matters. Counting the milk on its way, and letting every stage see what shoppers buy, keeps the chain steady and saves the milk.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 11.2, the fact in Section 11.3, the two try-its and the fact in Section 11.6, and the key points at the end of Section 11.9
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 11.6), Technology in Practice (Section 11.7), Limitations (Section 11.8) and the two open challenges

The control question Three sensible managers each order milk from the stage above through a delay, and a one-week promotion at the till reaches the plant as a swing many times larger. How far does the classical analysis of the ordering rule predict what a perishable, capacity-limited chain does, and which cure removes the amplification: a different rule, shared information or one planner for the whole chain?

Learning objectives

- Explain why a stage that ignores the milk already on order orders the same milk twice, and name the term of the rule that fixes it.
- Compute the frequency response $|\Gamma(e^{j\omega})|$ of an ordering rule at chosen frequencies and judge whether a chain of such stages is string stable.
- Compute the variance floor of order-up-to under local and shared information and compare it with a measured ratio.
- Trace where the farms' milk goes when the orders collapse, and price the waste, the drain and the dryer.
- Compare the shipped rule, the supply-line-aware rule, shared sales and the centralized MPC on the 26-week challenge.

Prerequisites Sections 1.4.8 (disturbance propagation and string stability), 1.6.1 (engine, seed and Reset), 1.6.3 (MPC option) and 1.6.4 (scenarios, scorecard and export).

The app https://apmonitor.com/apps/milk-distribution-game/

See also Volume 1, Chapter 10 for a stock whose level is set by adapting inflows decided on a calendar with a quantum and a cap, and Volume 1, Chapter 9 for an inventory that is a pure integrator. Volume 7, Chapter 14 for string stability of a line of cars, the same criterion applied to accelerations.

11.1 Introduction

An inventory is an integrating state, and an order is an input that reaches it after a delay. A shortage observed today says nothing about how much milk is already on its way, so a manager who orders again for the same gap orders the same milk twice. When every stage of a distribution chain corrects its own stock from the orders it receives, the corrections travel upstream as apparent demand, and the variance of the orders grows from stage to stage, the bullwhip effect [1]. Treating this as a control problem is older than the name, because production control was written as a servomechanism problem in 1952 [2] and periodic replenishment as a discrete servo problem in 1955 [3]. The inventory and order based production control system of the production-control literature is a feedback loop with a demand-averaging time, a feedback gain and a production delay [4], the bullwhip itself has been measured and avoided with control theory [5], and model predictive control (MPC) has been applied to supply chain optimization and to demand networks [6, 7]. The plant here is a chain of stocks and pipelines, the actuator is an order rate and the controller is a decision rule, so the tools of feedback control apply to it directly.

A common belief holds that sensible managers at every stage cannot together produce a crisis, and the simulator puts that belief to a test. In its second case a one-week promotion doubles what shoppers buy, from 20 to 40 kL/d, and every manager runs the same anchoring-and-adjustment rule. The depot then orders 65.5 kL/d and ten days later nothing, the plant's daily orders vary 12.70 times as much as the shoppers' demand, 32.5 kL of finished milk ages past its code date on the supermarket's shelf and 30.5 kL of raw milk goes down the drain when the dryer is full. Counting every liter already on order pours none. With order-up-to at every stage and every assumption of the analytical theory in force, the plant's variance ratio stays above its floor for every forecast window measured, 11.36 against 8.58 at a week, and sharing the supermarket's sales takes it to 2.73. A block diagram of three loops shows none of these consequences, because the cost of a policy lies in the ages of the cartons, the censoring of demand by an empty shelf and the farm milk that the trough in orders leaves unbottled.

The structure is general. A cascade in which each element answers its neighbor with the same local feedback passes a wave on multiplied by the element's gain at that frequency, and the requirement that no frequency grows along the cascade is string stability, a property defined for infinite strings of

interconnected nonlinear subsystems [8]. Section 1.4.8 applies the same criterion to vessels whose level loops pass flow variation downstream. A line of cars under adaptive cruise control is one such cascade, studied in Volume 7, Chapter 14, and a distribution chain is another. Perishable stocks add a deadline to every unit, and ordering policies for stock with a fixed life were reviewed with blood banks as an application [9]. Case studies of fresh chains for milk, fish and poultry found that their performance improves with more efficient sharing of demand and shelf-life data [10], and dairy planning joins an uncertain raw-milk supply with products as different as cheese, yogurt, powdered milk and ultra-high-temperature milk [11].

However, the record of the problem comes from board games and behavioral experiments [12–14] and from linear analysis [5, 15], so a student seldom has a chain on which the frequency-domain criterion, the analytical variance floor and the waste of a perishable product can be measured together while the rule, the information or the calendar changes one at a time. This chapter measures, on a perishable dairy chain where the order a manager controls can empty a shelf or pour milk away days later, how far the classical analysis of the ordering rule describes what the nonlinear chain actually does. The ordering rule, the order-up-to policy, the variance floor and the centralized MPC are established results. The analysis is a set of measurements on one simulator, including a sweep of the forecast window that places order-up-to at or above its floor every time. The simulator is an age-resolved hourly chain, with three decision calendars and a daily constrained predictive controller, and every equation, constant, delay and limit is stated in the sections that follow. Nine cases, a constant-flow reference and a 26-week challenge report service, waste, cost and order variance together.

Section 11.2 follows the problem from servomechanism theory to the beer game and to fresh food chains. Section 11.3 sets out the chain, its ages, the plant's balance and the cost ledgers, and Section 11.4 the ordering rules, the decision calendar and the centralized planner. Section 11.5 derives the rule's frequency response and the variance floor of order-up-to and measures the floor where its assumptions hold. Section 11.6 reports the nine cases, the challenge and the constant-flow reference, Section 11.7 describes the real stock records, orders and refrigerated trucks, Section 11.8 states the model's limits and Section 11.9 summarizes. The cases and the challenge carry completion rules, and the exercises of Section 11.10 are built on them.

11.2 From Servomechanisms to the Beer Game and Fresh Chains

Inventory control met feedback theory early. The problem of controlling the production rate of a single product was stated in 1952 in terms of servomechanism theory, with Laplace transforms introduced to study the stability and steady-state behavior of the system and a cost criterion constructed for evaluating alternative decision rules [2]. Three years later the simplest periodic replenishment case was analyzed with discrete variable servo theory, and a reorder rule was derived that minimizes the variance of the inventory balance, with that variance written explicitly in terms of the sales forecast error and the system constants [3]. The rules of this simulator are periodic replenishment rules of the same kind, decided once a day.

The game the simulator descends from began with appliances, not beer, as the box below tells. Forrester's article of 1958 described the structure and behavior of a three-stage system, with the oscillation, amplification and phase lag that the game still produces [12, 16]. The Chain view of the simulator is drawn as the game's board, with crates on docks, trucks on a road and order slips moving the other way.

The game became a research instrument in 1989. Subjects managing a simulated distribution system with multiple actors, feedbacks, nonlinearities and delays produced aggregate behavior that diverged systematically from the optimum, and an anchoring-and-adjustment heuristic for stock management, tested econometrically, explained their orders well. The estimates identified misperceptions of feedback, among them a failure to account for the supply line [12, 13]. Production engineers wrote the same loop as an inventory and order based production control system and

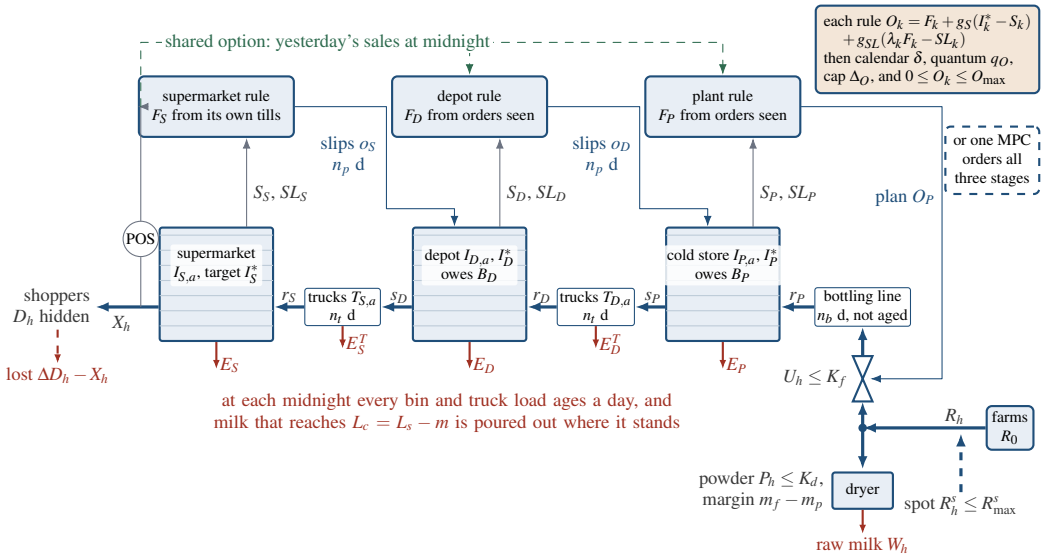

Figure 11.1 The dairy chain as stocks, delays and valves, not to scale. Each stock is a stack of daily age bins $I_{k,a}$ with its target I_k^*, and the depot and the plant carry what they owe downstream, B_D and B_P. Trucks $T_{S,a}$ and $T_{D,a}$ spend n_t days on the road, order slips o_S and o_D spend n_p days in processing, and bottled milk U_h spends n_b days on the line before it enters the cold store. The farms' milk R_h, with optional spot milk R_h^s, is split by the bottling valve, the dryer and the raw-milk drain. Shoppers take the oldest carton first except for a share f_L who take the freshest, and the point-of-sale (POS) record X_h is a censored measurement of the hidden demand D_h. Expired milk E_k and E_k^T leaves every stock and truck. Each rule sees its own stock, its supply line and its forecast input, with the shared option sending yesterday's sales to the depot and the plant at midnight.

showed that its feedback path is essential, that feedforward alone cannot give adequate performance, and that one good compromise makes the time to adjust inventory, the demand averaging time and the production delay comparable [4]. Later experiments found the bullwhip effect persisting after batching, price variation and demand estimation had been removed, with decision makers consistently underweighting the supply line, and found that sharing inventory information helped only somewhat [14].

> ⧖ HISTORY **1956 to 1984: From an Appliance Factory to the Beer Game**
>
> In 1956 and 1957 managers of a household appliance factory in Kentucky were troubled by production that swung from three shifts, seven days a week, to laying off half the people two or three years later. A hand simulation with columns for employees, production rate, backlog and inventory, one line per week, showed that the system created the instability even with a constant demand from the market. That notebook exercise became the production-distribution game, used in a university summer session from 1958 with retail, distributor and factory stages. Beer first appeared as its context in 1973, and the name Beer Distribution Game was first used in 1984. The session that one school plays each year with its incoming business class often has more than 400 participants [12].

The name and the information view arrived in 1997. Orders were shown to be a distorted signal of sales, with the variance of orders larger than that of sales and growing upstream, from four sources, which are demand signal processing, the rationing game, order batching and price variation [1]. Three years later the effect of forecasting and lead time was quantified for a retailer and a manufacturer, extended to chains of several stages with and without centralized customer demand information, and

Table 11.1 Signals of the Chain by Role

Symbol	Role	Unit
I_S, S_D, S_P	CVs, shelf stock and the net stocks $I_k - B_k$	kL
I_S^*, I_D^*, I_P^*	SPs, 3, 2 and 2 days of D_0	kL
O_S, O_D, O_P	MVs, two order rates and the bottling plan, 0 to 80	kL/d
D_h, R_h	DVs, shoppers' demand and the farms' milk	kL/d
—	DVs, promotion, heat, holiday, road and dryer events	events
$I_{k,a}$, $T_{k,a}$	states, age bins on docks and trucks	kL
B_k, SL_k	states, backlogs and supply lines	kL
X_h	measured, the tills' sales	kL
o_k^{in}	measured, orders arriving upstream	kL
F_k	estimates, each stage's forecast	kL/d
U_h, P_h, W_h	outputs, bottled, powder and raw drain	kL/d
C, ρ_k	outputs, cost ledgers and variance ratios	k$, —

shown to be reducible but not eliminable by centralizing that information [15]. Control theory then supplied a way to measure the effect and to avoid it [5], and the frequency response of Section 11.5 is drawn in that spirit.

Perishable stock has a theory of its own. Ordering policies for inventory with a fixed life and for inventory subject to exponential decay were reviewed in 1982, for deterministic and stochastic demand, with a brief review of their application to blood bank management [9]. Case studies of fresh food chains in the Nordic countries, milk among them, found that their performance improved with more efficient sharing of demand and shelf-life data, together with changes in the chain's structure that sped up deliveries [10]. Planning a real dairy chain now treats raw-milk supply and product demand as uncertain together, and in one two-stage stochastic formulation the value of the stochastic solution rose to 21.1 % above the deterministic one [11]. On the control side, chemical engineers posed supply chain optimization as model predictive control [6] and a framework of the same kind managed multi-product, multi-echelon demand networks [7]. The simulator stops short of an industrial plan. It keeps the board game's structure and gives every order a perishable product and a supply that cannot be turned away.

11.3 The Chain and Its Ledgers

The simulated chain is one product, three stocks and the farms that supply them, drawn in Figure 11.1 as the application draws it, milk moving from the farms at the right to the shoppers at the left and orders moving the other way. The supermarket (S) sells to shoppers and orders from the depot (D), the depot orders from the plant (P), and the plant sets a bottling plan and takes the farms' milk, which it cannot refuse. Table 11.1 classifies the signals. Volumes are thousand liters (kL), flows kL/d, money thousand dollars, and a kiloliter at a dollar a liter is $1k. The engine's clock runs in weeks on a one-hour step, so a day is exactly 24 steps.

11.3.1 One Hour in the Chain

Each hour the chain advances in a fixed order. Trucks and the bottling line deliver whatever is due, and orders that have finished processing reach the stage upstream. The shoppers then buy, the depot and the plant ship what they owe oldest first, the orders decided for this hour enter processing, and the plant splits the farms' milk. Every queue entry carries the step at which it arrives, so a road closure delays a truck without losing its milk. With $\Delta = 1/24$ day, the conservation laws on the aggregate

ledgers are Equations (11.1) to (11.4).

$$I_{S,h+1} = I_{S,h} + r_{S,h} - X_h - E_{S,h} \tag{11.1}$$
$$I_{k,h+1} = I_{k,h} + r_{k,h} - s_{k,h} - E_{k,h}, \quad k \in \{D,P\} \tag{11.2}$$
$$B_{k,h+1} = B_{k,h} + o^{\text{in}}_{k,h} - s_{k,h}, \quad k \in \{D,P\} \tag{11.3}$$
$$SL_{k,h+1} = SL_{k,h} + o_{k,h} - r_{k,h} - E^T_{k,h} \tag{11.4}$$

where $I_k = \sum_a I_{k,a}$ is milk on the dock, r_k the receipts, s_k the shipments, E_k and E^T_k the milk that expires on the dock and on the trucks bound for stage k, and o^{in}_D and o^{in}_P the supermarket's and the depot's orders placed n_p days earlier. The volumes are per step, not rates. Orders placed are $o_{k,h} = \Delta \max(0, O_k)$ for $k \in \{S,D\}$ and the volume actually bottled, $o_{P,h} = \Delta U_h$, at the plant. The supply line of the supermarket is its orders in processing, the depot's backlog B_D and the trucks on the road, that of the depot is the same one stage up, and that of the plant is the bottling line. The flows are

$$\begin{aligned} X_h &= \min(\Delta D_h,\ I_{S,h} + r_{S,h}), \\ s_{k,h} &= \min(B_{k,h} + o^{\text{in}}_{k,h},\ I_{k,h} + r_{k,h}), \end{aligned} \tag{11.5}$$

where D_h is the shoppers' demand rate and a shipment rides n_t days to the stage below. Shipments leave oldest first, and the shoppers' purchase takes a share f_L from the youngest bins and the rest from the oldest. A sale the shelf cannot make is lost, $\Delta D_h - X_h$, while an upstream order that cannot be filled stays in B_k until milk arrives.

11.3.2 Ages and the Code Date

An age bin is a state, not a label attached after the run. Receipts enter the bin of the age they carry, and at every midnight each dock and each truck load shifts by one day, as shown in Equation (11.6).

$$\begin{aligned} I^+_{k,a} &= I_{k,a-1}, & I^+_{k,0} &= 0 \\ E_k &= \sum_{a \ge L_c} I^+_{k,a}, & L_c &= L_s - m \end{aligned} \tag{11.6}$$

where L_s is the shelf life, m the sell-by margin and L_c the code date, 11 days by default, and the milk at or past it is poured out where it stands, on a dock or on a road. Milk that spoils on a truck also leaves the receiving stage's supply line in Equation (11.4). The bottling line holds milk that has not yet aged, so a liter enters the cold store at age 0 one day after bottling and is pulled at the code date counted from that arrival. Twenty-four bins hold ages 0 to 23 days.

11.3.3 The Plant's Balance

The farms' milk R_h arrives every hour whatever the plant plans. With bottling plan O_P, line capacity K_f, dryer capacity K_d (zero during a dryer outage) and spot milk up to $R^s_{\max}$ when the option is on, the plant splits its milk as shown in Equations (11.7) to (11.10).

$$R^s_h = \min\left(R^s_{\max},\ [\min(O_P, K_f) - R_h]_+\right) \tag{11.7}$$
$$U_h = \min(O_P,\ K_f,\ R_h + R^s_h) \tag{11.8}$$
$$P_h = \min(K_d,\ R_h + R^s_h - U_h) \tag{11.9}$$
$$W_h = R_h + R^s_h - U_h - P_h \tag{11.10}$$

where $R^s_h = 0$ with spot milk off, U_h is bottled, P_h dried into powder and W_h poured away raw, all in kL/d. Powder is a product, not waste, so only its lower margin is charged. The identity $R_h + R^s_h = U_h + P_h + W_h$ holds at every step to 2.2×10^{-16} by the application's own check.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.3.4 Shoppers, Herds and Events

Demand and the farms' milk follow a calendar. At day d the calendar week is $y_d = w_0 + (d + \frac{1}{2})/7$, a seasonal bump is $G_c(y) = \exp\{-\frac{1}{2}[\mathrm{wrap}_{52}(y-c)/5]^2\}$ with the distance wrapped around a 52-week year, and the rates are Equations (11.11) and (11.12).

$$D_h = D_0\,\omega_d\,[1-\delta_s G_{31}(y_d)]\,e_h\,\max(0,1+\sigma\varepsilon_d) \tag{11.11}$$

$$R_h = R_0\,[1+\phi_f G_{19}(y_d)-\phi_t G_{43}(y_d)]\,e_h^R \tag{11.12}$$

where $\omega_d = 7/(5+2w_e)$ on Monday to Friday and $7w_e/(5+2w_e)$ on Saturday and Sunday keeps the weekly mean at one, e_h is 2 during a seven-day promotion, 1.2 during a 14-day heat wave and 0.85 during a 14-day school holiday, $e_h^R = 0.9$ during the heat wave and 1 otherwise, σ is the demand-noise slider divided by 100 and ε_d a standard normal deviate drawn once a day from the scenario's seeded generator. Two further sliders multiply D_h and R_h by $1+s/100$ with s from -30 to 30, a road closure holds every truck bound for the supermarket for a day, and a dryer outage sets $K_d = 0$ for two days. An event begins at the first step at or after its scheduled time, and the run starts on a Monday.

11.3.5 Ledgers and Scores

The cost ledgers integrate, per step, the increments of Equation (11.13).

$$\begin{aligned}
\Delta C_h &= c_h \Delta \textstyle\sum_k I_k, \qquad & \Delta C_l &= c_l \Delta (B_D + B_P)\\
\Delta C_b &= c_b(\Delta D_h - X_h), & \Delta C_s &= c_s \Delta R_h^s\\
\Delta C_w &= c_w\big(\Delta W_h + \textstyle\sum_k E_k + \sum_k E_k^T\big)\\
\Delta C_{\mathrm{div}} &= (m_f - m_p)\Delta P_h
\end{aligned} \tag{11.13}$$

where the stocks are those on the three docks after the hour's receipts, sales and shipments, and expired milk is charged at the midnight it is poured. The chain's operating cost is $C_h + C_l + C_b + C_w + C_s$, and the total cost adds the margin C_{div} lost to the dryer. Service is cumulative sales over cumulative demand, waste is the finished milk poured out over the milk bottled, and the longest stockout counts consecutive hours in which a sale was lost. The amplification of stage k is a ratio of variances of daily totals, Equation (11.14).

$$\rho_k = \frac{\mathrm{Var}_{d \ge 7}\left(\sum_{h \in d} o_{k,h}\right)}{\mathrm{Var}_{d \ge 7}\left(\sum_{h \in d} \Delta D_h\right)} \tag{11.14}$$

where the first week is skipped, the plant's daily order is its bottling plan ΔO_P summed over the day rather than what the farms let it bottle, and the denominator is the true demand, lost sales included.

> ✦ DID YOU KNOW? **Milk Lost Before It Is Drunk**
>
> In 2010 the United States had an estimated 53.8 billion pounds of fluid milk at the retail level. About 12 percent of it was lost in stores and 20 percent by consumers. Overstocking, because the number of customers is hard to predict, is one listed cause of loss in stores [17].

11.3.6 Start, Integration and Checks

Every case starts lined out. Each stock is set to its target at age 0, every queue is filled at a common steady rate, the chain runs 21 days at that rate with every order equal to it so that the age profiles become periodic, and then the clock and the cumulative ledgers return to zero. At 20 kL/d the supply lines hold 40 kL, 40 kL and 20 kL, the stocks 60 kL, 40 kL and 40 kL, and the oldest cartons are 9, 5 and 2 days old. The engine advances the chain one hour and then integrates the 26 aggregate states (the stocks, backlogs, supply lines, cumulative flows and six cost ledgers) by an Euler step over the same flows, which is exact because every flow is constant over its hour. The controller runs before each step and the historian logs every hour. Over every shipped case, initial milk plus farm intake and

Table 11.2 Teaching Values and Where Each Is Set: M in the Show Model Dialog, C in a Control Card or the Disturbance Panel, A Among the Advanced Controls, S in the Source. No Constant Is Hidden From the Reader

Symbol	Meaning	Default	Range		Symbol	Meaning	Default	Range	
D_0	shoppers' mean demand	20 kL/d	5–60	M	a_S	stock adjustment	0.26/d	0.02–1	C
w_e	weekend lift	1.3	1–1.8	M	β	supply-line weight, a_{SL}/a_S	0.34	0–1	C
δ_s	summer dip at week 31	0.10	0–0.3	M	θ	forecast smoothing	0.36/d	0.02–1	C
w_0	calendar week at the start	9	0–51	M	p	moving-average window	7 d	3–28	A
R_0	farms' mean milk	30 kL/d	10–60	M	δ	decision interval	1 d	hour, day, week	C
ϕ_f, ϕ_t	flush at week 19, trough at 43	0.15, 0.10	0–0.4, 0–0.3	M	q_O	order quantum	0.5 kL/d	0–5	A
$R^s_{\max}$	spot milk available	10 kL/d	0–30	M	Δ_O	largest move per decision	40 kL/d	1–80	A
m_f, m_p	fluid and powder margins	0.40, 0.15 $/L	0.1–1, 0–0.5	M	I^*_S, I^*_D, I^*_P	stock wanted, days of D_0	3, 2, 2	0.5–6	C
c_h	holding	0.01 $/L/d	0–0.05	M	K_f	bottling capacity	40 kL/d	10–60	A
c_l	late order, upstream backlog	0.02 $/L/d	0–0.2	M	K_d	dryer capacity	20 kL/d	0–40	A
c_b	lost sale	0.40 $/L	0–2	M	n_p, n_t	order processing, transit	1, 1 d	1–3	A
c_w	milk poured out	0.60 $/L	0–2	M	n_b	bottling to the cold store	1 d	1–2	A
c_s	spot premium	0.20 $/L	0–1	M	L_s, m	shelf life, sell-by margin	14, 3 d	7–21, 0–5	A
σ	demand noise, day to day	0 %	0–30	C	f_L	shoppers taking the freshest	0 %	0–100	A
q, w	MPC stock and move weights	2, 2	0.1–50, 0–50	C	N	MPC horizon	14 d	7–28	A
$O_{\max}$	station limit on every order	80 kL/d	—	S	—	step, age bins, warm-up	1 h, 24, 7 d	—	S
—	line-out, seasonal width	21 d, 5 wk	—	S	—	default seed, cases 4, 6 and challenge	20261004, 11, 7	—	S

spot milk equals milk held on docks, trucks and the line plus sales, powder, raw discards and expired milk to 7.0×10^{-11} kL, and the age bins sum to the aggregate stocks and the supply lines to their recount from the queues within the same tolerance. Seeding, speed and Reset are the framework's of Section 1.6.1. At half the step the challenge's cost under the suggested tuning changes by a relative 5.9×10^{-6}. The application passes its 42 checks.

11.3.7 Assumptions

One product and three stocks. One supermarket, one depot and one plant stand for an aggregated chain. There is no network of stores, no substitution between products or stores and no returns, and a lost sale is lost for good.

A fixed code date for quality. Milk is pulled at the same age everywhere, trucks included. Temperature, handling and microbial growth, which decide real shelf life, are absent, so this is a logistics

model and not a food-safety model.

> ⚠ **OPEN CHALLENGE** **A Shelf Life That Updates Itself**
>
> Most shelf-life models for fresh dairy products estimate the shelf life once, from the product's condition at the start, while the product meets changing temperatures in distribution, at retail and at home. Industry would gain from systems that update the remaining shelf life as sensor readings, measurements or temperature records arrive. Such a system needs a validated model of deterioration, an estimate of the product's present state and its storage history, and uncertainty is seldom treated. Progress would be multi-batch data, realistic time and temperature histories and validation where the prediction is used [18]. The chapter's fixed code date is the case with none of these.

A smooth, compulsory supply. The farms' milk follows a seasonal curve with no day-to-day variation and must be taken in full, and every liter becomes bottled milk, powder or a raw discard. Cream separation, packaging losses and quality-dependent yield are absent. The milk evaporator of Chapter 4 and the milk powder spray dryer of Volume 5, Chapter 10 control unit operations that this plant reduces to a capacity and a margin.

Daily aging on an hourly clock. Flows move every hour, but ages move at midnight and the bottling day is not aged, so the effective selling window is a whole number of days measured from arrival in the cold store.

Teaching values. Every price, capacity, shelf life, gain and limit, the challenge's among them, is chosen to expose a mechanism. Real values come from a dairy's own costing, its cold-chain and food-safety rules and its contracts with retailers, as in planning studies of real dairy chains [11], and the cost ledger excludes capital, energy and the revenue from fluid sales, so no total here is a profit estimate.

11.4 Ordering Rules on a Decision Calendar

Each stage decides its order on the chain's calendar, through the same quantum and the same cap whether a rule or a person decides. The controller executes every hour, and an execution is a decision only when the calendar says so.

11.4.1 Anchoring and Adjustment With the Supply Line

Every automatic stage runs one rule by name, anchoring and adjustment with supply-line accounting, as shown in Equations (11.15) and (11.16).

$$g(a,\delta)=\frac{1-(1-a)^{\delta}}{\delta},\qquad a_{SL}=\beta a_S \tag{11.15}$$

$$O_k^{\mathrm{req}}=\Big[F_k+g(a_S,\delta)(I_k^{*}-S_k)+g(a_{SL},\delta)(\lambda_k F_k-SL_k)+b_k\Big]_{+} \tag{11.16}$$

where F_k is the stage's forecast, S_k its stock net of what it owes ($S_S=I_S$), $\lambda_k F_k$ the supply line it wants, with $\lambda_S=\lambda_D=n_p+n_t=2\,\mathrm{d}$ and $\lambda_P=n_b=1\,\mathrm{d}$, δ the days between decisions and b_k a transfer bias that is zero in steady use. The fractions are per day, so a weekly or hourly calendar closes the same share of each gap per day, compounded, and $g(a,1)=a$. The application's source names the structure APIOBPCS after the production-control literature, and Volume 4, Chapter 8 copies the same function for its production loop. In control terms the stock term is a proportional controller on an integrating plant with a dead time, and the supply-line term is a second proportional term on the controller's own pipeline. With β small, the loop orders again for a gap already on its way. The shipped human rule has $a_S=0.26/\mathrm{d}$, $\beta=0.34$ and $\theta=0.36/\mathrm{d}$, the shape of the rule fitted to beer-game players [13] with teaching values applied per day, not the published weekly estimates. The supply-line-aware rule

sets $\beta = 1$, and order-up-to sets $a_S = a_{SL} = 1$ on a p-day moving average. An advanced option runs the chosen rule at one stage only and the human rule at the other two.

The forecasts are updated as their data arrive, and a decision samples them. The supermarket observes its own sales every hour, which undercount demand while the shelf is empty, and the depot and the plant observe the orders that reach them. With shared information the two upstream stages forecast instead from the supermarket's sales of the day just ended, published at midnight. Exponential smoothing applies $F_k \leftarrow F_k + \gamma(x - F_k)$, with x the observation as a daily rate and $\gamma = 1 - (1-\theta)^{1/24}$ for an hourly observation or $\gamma = \theta$ for a published day, and order-up-to takes the mean of the last $24p$ hourly observations or the last p published days.

11.4.2 Shaping a Decision

The calendar and the shaping come from the framework's `shared/src/framework/core/cadence.js`, which also shapes the quota decisions of Volume 1, Chapter 10. A calendar decides at every execution (hourly), at every midnight (daily, as shipped) or at every seventh midnight from the start, a Monday (weekly), so every instant falls on the hourly step and the realized intervals are exactly 1, 24 and 168 steps. At a decision the request becomes the order in force as shown in Equation (11.17).

$$\begin{aligned} \tilde{O} &= q_O \,\mathrm{round}\left(O_k^{\mathrm{req}}/q_O\right) \\ O_k &= \Pi_{[0,O_{\max}]}\left(O_k^- + \mathrm{sat}_{\Delta_O}(\tilde{O} - O_k^-)\right) \end{aligned} \tag{11.17}$$

where O_k^- is the order in force before the decision, ties round away from zero, sat caps the move at the largest whole number of quanta not above Δ_O, and Π clips to the station's limits pulled onto the quantum grid. Between decisions the order holds. With $q_O = 0.5\,\mathrm{kL/d}$, $\Delta_O = 40\,\mathrm{kL/d}$ and $O_{\max} = 80\,\mathrm{kL/d}$ as shipped, the quantum and the cap are the only anti-saturation the rules need, because Equation (11.16) has no integral or derivative term to wind up or to filter. The plant's plan is further limited by Equation (11.8), and a request above the station limit or the line's capacity is flagged as saturated.

11.4.3 Manual Stations and Transfer

Each stage has its own MANUAL and AUTO switch, and a master switch throws all three. A station in MANUAL holds the learner's value, decided on the same calendar and shaped the same way. Switching to MANUAL puts the order in force on the station, so nothing moves. On the way back to AUTO the rule sets $b_k = O_k^- - O_k^{\mathrm{raw}}$, with O_k^{raw} the value of Equation (11.16) before the floor at zero, so its first request is the order in force, and later decisions multiply b_k by $(1 - \max(a_S, 0.05))^{\delta}$, the manual reset of a proportional controller decaying at the rule's own rate. The PID loops of Section 1.6.2 reach the same end through their running sum, which this rule does not have. While a stage is automatic its manual station follows the rule's request before the quantum, the cap and the limits, so that a station never ratchets on a shaped value. In the application's check a transfer from AUTO to MANUAL and back moves no order of any of the four rules.

11.4.4 Seasonal Feedforward

An advanced option, off as shipped, gives the anchor a calendar. A stage that forecasts from the shoppers' sales smooths their level with the known weekly and summer pattern divided out and multiplies it by the pattern on the day its milk will be sold, 2, 5 and 7 days ahead at the supermarket, the depot and the plant. A stage that forecasts from the orders it receives knows no calendar for them and is unchanged. On the constant-flow reference the option moves no order at all.

11.4.5 One Planner for the Chain

The centralized controller replaces the three rules by one quadratic program (QP) at each decision, built with the framework's condensed linear MPC, `shared/src/framework/core/linmpc.js`, of

Section 1.6.3. Its internal model steps one day at a time, as shown in Equation (11.18) for the default one-day delays.

$$\begin{aligned} I_S^{j+1} &= I_S^j + z_{S2}^j - \hat{D}_j, \quad I_D^{j+1} = I_D^j + z_{D2}^j - z_{S1}^j \\ I_P^{j+1} &= I_P^j + z_B^j - z_{D1}^j, \quad z_{S1}^{j+1} = u_S^j,\ z_{S2}^{j+1} = z_{S1}^j \\ z_{D1}^{j+1} &= u_D^j,\ z_{D2}^{j+1} = z_{D1}^j, \quad z_B^{j+1} = u_P^j \end{aligned} \tag{11.18}$$

where u are the three daily orders, z_{S1} and z_{D1} hold an order in processing, z_{S2} and z_{D2} the same order on its truck, z_B the bottling line, and the stocks are net, so a negative value is a backlog. At each solve the eight states are read from the chain itself, every queued order, backlog, truck and bottle put in the slot for the day it will move. The demand forecast $\hat{D}_j$ is a level smoothed at θ from the published sales divided by the known weekly and seasonal pattern, times that pattern on day j, so the planner reads the shared sales whatever the information setting. The farms' milk $\hat{R}_j$ is today's intake carried along its known season. No promotion or heat wave is previewed. The horizon of $N = 14$ days is split into eight move blocks, four of one day, two of two days and two of three, or into weeks on the weekly calendar, and a stage in MANUAL is fixed at its station. The QP is Equation (11.19), the objective of Equation (1.16) with a terminal term and linear prices on the slacks.

$$\begin{aligned} \min \ & \sum_{j=1}^{N} \sum_k q'(I_k^j - I_k^*)^2 + 5q' \sum_k (I_k^N - I_k^*)^2 \\ & + w' \sum_j \sum_k (u_k^j - u_k^{j-1})^2 + \sum_i (\varepsilon s_i^2 + \sigma_i s_i) \\ \text{s.t. } & 0 \le u_S, u_D \le O_{\max}, \qquad 0 \le u_P \le K_f \\ & |u_k^j - u_k^{j-1}| \le \Delta_O, \qquad u_P \le \min(K_f, \bar{R} + R_{\max}^s) \\ & I_k^j \le \bar{F}\, c_k + s_k^f, \qquad I_k^j \ge -s_k^0 \\ & z_B^j \ge \hat{R}_j - K_d - s^d, \qquad s \ge 0 \end{aligned} \tag{11.19}$$

where $q' = 10^{-3} q = 0.002$ and $w' = 10^{-3} w = 0.002$ convert the sliders to thousand dollars, the terminal day carries six times the stage weight in all, u^j are the inputs of move block j with u^0 the orders in force, $\bar{R}$ the forecast intake averaged over a block, with $R_{\max}^s$ only when spot milk is on, and $\bar{F}$ the mean forecast over the horizon. The freshness ceilings c_k are the code date less the days of target stock and transit that the milk spends at the other stages, at least one day, 5, 4 and 4 days of forecast at the default targets and transits, a linear stand-in for the expected discards of the age bins. One slack per soft band carries a linear price σ_i, c_w for the three ceilings and the dryer floor, c_b for the supermarket's floor at zero and $c_b/2$ for the two upstream floors, and a small quadratic $\varepsilon = 10^{-3}$ keeps the program strictly convex. A band-price setting beside the planner's weights multiplies every σ_i and is one as shipped. The program has 31 variables (24 block inputs and 7 slacks) and 209 constraint rows, and the dual active-set method of Goldfarb and Idnani [19] solves it, as in Section 1.6.3. Only the first day's orders are applied, shaped by Equation (11.17). A failed solve holds the orders in force and counts the failure, and the first automatic decision after a manual transfer holds the order in force. The planner tracks stocks with economic penalties and is not a profit optimum.

11.4.6 Plant, Analysis and Planner Models

The simulation model is Equations (11.1) to (11.14), with every carton aged, every sale censored and every limit hard. The ordering rules hold no model at all beyond the lead times in λ_k. The analysis model of Section 11.5 is the rule linearized, without the floor at zero, the quantum, the cap or any expiration, and it is used only to draw $|\Gamma(e^{j\omega})|$ and to compute the variance floor. The MPC's internal model is Equation (11.18), which knows the delays and the limits but neither the ages of the cartons

nor the censoring of sales, and whose demand forecast knows the calendar but no event. No model is identified from data.

11.5 The Rule's Frequency Response and the Floor

A chain of stages that each answer the stage below is a platoon of inventories, and the application draws its criterion as Volume 7, Chapter 14 draws the cars' and as Section 1.4.8 draws a level loop's. Linearized with daily decisions and no floor at zero, Equation (11.16) with net stock $S_t = S_{t-1} + O_{t-1-\lambda} - d_{t-1}$ and supply line $SL_t = O_{t-1} + \cdots + O_{t-\lambda}$ gives the ratio of a stage's orders to the orders it receives, Equation (11.20).

$$\Gamma(z) = \frac{(1 + a_{SL}\lambda)Q(z)H_F(z) + a_S z^{-1}}{Q(z) + a_S z^{-(\lambda+1)} + a_{SL}Q(z)\sum_{i=1}^{\lambda} z^{-i}} \tag{11.20}$$

where $Q = 1 - z^{-1}$, $z = e^{j\omega}$ with ω in radians per day, and H_F is the forecast filter, $\theta z^{-1}/[1-(1-\theta)z^{-1}]$ for exponential smoothing or $p^{-1}\sum_{i=1}^{p} z^{-i}$ for a moving average. $\Gamma(1) = 1$ for every rule, because a stage must in the end order what it sells, so a lasting change passes on whole and a chain is string stable only if $|\Gamma(e^{j\omega})| \le 1$ at every frequency [8]. For order-up-to the denominator is exactly one and $\Gamma = z^{-1} + (1+\lambda)QH_F$.

For white demand the variance ratio of stage k is the mean of $|T_k|^2$ over frequency, Equation (11.21), by Parseval's theorem.

$$\frac{\mathrm{Var}(O_k)}{\mathrm{Var}(D)} = \frac{1}{\pi}\int_0^{\pi} |T_k(e^{j\omega})|^2 \,\mathrm{d}\omega, \qquad T_k = z^{-1}\Gamma_k T_{k-1} \tag{11.21}$$

where T_k is stage k's orders over the shoppers' demand, $T_1 = \Gamma_1$ at the supermarket, and the factor z^{-1} is the day by which each upstream stage sees the orders below it late. With shared sales the forecast of each upstream stage is driven by demand instead of by the orders it receives. For order-up-to on a p-day moving average, with $L = \lambda + 1$ the lead time plus the review day, three days at the supermarket and the depot and two at the plant, the classical expressions are Equation (11.22).

$$\begin{aligned} \rho_k^{\mathrm{loc}} &\ge \prod_{i \le k}\left(1 + \frac{2L_i}{p} + \frac{2L_i^2}{p^2}\right) \\ \rho_k^{\mathrm{sh}} &= 1 + \frac{2L_1}{p} + \frac{2\sum_{i \le k} L_i^2}{p^2} \end{aligned} \tag{11.22}$$

where the first line is the floor under local information of Chen and coauthors [15], exact for one stage, and the second is what this chain gives with shared sales. The application's checks reproduce the one-stage factor from Equation (11.21) within 1.8×10^{-15} at five windows from 1 to 28 days, and the shared expression from the linear chain within a relative 3.1×10^{-15} at the seven windows of the sweep below. The shared expression is not the centralized value $1 + 2\sum L/p + 2(\sum L)^2/p^2$ of the same reference, 5.90 at the plant for $p = 7$, because here the orders reach each stage a day late while the sales reach every stage at the same midnight. Each stage then adds its own L^2 once instead of multiplying the stage below, so sharing does not remove the squared lead-time term.

The three rules have three different peaks, as shown in Figure 11.2(a). The shipped human rule peaks at 1.89 at 0.395 rad/d, a wave of 15.9 days, so any swing near that period leaves each stage larger than it arrived. Counting the supply line lowers the peak to 1.56 at 0.324 rad/d, a 19.4-day wave, and the challenge's suggested $a_S = 0.3$ with $\beta = 1$ peaks at 1.60. Order-up-to on a 7-day average peaks at $13/7 = 1.857$ at the two-day wave, $\omega = \pi$, with ripples that touch one at multiples of $2\pi/7$. None of the three is string stable, which is one reason the cures that work in practice change the information rather than the rule alone. The curve is not drawn for the MPC or for hourly and weekly calendars, where the linearization would misdescribe the controller.

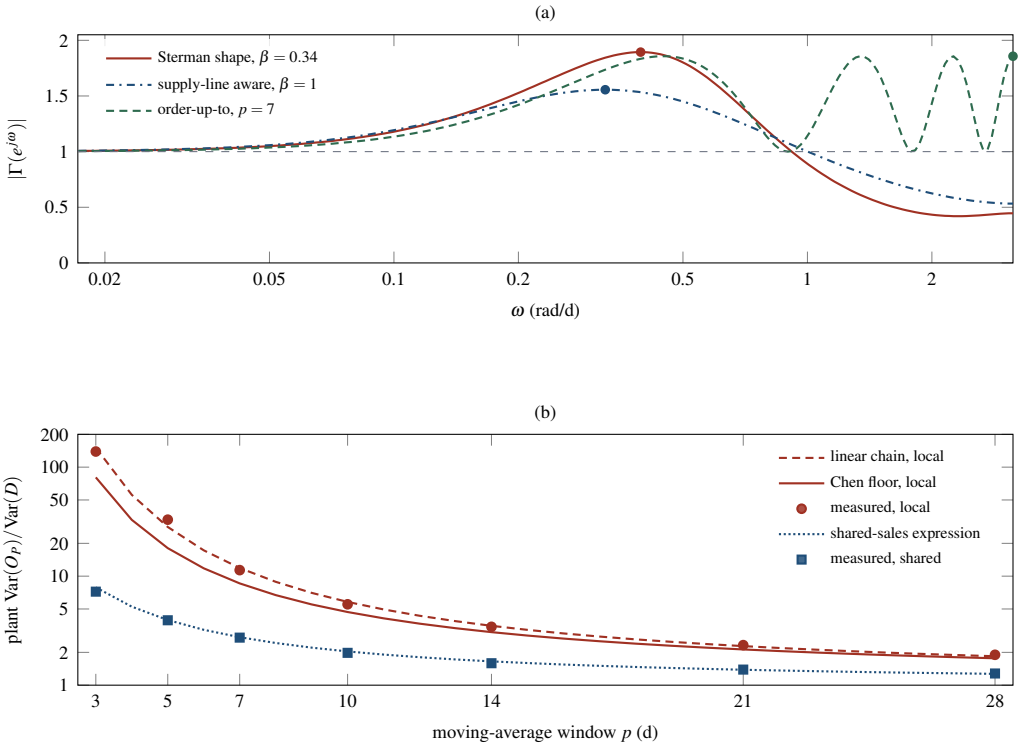

Figure 11.2 The rule and the floor, from the application's own functions. (a) $|\Gamma(e^{j\omega})|$ of the depot's rule, $\lambda = 2$ d, for the shipped human rule (solid), the same rule with $\beta = 1$ (dash-dotted) and order-up-to on a 7-day average (dashed), peaks marked by dots and unity by a thin dashed line. (b) The plant's variance ratio against the forecast window under the assumptions of case 4, with the Chen floor of Equation (11.22) (solid), the linear chain's own local ratio from Equation (11.21) (dashed), the shared-sales expression (dotted), and the ratios measured over 104 weeks at seed 11 (dots local, squares shared).

The floor holds where its assumptions hold. Case 4 runs Chen's chain, with order-up-to at every stage, no quantum and a cap of 80 kL/d, a 60 kL/d bottling line and a 40 kL/d dryer, a 21-day shelf life pulled at its date, spot milk on so that no plan is short of milk, and 2 % white daily noise on a flat calendar. Swept over $p = 3$ to 28 days for 104 weeks at seed 11, the plant sits at or above its floor every time, from 139.1 against 80.6 at $p = 3$ through 11.36 against 8.58 at $p = 7$ to 1.90 against 1.77 at $p = 28$, as shown in Figure 11.2(b). The floor is a bound, not the expected value. The linear chain's own ratio at $p = 7$ is 11.89, because the orders that reach the depot carry the supermarket's forecast and are correlated from day to day. Over seeds 1 to 8, which only a headless run can choose because the case fixes its seed, 104 weeks give a mean of 11.80 and a minimum of 11.22. At 104 weeks all eight seeds sit above the floor at every window, while at 52 weeks single seeds dip by up to 1.5 % below it at $p = 21$ and 28, 2.12 against 2.13 and 1.74 against 1.77, where the expected value exceeds the floor by only 4 to 7 %. With shared sales the same sweep falls below the local floor at every window and within 9 % of the shared expression, 7.21 against 7.89 at $p = 3$, 2.73 against 2.76 at $p = 7$ and 1.28 against 1.27 at $p = 28$. In the perishable, capacity-limited chain of the other cases orders are clipped, sales censored and milk poured out, and the floor there is a reference, not a bound.

One plausible expectation fails by measurement. With $a_S = 1$ and $\beta = 1$ the plant's order might be expected not to overshoot at a step, but a lasting 10 % demand step on the constant-flow reference lifts it to 34.9 kL/d before it settles at 22, because the supply line wanted, $\lambda_k F_k$, rises with the forecast

Table 11.3 The Cases and the Challenge

#	Change from the defaults	Teaching point
1	supermarket in MANUAL, flat calendar, promotion at week 2	the lead time by hand
2	flat calendar, promotion at week 2	the bullwhip and its milk
3	case 2 with $\beta = 1$	count the supply line
4	order-up-to under Chen's assumptions, 2 % noise, seed 11	a floor no tuning beats
5	$\beta = 1$, shared sales, seasons from week 9	the push and the dryer
6	case 4 with shared sales	information moves the floor
7	case 2 with $L_s = 10$ d, $f_L = 50\,\%$	perishability multiplies waste
8	MPC, shared sales, $K_f = 30$, $K_d = 15$ kL/d, seasons, promotion at week 2	limits planned for
9	$\beta = 1$, shared, $a_S = 0.3$, $R_0 = 22$, weekend lift, heat at week 2	who carries the shortage
ch	seasons, 5 % noise, seed 7, promotion at 10, heat at 20	several limits at once
ref	flat calendar, nothing moves	the equilibrium check

Table 11.4 Measured Cases. Total Cost Is the Operating Cost Plus the Margin Lost to the Dryer

Case and setting	wk	ρ_P	poured (kL)	raw (kL)	lost (kL)	outage (h)	total (k$)
1, order held at 20	8	0.00	0.0	0.0	80.0	96	226.1
2, as shipped	8	12.70	32.5	30.5	31.0	41	254.7
3, $\beta = 1$	8	7.61	0.0	0.0	27.0	44	205.8
4, $p = 7$	26	11.19	0.0	0.0	0.0	0	708.2
5, as shipped	16	0.21	0.0	0.0	0.0	0	528.8
5, $K_d = 12$	16	0.21	0.0	167.0	0.0	0	587.2
6, shared, $p = 7$	26	2.52	0.0	0.0	0.0	0	708.1
7, as shipped	8	8.95	261.4	0.0	24.8	43	298.7
8, MPC	8	0.35	0.0	0.0	0.0	0	206.7
8, $\beta = 1$ rule	8	2.11	13.1	8.4	1.5	4	227.7
9, spot off	8	6.15	0.0	0.0	0.0	0	78.8
9, spot on	8	0.74	2.4	0.0	0.0	0	118.2
ref	8	—	0.0	0.0	0.0	0	218.4

ρ_P the plant's ratio, Equation (11.14), poured the finished milk past its code date.

and must be filled once. What $\beta = 1$ removes is the oscillation. The plant's net stock crosses its target twice and overshoots it by 2.73 kL, and its order never falls below 21.3 kL/d after the peak, while at $\beta = 0.34$ with the same a_S the stock crosses its target 19 times and the order swings between 0 and the 80 kL/d limit.

11.6 Nine Cases on One Chain

Table 11.3 lists the cases. Each starts from the defaults of Table 11.2, the shipped human rule at every stage, local information and daily decisions, and changes only what its row lists. The figures below come from the engine run headless on each case's seed, over the window named, and Table 11.4 collects them. A variant that a case's briefing reaches by changing a setting and pressing Reset is measured that way, and the same change made before the first Run gives the same run. Integrated errors appear on the scorecard but are not quoted, because they keep accruing with the window.

11.6.1 The Lead Time by Hand

In case 1 the learner sets the supermarket's order while the depot and the plant run the human rule, and at the midnight of day 14 a promotion doubles demand to 40 kL/d for seven days. An order raised at that midnight is processed on day 15 and its first milk reaches the shelf in the first hour of day 16, while one raised at noon on day 14 waits for the next midnight and arrives on day 17. Left at 20, the 60 kL shelf empties early on day 17 and 80 kL of sales are lost, with the shelf empty for 96 hours in a row. Raised to 40 at the promotion's midnight and set back to 20 on day 19, it loses 6.5 kL, because

the depot behind it cannot ship the whole jump at once, and two days at 60 followed by three at 40 lose nothing. Because every order of the supermarket stays at 20 in the first run, the stages above see no variation at all and their ratios are zero. A turn-based option pauses the run at every midnight, and the application's check confirms 84 pauses in 84 days over four runs on three calendars, so a class can play it as the board game.

> **TRY IT IN THE APP** **1 · Play the Supermarket**
>
> Set Turn-based in Advanced so the run pauses at every midnight. At the midnight that opens the promotion on day 14, raise Supermarket order to 40, and set it back to 20 on day 19. The first extra milk reaches the shelf on day 16, and the shelf loses 6.5 kL of sales against 80 kL with the order left at 20.

11.6.2 The Bullwhip and Its Milk

Case 2 gives every stage the human rule. The supermarket's order peaks at 49.5 kL/d on day 17, the depot's at 65.5 kL/d on day 21 before falling to zero on day 31, and the plant's plan reaches the station's 80 kL/d limit on day 23 while the farms' 30 kL/d is all it can bottle. By day 40 it bottles 1.0 kL/d. Over eight weeks the ratios are 1.90, 6.36 and 12.70 at the supermarket, the depot and the plant. The overshoot ages on the supermarket's shelf until 32.5 kL, 2.54 % of what was bottled, is poured out at the four midnights that open days 44 to 47, and in the trough the dryer runs full for 144 hours and 30.5 kL of raw milk goes down the drain. The shelf is empty for 65 hours in all and 41 in the longest stretch, and the depot and the plant owe the stage below them for 270 and 430 hours. Figure 11.3 shows the trough at 5.75 weeks. A sweep in the application's checks makes the price explicit, because over 26 weeks of this case the plant's ratio and the milk poured out rise together as a_S goes from 0.1 to 1, from 2.5 and 18.5 kL to 72.1 and 141.5 kL, and as β falls from 1 to 0, from 7.9 and nothing to 29.8 and 183.0 kL.

> **TRY IT IN THE APP** **2 · The Bullwhip**
>
> Load the scenario, press Run and watch the depot's order on the trends through week 8. The one-week promotion at the till becomes a depot order of 65.5 kL/d on day 21 and nothing at all on day 31. Then open the Bullwhip view, where the plant's bar reads 12.70 against the shoppers' 1.

Case 3 changes only β, to 1. No milk is poured out, finished or raw. The depot's orders stay between 7.0 and 57.0 kL/d, the plant never bottles less than 14.5 kL/d, and the plant's ratio by week 8 is 7.61. The shelf still empties, losing 27.0 kL of sales against 31.0, and the plan still reaches 80 kL/d, because the backlog behind the promotion is real and the farms let the plant bottle only 30 kL/d. Counting the pipeline removes the double ordering, not the lead time or the farms' ceiling. The calendar matters as much as the gains. Deciding every hour instead of every midnight cuts case 2's lost sales from 31.0 to 8.4 kL and its longest stockout from 41 to 28 hours, but raises the plant's ratio from 12.70 to 13.89, while deciding weekly on Mondays orders the promotion week at the old rate and loses 80.0 kL, the board game's slowness.

Pause. Try Exercise 11.1 before reading on.

11.6.3 The Floor and Shared Sales

Over 26 weeks case 4 measures 2.13, 5.31 and 11.19 against floors of 2.22, 4.95 and 8.58. The supermarket lands on its value within the scatter of a 26-week estimate, because for one stage the expression is exact, and the depot and the plant land above theirs. At $p = 14$ the plant measures 3.34 against 3.07 and at $p = 28$ 1.90 against 1.77. Nothing is poured away and no sale is lost, because the shelf life is long and the capacities ample. This is the honestly infeasible case. With local information no forecast window gets under the floor, and the floor approaches one only as the window grows without end. Case 6 changes only the information, every stage forecasting from the published sales.

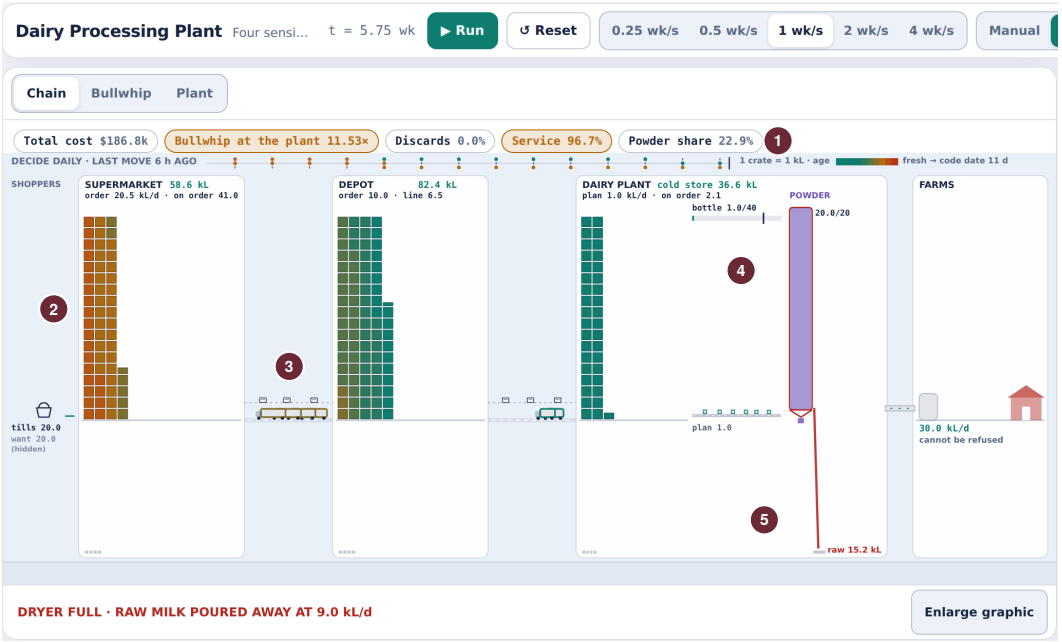

Figure 11.3 Case 2, the bullwhip, at $t = 5.75$ wk (day 40, 6 a.m.) in the Chain view, run in AUTO with nothing changed and an 8-week trend window, cropped to the header and the Chain view. The badges (1) read total cost $186.8k, bullwhip at the plant 11.53×, discards 0.0 %, service 96.7 % and powder share 22.9 %, and the status line reads DRYER FULL, raw milk poured away at 9.0 kL/d. The supermarket (2) holds 58.6 kL with 41.0 kL on order, every carton 7 to 9 days old against an 11-day code date, order slips and trucks (3) ride the lanes between the stages, the depot holds 82.4 kL and orders 10.0 kL/d, and the plant's cold store holds 36.6 kL while its line (4) bottles 1.0 kL/d of its 40, so its dryer takes 20 of its 20 and the drain (5) has taken 15.2 kL of raw milk. The decision strip reads daily, last move 6 h ago. In the app's trends, cut from this picture, the promotion at week 2 appears in the shoppers' demand, the depot's order climbs to its peak and falls to zero, and the plant's net stock falls far below its target while its bottling plan rests at the station limit. The app's scorecard, also cut, gives an operating cost of $117.6k, a powder diversion loss of $69.2k, a longest supermarket stockout of 41 h and 15.2 kL of raw milk discarded.

The ratios become 2.13, 2.51 and 2.52, the supermarket unchanged because it always forecast from its own tills, and the plant far under its local floor and near the shared value of 2.76. The total cost barely moves ($708.2k against $708.1k), because with a 21-day shelf life the extra variance costs only holding.

🕮 **Pause.** Try Exercise 11.2 before reading on.

11.6.4 Farms That Push

Case 5 starts in calendar week 9 with the weekend lift, the summer dip and the herds' season on, every stage counting the supply line and reading the shared sales. Over 16 weeks the farms send 3700 kL, peaking at 34.5 kL/d on day 69, while bottling stays between 18.0 and 22.0 kL/d. All of the difference goes to powder, 1486 kL or 40.2 % of the intake, a margin of $371.6k lost to the dryer. Without the flush the same run dries 1146 kL for $286.5k, so the flush itself costs $85.1k. With the dryer cut to 12 kL/d it is full for 2112 of the 2688 hours, 167.0 kL of raw milk is poured away, and the operating cost rises from $157.2k to $257.4k. The plant then behaves as a ratio controller on a flow it cannot change, splitting it between two outlets with limits, and the dryer is its surge capacity,

sized by the flush rather than by the average. The seasonal feedforward makes the supermarket chase the weekly pattern, raising its ratio from 0.94 to 2.65 and lowering the total by \$0.9k, because the shelf absorbs a weekly pattern almost free.

11.6.5 A Shorter Shelf Life

Case 7 is case 2 with a shelf life of 10 days instead of 14, pulled three days before its date, and half the shoppers taking the freshest carton. Over eight weeks 261.4 kL is poured out, 17.0 % of what was bottled, against case 2's 32.5 kL. The two causes compound. At 10 days with every shopper taking the oldest first the chain pours out 151.5 kL (10.5 %), and at 14 days with half of them taking the freshest it pours out 71.6 kL (5.3 %). Service barely moves, 98.0 % against 97.5 %, because the waste comes from the overshoot aging on the shelf. At a 7-day shelf life the milk passes its 4-day code date before it reaches the shelf and 88.4 % of what is bottled is poured out, which teaches nothing about the rule, so the case uses 10 days.

11.6.6 Planning Around the Limits

Case 8 hands every order to the MPC, with the line cut to 30 kL/d and the dryer to 15 kL/d so that both limits bind through the promotion, and the seasons on. It solves all 56 midnights of eight weeks without a failure, in a mean of 9.6 active-set iterations and at most 27. It ramps the supermarket's order to 47.5 kL/d on day 17, holds the plant at its 30 kL/d line for 312 hours, keeps the dryer at no more than 14.6 of its 15 kL/d, and serves every shopper with nothing poured away, at an operating cost of \$68.7k. The supply-line-aware rule on the same chain and the same shared sales asks for up to 55.5 kL/d against the line, overflows the dryer for 192 hours, empties the shelf for 4 hours (1.5 kL of lost sales) and pours away 13.1 kL of finished and 8.4 kL of raw milk, at \$92.9k. The planner meets the limits in its prediction, while the rule meets them as backlog and overflow.

11.6.7 Heat and Spot Milk

Case 9 is a plant that sells nearly all its milk fluid, the farms sending 22 kL/d to shoppers who buy 20 on average with a weekend lift, and a two-week heat wave at week 2 that raises demand 20 % and cuts the herds' milk 10 %. The stocks carry the shelf through, and no sale is lost either way. Without spot milk the shortage lands upstream, the plant owing the depot for 423 hours, up to 32.8 kL. With spot milk on the plant buys 69.8 kL and owes nothing, but the run costs \$118.2k against \$78.8k, the premium plus farm milk displaced into powder afterward (99 kL dried against 35). The override decides who carries the shortage, and here the safety stock carried it for free.

11.6.8 Twenty-Six Weeks Under Three Limits

The challenge runs 26 weeks from calendar week 9 through the spring flush, with a promotion at week 10, a heat wave at week 20 and 5 % daily noise at seed 7. Its target asks, at week 26, for a total cost of at most \$810k, finished milk poured out at most 3 % of bottling and no supermarket stockout longer than 24 hours. Table 11.5 compares seven configurations. The shipped rule fails all three clauses, at \$898.2k, 3.72 % and 44 hours, missing the limits by \$88.2k, 0.72 points and 20 hours. Each cure alone fails. Counting the supply line costs \$817.8k and still leaves a 44-hour stockout, and sharing the sales alone costs \$820.8k. Both together with $a_S = 0.3$, the suggested tuning, cost \$791.3k, pour out 1.28 % and leave a longest stockout of 6 hours, inside the limits by \$18.7k, 1.72 points and 18 hours, and at a half-hour step the result is the same to \$0.1k. The MPC with shared sales does better still, \$774.4k, 0.47 % and no stockout, over 182 solves without a failure. The diversion is \$468k to \$517k of every total, because the farms send half again what the shoppers buy. Weekly decisions break the suggested tuning, at \$916.6k, 4.82 % and 56 hours, while hourly decisions pass at \$789.2k with no stockout. The limit is \$810k rather than the \$820k first calibrated, because at \$820k sharing the sales alone missed by only \$0.8k. The verdicts do not hinge on seed 7. Over seeds 1 to 10, run headless because the case fixes its seed, the same four configurations fail and the same three pass on every seed, and the MPC and the hourly calendar cost less than the suggested tuning on every one,

Table 11.5 The 26-Week Challenge in Seven Configurations

Setting	chain (k$)	div. (k$)	total (k$)	waste (%)	outage (h)	ρ_S, ρ_D, ρ_P	target
shipped	430.1	468.1	898.2	3.72	44	1.43, 4.07, 9.51	not met
$\beta = 1$	307.9	509.9	817.8	1.49	44	1.01, 1.81, 4.92	not met
shared	325.1	495.7	820.8	2.29	20	1.72, 1.82, 1.78	not met
suggested	280.1	511.2	791.3	1.28	6	1.53, 1.25, 1.50	met
MPC, shared	257.5	517.0	774.4	0.47	0	1.00, 0.82, 0.67	met
hourly	277.0	512.1	789.2	1.14	0	1.30, 1.42, 2.29	met
weekly	423.2	493.4	916.6	4.82	56	0.70, 0.30, 0.21	not met

Suggested: $\beta = 1$, shared sales, $a_S = 0.3$, also run hourly and weekly.
Service 98.97 % shipped and 99.94 % suggested.

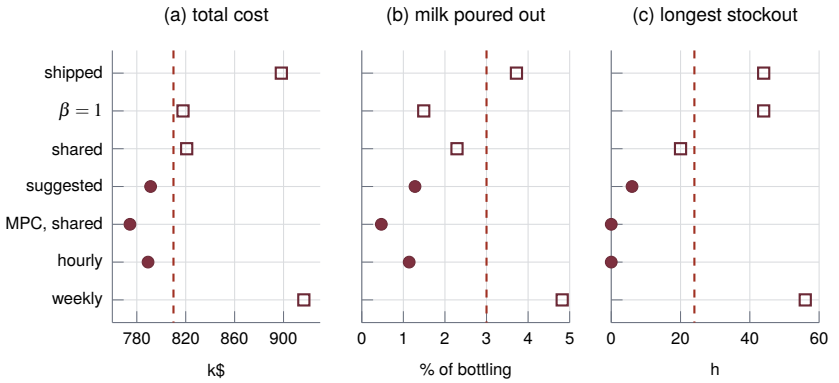

Figure 11.4 The 26-week challenge of Table 11.5 in its seven configurations against the three limits of the target (dashed): (a) a total cost of $810k, (b) 3 % of bottling poured out and (c) a 24-hour stockout. Filled circles mark the settings that meet the target and open squares those that do not. Notice that counting the supply line alone passes only (b), sharing the sales alone passes (b) and (c) but fails (a), and only their combination, the planner and the hourly calendar pass all three.

although the MPC's longest stockout reaches 12 hours on seed 4. The challenge is judged on the case as given, so a changed model parameter or disturbance setting, or an extra disturbance button, voids the attempt. Figure 11.4 draws the seven configurations against the challenge's three limits.

> ✦ DID YOU KNOW? **The Smallest Bullwhip, the Worst Result**
>
> From the simulator (Table 11.5). Deciding weekly gives the challenge its smallest variance ratios, 0.70, 0.30 and 0.21 from the supermarket to the plant. It is also the costliest of the seven settings, at $916.6k with 4.82 % poured out and a 56-hour stockout. A small ratio alone does not make a good chain.

In the constant-flow reference every order stays at exactly 20 kL/d for eight weeks and every stock returns to the same value at every midnight, while the plant dries a third of its milk for $140.0k of margin.

11.7 Technology in Practice

What Measures It

Each rule of this chapter reads three things: its stage's stock, the milk it has on order, and a forecast input, either the orders that reach it or the supermarket's published sales. In a real store these come from one stock record that software keeps from the deliveries received and the sales scanned at the tills, and the same software forecasts demand and triggers replenishment [20]. The model's

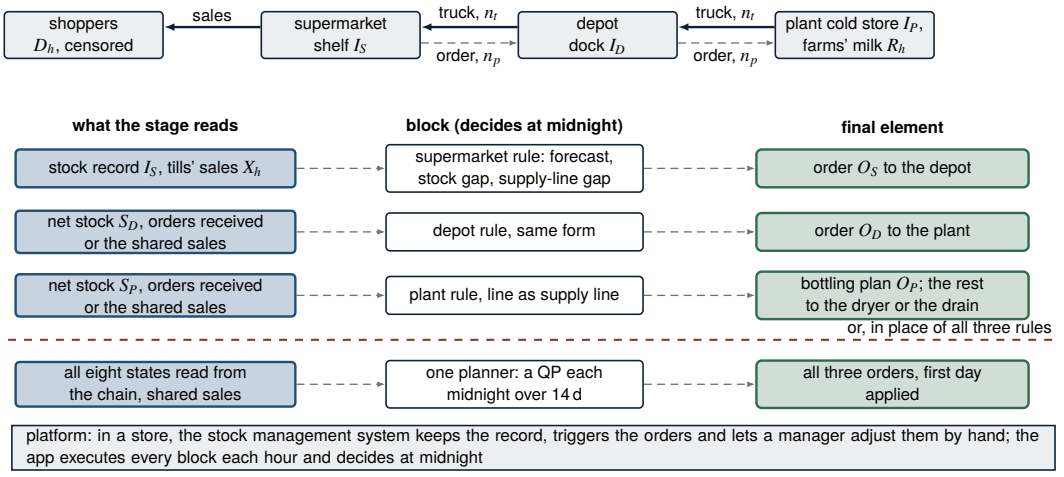

Figure 11.5 The ordering loops of Section 11.4 on real hardware, from what each stage reads to the order it sets, under the chain whose milk and orders move in opposite directions. Notice that the three decentralized rules each see one stock and one input, while the planner below the dashed line reads the whole chain and sets every order at once.

supermarket reads its shelf exactly, and its tills undercount demand only while the shelf is empty, as Section 11.4 states. A real record also drifts away from the shelf, and a store corrects it only by counting [20].

> **TECHNOLOGY** **Sensor: The Stock Record and the Tills**
>
> A store's stock is not read by a gauge on the shelf. Software keeps a stock record from the deliveries received and the sales scanned at the tills, and the tools that forecast demand and trigger replenishment read that record and the sales. A gap between the recorded and the actual stock is called inventory record inaccuracy, and earlier studies found 20 to 70 percent of records wrong. When the record shows more than the shelf holds, the system orders too late, and the shelf can empty without the system knowing. When it shows less, the system orders milk that is not needed. In a study of grocery stores, perishable items had the larger inaccuracy, and the stores corrected their records with physical counts up to four times a year [20].

What Moves It

The final element of every loop is an order: the supermarket's and the depot's order rates and the plant's bottling plan, each rounded to a quantum and capped, as Figure 11.5 draws them. An order becomes milk on refrigerated trucks, which the model reduces to a fixed transit of n_t days with every carton aging by the calendar. The Grade "A" Pasteurized Milk Ordinance governs that cold chain, from cooling the milk before it is filled to the temperature on the delivery vehicles [21].

TECHNOLOGY Final Element: The Order and the Refrigerated Truck

The same stock management system releases the order: it triggers orders automatically, and a store manager may adjust an order or trigger one by hand [20]. The milk the order calls for moves in refrigerated vehicles. The Grade "A" Pasteurized Milk Ordinance requires every vehicle that carries pasteurized milk to be built and operated so that the milk stays at 7 °C (45 °F) or less, and every refrigerated room that stores it to have a thermometer in its warmest zone [21]. A truck that runs warm does not spill the milk. Higher storage temperatures speed microbial growth in pasteurized milk and shorten the time it takes to spoil [18]. In this chapter's model no truck runs warm.

Where the Controller Runs

In a store the ordering rule runs inside the stock management system, which triggers orders automatically while the managers keep the option to adjust an order or to trigger one by hand [20], the role of the MANUAL stations of Section 11.4. No source read for this section prints the cycle of such a system, so none is given here. The app executes every rule each hour and lets a decision through only at midnight on the daily calendar, as shipped, or every hour or every Monday on the other two. The chapter's structure decides where the rules sit: three decentralized rules, one per stage, each with its own stock, supply line and forecast, or one planner that reads all eight states and the shared sales at each midnight and sets the three orders together.

What Is Modernizing It

Two developments bear on this chapter. Sensors and models aim to predict the shelf life left in fresh dairy products: spectroscopic, imaging and sensor-based methods give analytical fingerprints of quality, and chemometric and machine-learning methods relate those signals to quality and to microbial behavior [18]. The model instead pulls every carton at one code date. Radio-frequency identification, used during stock counts or to track stock movements, has lowered record inaccuracy substantially but not completely, and a grocery study found that counts correcting records that overstate the shelf raise sales, so the tags serve as a complement to targeted counts rather than a substitute [20]. Both feed the rules of this chapter better measurements: a stock that is really on the shelf and an age that the milk really has.

IN PRACTICE Date Codes Set by Practice, Not One Rule

In the United States, federal law generally does not require or regulate date labels, infant formula excepted. A review of 2013 found that most states that require dates apply them to specific foods such as milk and dairy, very few define what the words mean, and virtually none set how the date is determined. A sell-by date serves the retailer's stock control [22]. On the shelf, perishable goods add first-in, first-out rotation, date checks, markdowns and the removal of expired goods to a store's routines, and each step is a further chance for a record error [20].

11.8 Limitations

The chain is one product, one store, one depot and one plant, and its managers run one rule with teaching gains, while real people vary from one another and from round to round. Ages move once a day, the bottling day is not aged, and quality is a single code date rather than a function of temperature and handling. The farms' milk is smooth and compulsory, the spot market unlimited up to a fixed amount at a fixed premium, and the cost ledger leaves out capital, energy and fluid revenue. The MPC's daily prediction is linear, knows no carton ages and replaces expiry by freshness ceilings, so its comparison with the rules is a seeded teaching result, not evidence of general optimality. The variance floor holds only under the assumptions of case 4, and a single run's ratio is an estimate that scatters around its expected value, by up to 1.5 % below the floor at 52 weeks. Every threshold, the challenge's three among them, is a teaching value.

OPEN CHALLENGE **Stock Records Nobody Has Counted**

Every rule and the planner of this chapter read their stocks exactly. A real one reads a record that drifts from the shelf, and the drift cannot be seen when it occurs: it is measured only afterward, by a deliberate count. In one grocery study a store-wide count lifted sales by about 11 percent, all of it on items whose record had shown more stock than the shelf held, and more on perishable items. Radio-frequency tags improve visibility but do not remove the errors, so the study expects most from selective technology combined with focused counts and disciplined routines [20].

11.9 Conclusion

The dairy chain gives an ordering mistake a physical destination, a shelf of cartons near their code date or a drain under a full dryer. With the shipped human rule a one-week promotion becomes a depot order of 65.5 kL/d and then nothing, a plant ratio of 12.70, 32.5 kL of milk poured out and 30.5 kL of raw milk drained, and counting the supply line pours none. Order-up-to under the analytical assumptions sits at or above the local floor at every window from 3 to 28 days, 11.36 against 8.58 at a week, while shared sales take it to 2.73. A 10-day shelf life multiplies case 2's waste by eight, and the MPC holds a cut bottling line for 312 hours without a lost sale. In the challenge the shipped rule costs $898.2k with a 44-hour stockout, each cure alone fails, and both together pass at $791.3k and 6 hours. The chain is aggregated and its values are for teaching. A fuller model would add a network of stores, quality driven by temperature, a stochastic farm supply and planners that see the ages of their stock. The simulator makes the bullwhip a measurement of milk, money and hours rather than a ratio alone.

KEY POINTS

- A manager who forgets the milk already on order orders the same milk twice, and the error grows from stage to stage.
- With the shipped human rule a one-week promotion becomes a depot order of 65.5 kL/d and then nothing, a plant ratio of 12.70 and 32.5 kL of milk poured out, and counting the supply line pours none.
- Order-up-to under the analytical assumptions sits at or above the local floor at every window, 11.36 against 8.58 at a week, while shared sales take it to 2.73.
- In the challenge each cure alone fails, and both together pass at $791.3k with a 6-hour stockout.

Where else this applies. The product startup of Volume 4, Chapter 8 sets its production with the same anchoring-and-adjustment rule, its supply-line term counting the work in process, decided on a quarterly calendar as shipped.

References

[1] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[2] H. A. Simon, "On the application of servomechanism theory in the study of production control," *Econometrica*, vol. 20, no. 2, pp. 247–268, 1952.

[3] H. J. Vassian, "Application of discrete variable servo theory to inventory control," *Journal of the Operations Research Society of America*, vol. 3, no. 3, pp. 272–282, 1955.

[4] D. R. Towill, "Dynamic analysis of an inventory and order based production control system," *International Journal of Production Research*, vol. 20, no. 6, pp. 671–687, 1982.

[5] J. Dejonckheere, S. M. Disney, M. R. Lambrecht, and D. R. Towill, "Measuring and avoiding the bullwhip effect: A control theoretic approach," *European Journal of Operational Research*, vol. 147, no. 3, pp. 567–590, 2003.

[6] E. Perea-López, B. E. Ydstie, and I. E. Grossmann, "A model predictive control strategy for supply chain optimization," *Computers & Chemical Engineering*, vol. 27, no. 8–9, pp. 1201–1218, 2003.

[7] M. W. Braun, D. E. Rivera, M. E. Flores, W. M. Carlyle, and K. G. Kempf, "A model predictive control framework for robust management of multi-product, multi-echelon demand networks," *Annual Reviews in Control*, vol. 27, no. 2, pp. 229–245, 2003.

[8] D. Swaroop and J. K. Hedrick, "String stability of interconnected systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 3, pp. 349–357, 1996.

[9] S. Nahmias, "Perishable inventory theory: A review," *Operations Research*, vol. 30, no. 4, pp. 680–708, 1982.

[10] R. Kaipia, I. Dukovska-Popovska, and L. Loikkanen, "Creating sustainable fresh food supply chains through waste reduction," *International Journal of Physical Distribution & Logistics Management*, vol. 43, no. 3, pp. 262–276, 2013.

[11] A. Guarnaschelli, H. E. Salomone, and C. A. Méndez, "A stochastic approach for integrated production and distribution planning in dairy supply chains," *Computers & Chemical Engineering*, vol. 140, p. 106966, 2020.

[12] I. J. Martinez-Moyano, "History of the beer game," *System Dynamics Review*, vol. 40, no. 2, p. e1767, 2024.

[13] J. D. Sterman, "Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment," *Management Science*, vol. 35, no. 3, pp. 321–339, 1989.

[14] R. Croson and K. Donohue, "Behavioral causes of the bullwhip effect and the observed value of inventory information," *Management Science*, vol. 52, no. 3, pp. 323–336, 2006.

[15] F. Chen, Z. Drezner, J. K. Ryan, and D. Simchi-Levi, "Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information," *Management Science*, vol. 46, no. 3, pp. 436–443, 2000.

[16] J. W. Forrester, "Industrial dynamics: A major breakthrough for decision makers," in *The Roots of Logistics*, P. Klaus and S. Müller, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 141–172.

[17] J. C. Buzby, H. F. Wells, and J. Hyman, "The estimated amount, value, and calories of postharvest food losses at the retail and consumer levels in the United States," U.S. Department of Agriculture, Economic Research Service, Economic Information Bulletin EIB-121, Feb. 2014. [Online]. Available: https://www.ers.usda.gov/media/8679/eib-121.pdf

[18] L. Lima, A. I. Pereira, M. Pintado, and M. Carocho, "Fresh dairy quality in the era of smart food systems: Deterioration pathways, analytical technologies, and predictive modeling," *Foods*, vol. 15, no. 19, p. 3439, 2026.

[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] Y. Rekik, R. Oliva, A. A. Syntetos, and C. H. Glock, "Inventory record inaccuracy in grocery retailing: Impact of promotions and product perishability, and targeted effect of audits," *Journal of Business Logistics*, vol. 47, no. 3, p. e70079, 2026.

[21] *Grade "A" Pasteurized Milk Ordinance (PMO), 2025 Revision*, U.S. Food and Drug Administration and National Conference on Interstate Milk Shipments, 2025. [Online]. Available: https://www.fda.gov/media/193438/download

[22] E. Broad Leib, D. Gunders, J. Ferro, A. Nielsen, G. Nosek, and J. Qu, "The dating game: How confusing food date labels lead to food waste in America," Natural Resources Defense Council and Harvard Food Law and Policy Clinic, Tech. Rep. R:13-09-A, Sep. 2013. [Online]. Available: https://www.nrdc.org/sites/default/files/dating-game-report.pdf

11.10 Exercises

11.1 Concept check · 2 min

In case 2 the supermarket's shelf is still short a day after a large order left the store, although that order is already on a truck. The rule of Equation (11.16) orders more milk again. Which change to the rule stops it ordering the milk on the truck a second time, while it still corrects the stock gap?

(a) Lower the stock adjustment a_S, so that each day closes a smaller share of the stock gap

(b) Raise the supply-line weight β to 1, so that every liter on order is counted

(c) Lower the forecast smoothing θ, so that the forecast follows the orders more slowly

(d) Raise the order quantum q_O, so that small corrections round away to zero

11.2 Concept check · 2 min

In case 4, order-up-to on a 7-day moving average with local information, the plant's variance ratio over 26 weeks is 11.19 against a floor of 8.58 from Equation (11.22). Which single change takes the plant's ratio below the local floor computed for the setting in use?

(a) Raise the moving-average window from 7 to 14 days
(b) Lower every stage's stock target from 2 or 3 days of demand to 1 day
(c) Forecast every stage from the supermarket's shared sales
(d) Count the supply line with $\beta = 1$ at every stage

11.3 Concept check · 2 min

At 5.75 weeks in case 2 the plant bottles 1.0 kL/d while the farms deliver 30 kL/d, the dryer's capacity is 20 kL/d and spot milk is off. Where do the other 29 kL/d go?

(a) Back to the farms, which hold their milk until the bottling plan rises again
(b) Into the cold store, where they wait to be bottled when the orders return
(c) To spot buyers, who take the surplus at the premium c_s
(d) 20 kL/d to powder and 9 kL/d down the raw-milk drain

11.4 Theory and calculation · 5 min

Take the shipped human rule, $a_S = 0.26/\text{d}$ and $\beta = 0.34$, in Equations (11.15) and (11.16).

(a) Compute the gain $g(a_S, \delta)$ for daily, weekly and hourly decisions, $\delta = 1$, 7 and $1/24$ day. *Checkpoint:* 0.260, 0.1255 and 0.2992 per day.

(b) Show that a gap left unchanged for a week is closed by the same share on all three calendars, and give that share. *Checkpoint:* $1-(1-a_S)^7 = 0.878$.

(c) The supermarket forecasts 20 kL/d. Its stock is 10 kL below target, and the order it raised yesterday has put its supply line 10 kL above the $\lambda_S F_S$ it wants. Compute the daily order of the human rule and of the supply-line-aware rule, after the quantum of 0.5 kL/d. *Checkpoint:* 21.5 and 20.0 kL/d.

11.5 Theory and calculation · 6 min

Use Equation (11.20) with daily decisions.

(a) Show that $\Gamma(1) = 1$ for every choice of a_S, a_{SL}, λ and forecast filter, and say what it means for a lasting change in demand. *Checkpoint:* $\Gamma(1) = a_S/a_S = 1$.

(b) For order-up-to, $a_S = a_{SL} = 1$, show that the denominator is one. Then evaluate $|\Gamma|$ at $\omega = \pi$ for the depot, $\lambda = 2$, with a 7-day moving average, and give the period of that wave in days. *Checkpoint:* $13/7 = 1.857$, a 2-day wave.

(c) Evaluate $|\Gamma|$ for the same rule at $\omega = 2\pi/7$. *Checkpoint:* 1.

11.6 Theory and calculation · 6 min

Use Equation (11.22) for case 4: lead times plus the review day $L = 3$, 3 and 2 days at the supermarket, the depot and the plant.

(a) Compute the local floors of the three stages at $p = 7$. *Checkpoint:* 2.22, 4.95 and 8.58.

(b) Compute the plant's shared-sales value and the centralized value $1 + 2\sum L/p + 2(\sum L)^2/p^2$ at $p = 7$, and say why the chain's shared value is the smaller. *Checkpoint:* 2.76 and 5.90.

(c) Case 4's completion rule asks for a run at $p \geq 14$ with the plant's ratio at least 95 % of its floor. Compute the floor at $p = 14$ and check the measured 3.34 against the rule. *Checkpoint:* Floor 3.07; $3.34/3.07 = 1.09$.

11.7 App challenge · 8 min

Scenario 1, Play the supermarket, default seed. The supermarket's orders are in Manual at 20 kL/d, and the promotion starts at the midnight that opens day 14 (week 2.00).

(a) Before running, predict the day on which the first milk of an order raised at that midnight reaches the shelf, and the day for one raised at noon on day 14. *Checkpoint:* Day 16 and day 17.

(b) Press Run and leave the order at 20 for 8 weeks at 1 week per second (8 s). Read the Longest supermarket stockout tile, and the last value of the exported `service_fraction` column. Convert the service into lost sales, given that the shoppers want 1260 kL over the 8 weeks. *Checkpoint:* 96 h; 0.9365, so 80 kL (within 0.5 kL).

(c) Reload the case. In Advanced set Turn-based: pause at every midnight to On. Press Run at each pause. At the pause at day 14 set Supermarket order to 40, and at the pause at day 19 (week 2.71) set it back to 20. Then set Turn-based to Off and run to week 8. Read the same two values. *Checkpoint:* 24 h; 0.9948, so 6.5 kL (within 0.5 kL).

11.8 App challenge · 8 min

Scenario 2, The bullwhip, default seed, at 1 week per second.

(a) Run 8 weeks. Read the three stages' bars on the Bullwhip view (or the last row of `S_bullwhip`, `D_bullwhip` and `P_bullwhip`), the Raw milk discarded tile and the last value of `finished_waste_kL`. *Checkpoint:* 1.90, 6.36, 12.70; 30.5 kL; 32.5 kL (all within 0.1).

(b) In the Ordering policy card set Decide to Hourly, press Reset and run 8 weeks again. Read the plant's bar, the Longest supermarket stockout tile and the last `service_fraction`. *Checkpoint:* 13.89; 28 h; 0.9934.

(c) Explain why deciding every hour cuts the lost sales but raises the plant's ratio, and name the setting that would cut both. *Checkpoint:* The supply-line weight, $\beta = 1$.

11.9 App challenge · 8 min

The challenge, Twenty-six weeks (the tenth case in the list), default seed (7), at 4 weeks per second (6.5 s per run).

(a) Run the shipped managers for 26 weeks. Read the Total cost including diversion tile, the Longest supermarket stockout tile, the Challenge tile and the last `finished_waste_fraction`. *Checkpoint:* \$898.2k, 44 h, not met, 0.0372.

(b) Press Reveal suggested tuning, then Reset, and run 26 weeks. Read the same values. *Checkpoint:* \$791.3k, 6 h, met, 0.0128.

(c) Keep those settings, set Ordering rule to Centralized MPC, press Reset and run 26 weeks. Read the same values and the MPC solves / failures tile. *Checkpoint:* \$774.4k, 0 h, met, 0.0047; 182 / 0.

Give costs within \$0.1k.

11.10 Application area · 4 min

Case 5 runs 16 weeks (112 days) through the spring flush. The farms send 3700 kL in all, peaking at 34.5 kL/d on day 69, while bottling stays between 18.0 and 22.0 kL/d and 1486 kL is dried.

(a) Compute the mean powder rate and the dryer capacity needed at the flush peak, and say why a 12 kL/d dryer must pour raw milk away. *Checkpoint:* 13.3 kL/d; 12.5 to 16.5 kL/d.

(b) With the dryer cut to 12 kL/d, 167.0 kL of raw milk is poured away. Price it at c_w and compare with the rise in operating cost the chapter reports. *Checkpoint:* \$100.2k, the whole rise from \$157.2k to \$257.4k.

11.11 Application area · 4 min

In Table 11.5 the shipped managers lose the least margin to the dryer, \$468.1k against \$511.2k for the suggested tuning. The powder margin loss is $m_f - m_p = 0.25\,\$/\mathrm{L}$ and milk poured out costs $c_w = 0.60\,\$/\mathrm{L}$.

(a) Convert both diversion losses into kiloliters of powder and give the difference. *Checkpoint:* 1872.4 and 2044.8 kL; 172.4 kL.

(b) A plant manager concludes that the shipped managers use the farms' milk best. Explain where the milk the shipped chain did not dry went, and why that is the more expensive outlet. *Checkpoint:* Bottled and later poured out, or drained raw, at 0.60 $/L.

11.12 Review the agent · 5 min

An AI assistant reviews case 2 and writes: "The bullwhip comes from managers who react too slowly. Raise the stock adjustment a_S from 0.26 to 1/d at every stage, so each stage closes its stock gap in one day. The shelf will be empty less often, and the swings will die out sooner." Find the flaw, say how you would check it in the app, and propose a better change.

CHAPTER 12

Summary and Outlook



Open gaps 1 one-way coupling 2 decoupler cost 3 switching structures 4 downstream variables 5 weak directions 6 perishable chains

Notice that the gap of switching structures joins the reverse osmosis skid, the recycle plant and the Tennessee Eastman plant, the three units that held a limit only by changing what they controlled. Illustration, not to scale.

IN PLAIN TERMS **What Ten Plants Taught About Structure**

The main idea. Each chapter of this book ran a plant with several handles and several results to hold, from a seawater plant to a dairy chain. Before any tuning, someone had to decide which handle answers to which result, and where the plant's output is set.

The challenge. A wrong choice failed however it was tuned. A common test of the pairs missed a pull that runs one way only, and a fix for one result could cost another far down the line.

Why it matters. Testing each choice against a real upset, judged where the next unit feels it, told the designs that worked from those that failed.

Ten chapters put the structure of a control system on trial before any loop was tuned. Each plant had more than one actuator and more than one variable worth holding, and in each a wrong arrangement failed in a way no tuning could repair. A pairing drove a train to a valve stop, a fixed fresh feed flooded a column, and a full inverse lost the orbit to one reversed corrector. This chapter sets the results side by side with the measure that predicted each, draws a short guide for choosing a structure, 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.

12.1 What the Chapters Showed

Table 1.2 of Chapter 1 listed the structural question of each chapter and the measure it reports, without numbers. Tables 12.1 and 12.2 add the results. The first holds the unit-level chapters, where several actuators act on one unit, one line or one ring. The second holds the plant-level chapters, where the question is where production is set and where the upsets go. Each row names the difficulty, an arrangement that worked, one that failed, and the measure with the measured result, cited to the table, equation or section of the chapter that reports it. Figure 12.1 sets one like pair of numbers from each row side by side.

The columns of the two tables make five points.

The relative gain rejected the pairings that fail outright. Each pairing whose relative gain was negative failed whatever the tuning. The swapped pairing of Chapter 3, at $\Lambda_{12} = -0.517$, made 15.5 times the integrated error of the natural pairing and 13.6 t of off-specification paper against 2.1 t. The reversed pairing of Chapter 2 drove the valve to its stop. The array did not report the interaction that remained. The reverse osmosis gains are triangular, so $\lambda_{11} = 1.00005$ while the pump moves the recovery by 0.703 %/% of speed, and the cabin of Chapter 10 couples its two gas loops through

Table 12.1 Pairings, Decouplers and Inverses, Chapters 2 to 6

Chapter	Difficulty	What worked	What failed	Measure and result
2, reverse osmosis train	The valve's gain on permeate flow is –0.001 in $m^3/h/\%$ at design and changes sign either side. The pump moves recovery strongly.	Pump on production, valve on recovery, one static feedforward from the pump to the valve at the gain ratio 0.914.	Production on the valve. The train ends at 82.5 m^3/h against a 500 m^3/h setpoint (Section 2.6).	$\lambda_{11} = 1.00005$ (Equation (2.14)). Peak recovery error 0.937 points, 0.213 with the feedforward, 24.35 swapped (Table 2.3).
3, paper machine	Weight and moisture both respond to stock and steam. The scanning gauge adds a scan to the dead time, 21.9 s for moisture.	Stock on weight, steam on moisture. The stock-to-steam half of a static decoupler. A 10 s scan.	The swapped pairing, still walking 22 minutes after a dilution, at 72.8 g/m^2 and 13.12 % moisture.	$\Lambda_{11} = 1.517$ (Equation (3.25)). Moisture excursion 1.80 to 0.23 points with the decoupler, off-specification paper 2.38 to 4.54 t (Table 3.5).
4, milk evaporator	Solids are read at the end of the vapor chain, 380 s after a steam step (Section 4.4).	Steam on solids, a ratio feedforward compensated for feed solids, averaging level loops.	Feed on solids, which cuts throughput. Tight level loops, which pass boiling noise to the dryer.	One point of solids costs 0.12 t/h of feed or 0.030 t/h of steam. Shift: 0.64 t off and 155.5 t fed as shipped, none off and 160.0 t suggested (Table 4.7).
5, coating line	Line speed is the throughput handle. A speed change disturbs both spans and the coat weight, which a gauge reads 19 m downstream.	Pump feedforward from the speed and inertia compensation. For a late batch change, a lead on the pump and a 25 s trim.	A speed ramp with the pump flow held. A trim that waits for the gauge.	Scrap after a speed change 208.90, 22.30 with the feedforward and 12.80 m with both (Table 5.6). Case 9: 63.00 loops, 21.05 MPC, 20.55 m with the lead (Table 5.8).
6, storage ring	Singular values fall from 379 to 0.94, a condition number of 405.2, with a cliff at mode 48 (Section 6.4).	The inverse truncated at the cliff or regularized. A fast set, a monitor residual test and a feedforward.	The full inverse. It moves the magnets most and fails with one corrector wired backward.	Jitter 0.198 against 0.0164 µrad, orbit 0.09 against 0.21 µm. Reversed corrector: 132.6 to 136.0 µm full, 0.38 at 48 modes (Table 6.4).

one sum with a relative gain of one. A triangular plant needs its off-diagonal gains read one by one, because the array multiplies them together.

A decoupler is priced by its other loop. The static decouplers of Chapter 2 and Chapter 3 cut the disturbance they targeted by a factor of four or more, from 0.937 to 0.213 points of recovery and from 1.80 to 0.23 points of moisture. On the paper machine the stock-to-steam half carried almost all of the benefit, and the full decoupler raised the off-specification paper from 2.38 to 4.54 t, because the weight loop now carried a steam-driven bias while the production loop slowed the machine. On the reverse osmosis train the decoupler gain had a minimum, and a gain of 1.6 left a peak of 0.864 points, close to the 0.937 of no decoupling at all (Table 2.4). Section 1.4.6 found the same trade on the running example.

Where production is set decides where the upset goes. With the fresh feed fixed, 10 % more feed raised the recycle of Chapter 7 by 62.3 % in 12 h and left a vapor margin of 3.1 %. Setting production at the reactor holdup gave about the same production with the recycle 4.2 % lower, at the price of 14 % more holdup for each 10 % of production. The evaporator could hold its solids on steam or on feed at nearly the same speed, but the feed pairing spent throughput. On the coating line the speed was the throughput handle, and the scrap of a speed change fell from 208.90 to 12.80 m once the

Table 12.2 Production Handles, Inventories and Chains, Chapters 7 to 11

Chapter	Difficulty	What worked	What failed	Measure and result
7, reactor and column with recycle	The recycle magnifies a change of fresh feed. The column floods 10.7 % above design production (Section 7.4).	CS2, production set by the reactor holdup. In the challenge, CS2 with the level at 75 %.	CS1 with the fresh feed on flow. CS3 with a fixed recycle, which floods at 4.47 h.	Snowball elasticity 4.61 (Table 7.3). Recycle up 62.3 % on CS1, down 4.2 % on CS2 (Table 7.6). Challenge fails on CS1, passes on CS2 (Table 7.7).
8, Tennessee Eastman plant	Unstable open loop, a trip at 3.59 h. Gas recycle with an inert and a purge.	Inventory loops first, composition loops above them, pressure at 2800 kPa g. An override and a production cut on the A feed loss.	Inventory loops alone on a feed change. The decentralized structure as shipped on the A feed loss.	Trips at 5.84 and 8.40 h, none with composition loops or with the override and cut (Table 8.5). Challenge cost 215.7 at 2705 and 187.3 $/h at 2800 kPa g (Table 8.7).
9, peptide purification	The chromatography cut is judged at the filter. Late variant in the pool shrinks the crystals and slows filtration.	A downstream-aware cut that stops on the predicted late variant, or a fixed window at the times it realizes.	A cut that maximizes recovery at the pool's purity limit. The UV threshold rule.	Late variant 0.430 against 0.698 %, filter 46.8 against 62.0 min, 59.9 g acceptable against none. Purity-maximizing cut 43.4 g (Table 9.6).
10, sealed cabin	Oxygen and nitrogen loops share one total pressure, and the machines can only add gas. Eclipses limit power.	A power-aware schedule that keeps electrolysis out of the dark. A nitrogen balance that flags a leak.	The week without the balance, or on the adaptive schedule.	Relative gain one. 1 kPa more oxygen shuts the nitrogen valve 2.52 h later (Section 10.4). Week passed at 11 of 11 seeds, at 0 of 11 by either failure (Table 10.5).
11, dairy supply chain	Each stage orders from the next one up. Orders lag, milk spoils and the plant must take its farms' milk.	Counting the supply line ($\beta = 1$) and sharing the supermarket's sales. The MPC with shared sales.	The shipped human rule, $\beta = 0.34$.	Peak amplitude ratio 1.89 (Section 11.5). Plant variance ratio 12.70, 7.61 at $\beta = 1$ (Table 11.4). Challenge 898.2 k$ and 44 h of outage, 791.3 and 6 suggested (Table 11.5).

pump and the winders were driven from it. In Chapter 8 the inventory loops held the plant until a feed changed, and only the composition loops above them carried it through.

Weak directions carry noise and model error. The full inverse of Chapter 6 made the slow correctors jitter 12 times as much as a 48-mode truncation, for an orbit of 0.09 instead of 0.21 µm. The disturbance decided what the weak modes were worth. The girder drift put 86.0 % of its power in modes 45 to 48 and 7.3 % in the four weakest. With one corrector wired backward the full inverse lost the orbit at every bandwidth tried, while the truncation at 48 modes held 0.38 µm. A large condition number warns about the noise and about the model at once.

A chain is judged at its far end. Three chapters judged a decision by its effect downstream. The shipped ordering rule of Chapter 11 has a peak amplitude ratio of 1.89, at a wave of 15.9 days, so a swing near that period grows at each stage, and counting the supply line cut the plant's variance ratio from 12.70 to 7.61. Tight level loops in the evaporator passed the boiling noise on as a product flow with a standard deviation of 0.60 t/h, against 0.023 t/h for averaging loops (Section 4.6). The cut of Chapter 9 that maximized recovery at the pool's purity limit lost the batch at the filter. In each case the measure that predicted the result was computed at the end of the chain, not at the decision.

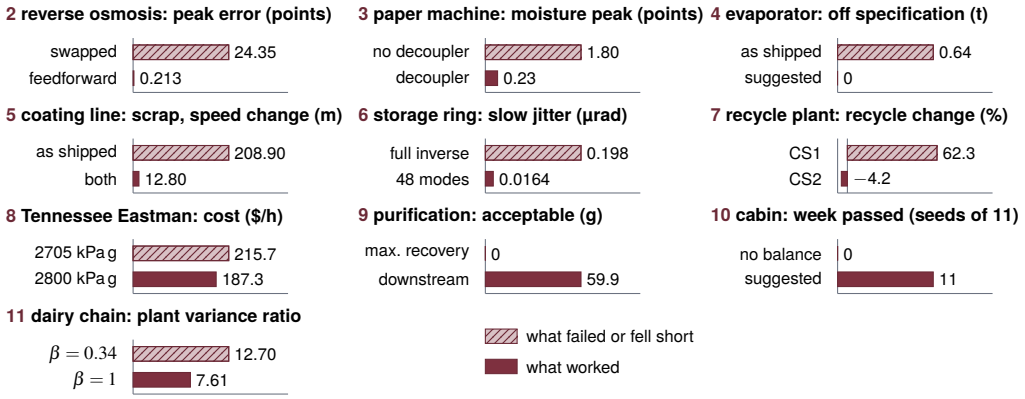

Figure 12.1 One like pair of numbers from each row, as the "Measure and result" column of Tables 12.1 and 12.2 prints them. Each panel has its own scale. The hatched bar is the structure or setting that failed or fell short and the solid bar the one that worked; a value of zero has a label and no bar. Lower is better on every panel except the purification and the cabin. Notice that no better number came from a faster loop alone: each came with a different pairing, feedforward, decoupler, production handle, inverse, cut rule, balance or ordering rule, or, on the Tennessee Eastman plant, a higher pressure setpoint.

12.2 Choosing a Structure

The chapters suggest an order of questions for a new plant. Each question names the chapter where its answer decided the result.

1. *How many degrees of freedom are there, and where is production set?* Count the valves that affect a steady state (Section 1.4.2), place the throughput handle, and point every level loop away from it (Section 1.4.3). Close the inventory loops first. The plant of Chapter 8 trips within hours without them.

2. *Does a recycle magnify a change of production?* Compute the snowball elasticity and the capacity margin (Section 1.4.4). Where the margin is small, set production at a flow inside the loop or at a holdup, as CS2 does in Chapter 7.

3. *What does the next unit respond to?* Hold the variable that the downstream unit feels, as the cut of Chapter 9 holds the predicted late variant rather than the pool purity.

4. *What is the relative gain at the operating point, and does any gain change sign across the range?* Use a central difference at the design point. Reject a negative relative gain. Where the array reads one, check whether the plant is triangular. The reverse osmosis valve gain runs from +0.938 at 46 % to −1.045 at 54 % opening (Table 2.2).

5. *Which half of a decoupler carries the benefit, and what does the other loop pay?* Run each half alone and sweep the gain, as Chapter 3 and Chapter 2 do.

6. *Is the gain matrix ill-conditioned?* Read the singular values and project the disturbance on the modes before choosing how many to correct (Chapter 6).

7. *Does a chain of inventories amplify?* Compute the amplitude ratio of each level loop or ordering rule (Section 1.4.8). Tune levels as buffers where a downstream unit feels the flow, and count what is in transit (Chapter 11).

A predictive controller paid for itself where it carried a delay or a constraint that the loops did not see, and even there simple elements often matched it. On the coating line it scrapped 21.05 m against

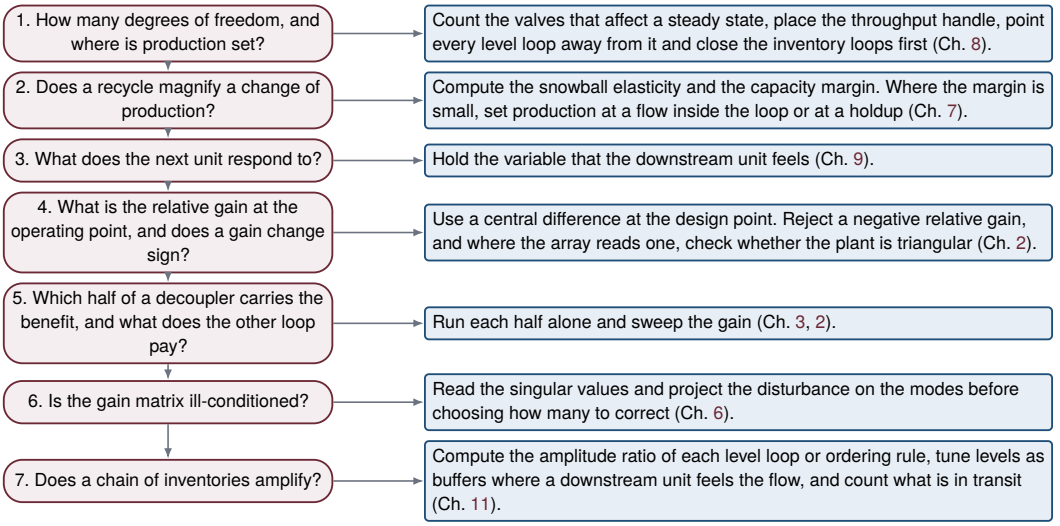

Figure 12.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the production handle and the recycle come before any pairing is chosen: a fixed fresh feed flooded the recycle plant's column, and the Tennessee Eastman plant tripped without its inventory loops.

63.00 m for the loops when a thinner batch and a faster order arrived together, and loops with a lead on the pump and a 25 s trim scrapped 20.55 m. On the reverse osmosis train it held the feed pressure at 69.81 bar where the loops spent 108.6 min above the 70 bar rating, and a pressure override on the recovery target held 69.87 bar (Table 2.8). On the paper machine's grade change it made 2.68 t of off-specification paper against 4.32 t for the shipped decoupler and lost to the decoupler retuned by the same rule, at 1.51 t (Table 3.7), chiefly because its model holds the speed that the production loop moves. In the evaporator it cut the integrated error of a throughput step from 54.92 to 20.28 % min and put 0.466 t off specification on a shift the loops passed (Table 4.8). Volume 9 treats that choice in full. Figure 12.2 draws the seven questions in order.

12.3 Where Practice Is Heading

Structure before algorithm. Plantwide control asks which variables to control, which to measure, which inputs to manipulate and which links to make between them, questions that most control theory assumes are answered at the outset [1]. A recent review argues that decomposing a control system into simple elements such as cascade, ratio, feedforward, decoupling and selectors is a powerful way to design control for complex processes with limited information. It challenges the view that model predictive control can solve every control problem and calls for more teaching of these elements [2]. The chapters of this volume point the same way on their own cases. The retuned decoupler of the paper machine beat the predictive controller, and the pressure override of the reverse osmosis train matched it.

Decomposition at plant scale. Large control and optimization problems can be decomposed by analyzing their network for its latent block structure, which then determines the decomposition used for distributed control [3]. Distributed predictive controllers built on recurrent neural network models have been shown, on a simulated chemical process network, to compute more efficiently than a centralized controller while performing very close to it [4]. In the editor's view, partitioning a

plant for such controllers is the structure question of this volume asked again, and the measures of Chapter 1 remain a way to check a proposed partition before it is trusted.

> ⧖ HISTORY **2004 to 2023: From a Design Procedure Back to Simple Elements**
>
> The history boxes of Chapter 1 end with a review of 2000. A procedure published in 2004 split the design of a plant's control system into a top-down analysis and a bottom-up design of the control layers, and its author wrote that more work was needed to understand how to decompose and coordinate those layers [5]. In 2011 an industrial and an academic author found that much of the research had still not gained enough acceptance to influence industry profoundly, and stressed the need for a formalized yet simple approach [6]. In 2023 a review proposed that far more academic research go into the theory of the simple control elements that industry already uses [2]. In the editor's reading, the ten chapters of this volume stand at that point: each structure that worked was built from such elements and checked by a run.

Benchmarks that repeat. The Tennessee Eastman plant has served as the benchmark in comparison studies of data-driven fault diagnosis and process monitoring [7]. A revision of its code found that changes in simulation software since the 1990s broke the repeatability of simulation results, and presented a modified code [8]. Chapter 8 met the same sensitivity in its own port, where the open-loop trip came at 3.47, 3.59 or 4.06 h depending on how the constants were rounded (Section 8.4). In the editor's view, a plantwide benchmark is useful only with its code, its numerical precision and its scoring published together.

Information along supply chains. When each company in a chain orders from the one upstream, the variance of orders can exceed that of sales and grows moving upstream, and actions to mitigate the distortion have been discussed since 1997 [9]. The dairy chain of Chapter 11 measured the same growth in milk, and sharing the supermarket's sales took the plant's variance ratio under order-up-to from 11.36 to 2.73 at a seven-day window (Section 11.5).

Feedback by mode on large arrays. Orbit feedback at light sources corrects a large array through the modes of its response matrix. A study of one feedback system reports that Tikhonov regularization, which lowers the gain of the weak modes, is used in the feedback systems of many light source facilities [10]. A controller with its own dynamics for each mode was implemented on a booster ring run with stored beam as a test bed, because the two extra matrix products per update are usually too heavy for a storage ring with more than a hundred actuators [11]. For beams of the order of 4 µm, a design study began to look beyond the classic PID regulator to model-based methods and stated that a reliable model is required [12]. The truncated and regularized inverses of Chapter 6 belong to the first of these steps.

Measurements that see what the next unit feels. Process analytical technology has become a pillar of process control for numerous small-molecule drugs, while the close chemical similarity of protein contaminants makes suitable methods hard to find. In-line soft sensors that combine spectroscopy with chemometrics have reached the needed detection limits for some contaminants at laboratory scale, and the review that reports them names closed-loop control as a next step [13]. On paper machines, a laser ultrasonic sensor of sheet stiffness was tested at three industrial sites and placed at the demonstration stage [14]. The soft sensor of Chapter 9 predicts the late variant at the column, while the filter decides the batch. In the editor's view, a sensor of what the next unit feels lets a structure hold that variable directly.

Treatment trains as control networks. A roadmap for water treatment observes that most systems need operators to adapt to changing water quality and component degradation, and that the data of their many sensors are used mainly to detect failures. It calls for sensing and control networks with

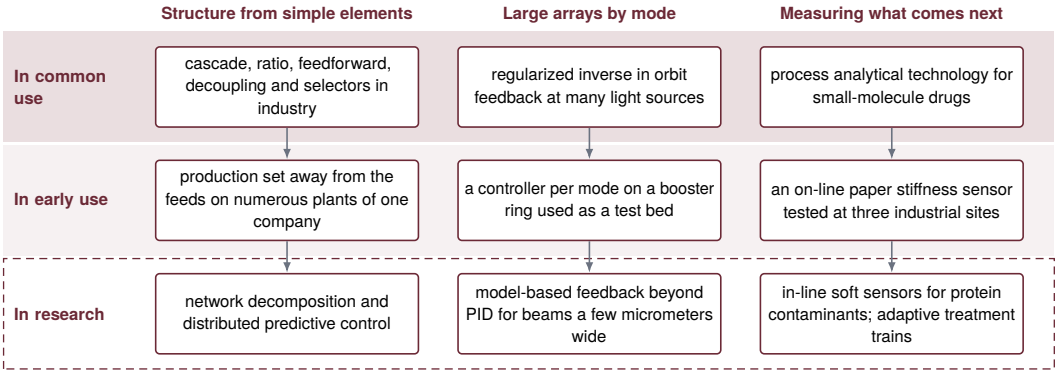

Figure 12.3 Developments in three threads. The top lane holds what the cited source calls commonly used or used at many facilities, the middle lane what it reports in service at some plants or on a test bed, and the bottom lane what it reports from simulation, at laboratory scale or as a need, in the editor's reading. Sources by thread: structure from simple elements [2–4, 6], large arrays by mode [10–12] and measuring what comes next [13–15]. Notice that each thread ends in a model that must itself be checked on the plant before a structure can rely on it.

models of the interdependencies between unit processes, and notes that the savings of predictive control are hard to quantify without actual facilities [15]. The reverse osmosis train of Chapter 2 models one such dependency, from the pump to the recovery, and leaves the pretreatment out. In the editor's view, the pairing and override questions of this volume return at the scale of a whole treatment train. Figure 12.3 sets these developments in three threads.

Agents that propose structures. A 2026 survey of language models in process systems engineering found genuine promise for tasks built on natural language, such as querying documentation, and found that applications needing real-time execution, constraint satisfaction or formal safety guarantees remain challenging [16]. An agent built to turn industrial specifications and process diagrams into optimization models for process control was evaluated on 50 scenarios, and its authors state that the evaluation assessed the formulation, not closed-loop industrial performance [17]. In the editor's view, agents will soon draft pairings, decouplers and inventory structures in quantity. The checks of this volume, a relative gain at the operating point, a snowball gain, a run of each decoupler half and a run on a disturbance, are how an engineer will accept or reject them.

12.4 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 [18]. 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.

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 [19]. 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.

12.4.1 How This Volume Was Made

Each of the ten apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace numbers to runs, and a referee report was written on the paper. For at least three papers the report is a self-review by the agent that wrote it, and for the reverse osmosis paper it covers only material added in a later round. Defects found along the way went into a log for each app. For this book one agent per chapter reran every number that an exercise uses. The editor set the scope of each round and the rules on what each agent could change, and statements that runs contradicted were left for the editor to decide. The record holds the following errors and the checks that caught them.

The defect logs. The ten logs hold 70 entries with a severity. Of these, 18 are marked as wrong behavior of an app (a wrong output, a blocker, a functional fault or a wrong credit), 47 as documentation, four as cosmetic and one as minor. The reverse osmosis log adds 15 numbered entries without a severity. Seven cases fit the topic of this volume.

A relative gain from the wrong derivative. The reverse osmosis app's briefing computed $\lambda_{11} = 0.99$ from a valve gain of -0.13, a one-sided difference taken away from the maximum. With the briefing's own gains the arithmetic gives 1.007, and the central difference at the design point gives a valve gain of -0.001 and $\lambda_{11} = 1.00005$. The same log found the pump-to-recovery gain quoted as 0.73 in five places where the model gives 0.7029, which moves the decoupler ratio from 0.95 to 0.914. Both were caught by recomputing each quoted number from the shipped model.

A pairing switch that was not bumpless. Choosing the swapped pairing before Run made the first scan command the concentrate valve from 50 to 88 %, the pump speed its new controller had been initialized with, and the permeate fell from 450 to 27.8 m^3/h. The controllers kept their state when the pairing changed. The defect showed only when the pairing was changed in automatic, which the shipped scenarios avoided by starting in manual, and it was reported with the scripts that reproduce it. Each loop is now seeded on the actuator it drives, no shipped scenario run moved, and the swapped row of the paper was measured again.

Two tunings compared as two handles. The first draft of the evaporator paper repeated the app's explanation that the feed-paired loop is slower because the feed must pass the sumps. The self-review identified both handles open loop, found loop gains of 1.32 and 1.0 for the shipped tunings, and reran the feed pairing at the steam loop's loop gain. It then reached the setpoint band only 0.2 minutes behind, 0.3 on the current engine. The speed difference was the tuning, and the lasting difference is the cost in milk.

A production loop that read the plant's truth. The paper machine's production loop computed its tonnage from the true dry fiber flow times a hidden retention, which no instrument reads. Two identical machines, one given a lower retention, differed by 1.36 t/h in measured tonnage before the gauge had posted a pass of the changed sheet. The read acted as a hidden feedforward through the speed. Once the loop measured the scanner's weight times the speed, every automatic result with the production loop in service moved. The decoupled grade move made 1.07 t of off-specification paper before the fix and 4.54 t after, and decoupling became the trade of Section 12.1. A second log, the coating line's, records a trim started from the true viscosity, where no shipped run moved. Both logs document the same probe, which changes a hidden value and asks whether the controller moves while the measurements do not. This is the check the shared section calls reading code for what it reads.

A rule described as its labels intend. The purification app labels a UV threshold "Stop below", and its texts describe a cut that collects until the detector falls below it. The referee pass found that the opening threshold is tested at every scan, so the stop threshold never acts. Settings of 0.10, 0.35 and 0.59 gave identical runs. A final validation over 31 seeds at two step sizes then found that the suggested challenge settings pass at 18 of 31 seeds at the default step. The paper states both, and the fix, which moves every run that uses the rule, awaits the editor.

A claim about every structure. The first draft of the recycle paper said that PI loops hold three levels and two purities under every structure. On CS3 the reflux holds the drum and the recycle purity floats. The self-review read the claim against the paper's own case table and narrowed it.

An offset that looked like a plant property. The agent writing Chapter 1 first clamped the controller outputs of the decoupled example rather than the applied actuators, and the loops held a steady offset that looked like a property of the plant. The catalyst-loss run exposed it, and now only the actuators are clamped.

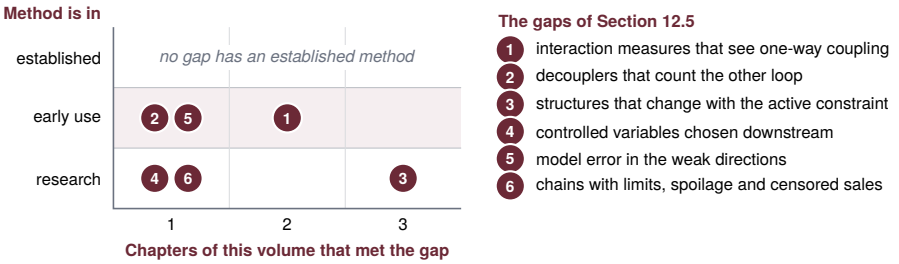

Figure 12.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 2 and 10), one for gap 2 (3), three for gap 3 (7, 8 and 2), one for gap 4 (9), one for gap 5 (6) and one for gap 6 (11). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap met by the most chapters, structures that switch with the active constraint, is still in research.

The book round. The ten chapter records list 162 rows of numeric verification, each a stated value beside its reproduction. Every chapter number rerun reproduced except one already recorded in the paper's notes, a worst posted moisture of 1.25 points on the paper machine's challenge that now runs at 1.244. Two statements contradicted by runs were reported, and both papers and chapters have since been corrected: the recycle product first crosses 97.5 mol % at 7.62 h while the column still floods, and the cabin's case 1 settles at 101.37 %. A printed saturation pressure of 1.002 kPa, which reads 1.0007 kPa in the model, was corrected with them, and the milk app now passes 54 checks where its chapter says 42.

Every error above was caught by a check that could fail: a derivative recomputed at the operating point, a structure switched in automatic, two handles compared at equal loop gain, a hidden value changed while the measurements stayed fixed, a threshold varied, a claim read against a table, a disturbance run as well as a setpoint step. None was caught by rereading the sentence that contained it. Four working rules follow from this record. Compute every structural measure from the model at the operating point and say how. Compare structures at equal loop gain and on a disturbance. Read what each controller reads. Let a person decide what changes.

12.5 Gaps in Multivariable and Plantwide Control

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow them in Section 12.5.1. 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 12.4 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. *Interaction measures that see one-way coupling.* The reverse osmosis train and the cabin both returned a relative gain of one while one loop disturbed the other strongly. A review of plantwide control lists the closed-loop disturbance gain and the performance relative gain array beside the relative gain as tools to analyze interactions in decentralized control [1], and a later review notes that single loops may perform poorly on some interactive processes, where decoupling or multivariable control should be considered [2]. The gap is hard because a measure that multiplies the two off-diagonal gains cannot tell one strong path from none. A routine measure, reported beside the array, is needed that gives the size and direction of one-way coupling and flags a gain that changes sign inside the operating window, as the concentrate valve's gain does.

2. *Decouplers that count the other loop.* The static decoupler of Chapter 3 put the right steam on at the wrong time, because its two paths take 44 s and 84 s, and it doubled the off-specification paper of the loop it did not help. A review of simple control elements suggests one-way decoupling, with

one element set to zero, to avoid instability and sensitivity to model error, and partial decoupling of only the important outputs [2]. The gap is hard because the cost falls on the loop the decoupler was not built for. Design rules that choose which half to build, add dynamics behind a late gauge and price the cost to the other loop are thin.

3. *Structures that change with the active constraint.* The recycle plant floods once production passes its vapor wall, and the Tennessee Eastman plant on its A feed loss needed an override and a production cut to survive. At maximum production sixteen settings all tripped within a day. The reverse osmosis train of Chapter 2 held its pressure rating through an override that gave up recovery at the limit. The Open challenge box after this list takes the gap further.

4. *Controlled variables chosen by their downstream effect.* The soft sensor of Chapter 9 predicts about 0.78 % late variant whatever the crude (Table 9.7), so its downstream-aware rule does no better than a fixed window. A review of plantwide control warns that outputs chosen by a controllability analysis may be those easy to control rather than those important to control [1], and a guideline on drug substance manufacture asks that the linkages between process steps be evaluated so that the cumulative generation and removal of impurities is controlled [20]. The gap is hard because the variable that matters is measured only after the next unit has acted. Choosing such variables needs measurements that see the upstream composition and a validated model of the downstream unit.

5. *Model error in the weak directions.* The full inverse of Chapter 6 broke on one reversed corrector, and a response measured through the plant held 0.05 µm. A study of one fast orbit feedback system found that the weak modes, though dominated by measurement errors, still carry information, and that measuring the closed-loop gain needs an excitation larger than the beam's noise, which itself disturbed the gain [10]. The gap is hard because the directions that most need checking respond least to a test the plant can tolerate. A way to measure and validate the weak directions of a large response during operation, without disturbing it, is still needed.

6. *Chains with limits, spoilage and censored sales.* No ordering rule of Chapter 11 is string stable, and in a perishable chain with capacity limits the variance floor is a reference, not a bound. A review of fresh dairy quality finds that most models estimate shelf life from the product's initial condition and seldom state their uncertainty [21], and a study of grocery stores reports that an error in the stock record cannot be observed when it occurs and is found only by a deliberate count [22]. The gap is hard because the sales a rule sees are clipped by the stock it already lacks. Design rules for chains that pour out milk and clip orders are thinner than for linear chains.

> ⚠ **OPEN CHALLENGE** **Structures That Switch With the Active Constraint**
>
> Gap 3 runs through the recycle plant, the Tennessee Eastman plant and the reverse osmosis train, which held a limit only by changing what they controlled. A perspective on plantwide control notes that there is usually no effort during design to build a control system that can run at maximum throughput [6]. A review of simple control elements finds that switching between active constraints is usually done with selectors and split range control, and that the analysis of such switched systems is mathematically very difficult [2]. In the editor's view the gap is hard because each switch changes the structure whose stability was checked. Progress would be a routine test that proves a structure through every switch it may make.

12.5.1 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 [23]. 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 [19]. 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 [24]. 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 [25]. 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 [26]. 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 arrangement of a control system, chosen before any tuning, set what tuning could reach: a wrong pairing failed whatever the gains.
- A relative gain of one did not prove the loops independent. A triangular plant needed its off-diagonal gains read one by one.
- A decoupler was priced by its other loop, and where production was set decided where the upset went.
- A chain was judged at its far end, and every error in the record was caught by a check that could fail, never by rereading.

References

[1] T. Larsson and S. Skogestad, "Plantwide control: A review and a new design procedure," *Modeling, Identification and Control*, vol. 21, no. 4, pp. 209–240, 2000.

[2] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

[3] W. Tang, A. Allman, I. Mitrai, and P. Daoutidis, "Resolving large-scale control and optimization through network structure analysis and decomposition: A tutorial review," in *Proc. American Control Conference (ACC)*, 2023, pp. 3113–3129.

[4] S. Chen, Z. Wu, D. Rincon, and P. D. Christofides, "Machine learning-based distributed model predictive control of nonlinear processes," *AIChE Journal*, vol. 66, no. 11, p. e17013, 2020.

[5] S. Skogestad, "Control structure design for complete chemical plants," *Computers & Chemical Engineering*, vol. 28, no. 1-2, pp. 219–234, 2004.

[6] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[7] S. Yin, S. X. Ding, A. Haghani, H. Hao, and P. Zhang, "A comparison study of basic data-driven fault diagnosis and process monitoring methods on the benchmark Tennessee Eastman process," *Journal of Process Control*, vol. 22, no. 9, pp. 1567–1581, 2012.

[8] A. Bathelt, N. L. Ricker, and M. Jelali, "Revision of the Tennessee Eastman process model," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 309–314, 2015.

[9] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[10] S. Kongtawong, Y. Tian, L. Yu, X. Yang, G. Wang, K. Ha, and T. Shaftan, "Numerical simulation of NSLS-II fast orbit feedback system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 997, p. 165175, 2021.

[11] S. Gayadeen, S. R. Duncan, and M. T. Heron, "Fast orbit feedback control in mode space," in *Proceedings of ICALEPCS2013*, San Francisco, CA, USA, 2013, pp. 1082–1085, paper THMIB07.

[12] P. Kallakuri, D. Paskvan, J. Carwardine, N. Sereno, and N. Arnold, "Modeling the fast orbit feedback control system for APS upgrade," in *Proceedings of IBIC2017*. Grand Rapids, MI, USA: JACoW, 2018, pp. 196–198.

[13] L. Rolinger, M. Rüdt, and J. Hubbuch, "A critical review of recent trends, and a future perspective of optical spectroscopy as PAT in biopharmaceutical downstream processing," *Analytical and Bioanalytical Chemistry*, vol. 412, no. 9, pp. 2047–2064, 2020.

[14] L. Kong, A. Hasanbeigi, and L. Price, "Emerging energy-efficiency and greenhouse gas mitigation technologies for the pulp and paper industry," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep., Nov. 2012.

[15] D. Sedlak, M. Mauter, J. Macknick, J. Stokes-Draut, P. Fiske, D. Agarwal, T. Borch, R. Breckenridge, T. Cath, S. Chellam, A. Childress, D. Dionysiou, D. Giammar, E. Hoek, S. Jiang, L. Katz, J. Kim, R. Kostecki, J. McCutcheon, Y. Polsky, Z. Stoll, and P. Xu, "National alliance for water innovation (NAWI) master technology roadmap," U.S. Department of Energy, Report DOE/GO-102021-5617, 2021.

[16] B. Gopaluni, V. Kotamraju, and S. Bhushan, "Large language models in process systems engineering: Opportunities, architectures, and industrial deployment challenges," arXiv preprint arXiv:2606.11589, 2026.

[17] Y. Ye, X. He, D. Boshoff, W. Kuo, and L. Li, "little m: An AI agent for industrial process optimization," arXiv preprint arXiv:2609.16680, 2026.

[18] 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.

[19] 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.

[20] International Conference on Harmonisation, "Development and manufacture of drug substances (chemical entities and biotechnological/biological entities) Q11," ICH harmonised tripartite guideline, Step 4 version, https://database.ich.org, May 2012.

[21] L. Lima, A. I. Pereira, M. Pintado, and M. Carocho, "Fresh dairy quality in the era of smart food systems: Deterioration pathways, analytical technologies, and predictive modeling," *Foods*, vol. 15, no. 19, p. 3439, 2026.

[22] Y. Rekik, R. Oliva, A. A. Syntetos, and C. H. Glock, "Inventory record inaccuracy in grocery retailing: Impact of promotions and product perishability, and targeted effect of audits," *Journal of Business Logistics*, vol. 47, no. 3, p. e70079, 2026.

[23] 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.

[24] 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.

[25] 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.

[26] L. Bainbridge, “Ironies of automation,” *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

12.6 Exercises

12.1 Discussion · 8 min

An agent proposes a pairing for a new two-by-two plant and reports a relative gain of 1.01, “so the loops will not interact and need no decoupler”. Using Chapter 2 and Chapter 10, explain why that conclusion does not follow, and list what you would check before accepting the pairing.

12.2 Discussion · 8 min

On the paper machine of Chapter 3 an operator reports that the static decoupler “made the eight-gram grade move worse”, and the control engineer reports that it “cut the moisture excursion by a factor of eight”. Using Table 3.5, show that both are right. Then recommend a configuration and say what you would measure to settle the argument.

12.3 Discussion · 7 min

The owner of the reactor and column plant of Chapter 7 wants 15 % more product and proposes to raise the fresh feed on the shipped structure, CS1. Using the chapter’s results, explain what will happen. Then propose where production should be set and the check to make before any change.

12.4 Discussion · 8 min

An agent retuning the orbit feedback of Chapter 6 proposes the full inverse because “it gives the smallest orbit, 0.09 against 0.21 µm”. Evaluate the proposal with the chapter’s results, and say what you would ask the agent to run before you accept any inverse.

12.5 Discussion · 10 min

Section 12.4.1 lists seven errors made by the agents that built this volume’s apps, papers and Chapter 1. For each, name the check that caught it. Then write a procedure, in five to eight steps, for an agent that will propose a change of control structure on a running plant (a new pairing, a decoupler or a moved throughput handle), stating what the agent may read, what it may write and what a person must approve.

12.6 Discussion · 7 min

A site standard asks for every level loop to be tuned tight “so that inventories stay on target”. Using Chapter 4, Chapter 11 and Section 1.4.8, argue when the standard is wrong, and what a supply chain adds to the argument.

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) 2 and 8 m^3/h, gains $1, 1, 8, -2$. (b) $\lambda = 0.2$. Pair the flow with the water and the concentration with the concentrate, $\lambda = 0.8$.

1.3 $\lambda_{11} = 2$. Diagonal pairing NI $= 0.5$, off-diagonal pairing NI $= -1$.

1.5 (c). It cancels the cross effect at steady state only.

1.7 Two: the feed rate and the steam rate.

1.9 $\lambda_{11} = 1$. $\sigma_1 = 1.618$, $\sigma_2 = 0.618$, condition number 2.618.

Chapter 2. Reverse Osmosis Loop Pairing and Decoupling

2.1 (d). Salt enters with the feed and leaves with the concentrate and the permeate, and the difference accumulates in the channel volume.

2.3 (b). The coupling from the pump to the recovery is the gain K_{21} of the lumped model, which scenario 3 measures.

2.5 (a) $D = 0.914$, residual -0.027 %/%. (b) 3.50 % and 2.46 points, against 2.30 measured. (c) 3.73 %, against 3.44 % measured.

2.7 (a) 45.94 % and 53.43 %. (b) 44.84 and 45.00 %, a peak deviation of 0.16 point. (c) 44.14 %, a deviation of 0.86 point below the setpoint.

2.9 (a) 69.81 bar, no time over the rating, 2415 m^3, 2.757 kWh/m^3. (b) 74.00 bar, 108.6 min over, 2415 m^3, 2.675 kWh/m^3. The MPC lowers its recovery target at the limit, for 3.0 % more energy.

2.11 (a) 70.24 g/L and 65.60 g/L. (b) 67.6 g/L, 0.962 and 1.030, against 0.958 and 1.025 in the table. (c) About 48.0 %.

Chapter 3. Paper Machine Control With a Scanning Gauge

3.1 (d). The posted value is the average of a traverse, centered half a period back, and it is then held for one more period, whose mean age is half a period.

3.3 (b). A negative relative gain means the steady gain of one loop changes sign between the other loop being in Manual and in Automatic.

3.5 (a) $\Lambda_{11} = 1.52$, $\Lambda_{12} = -0.52$. (b) 0.66 for the natural pairing and -1.93 for the swapped one. (c) $D_{12} = 0.0233$ and $D_{21} = 14.6$ (14.64 from the unrounded gains). The residual coupling is $+0.009$ points per percent at 14.5 and -0.887 at 29.25.

3.7 (a) At 140 s, to 80.14 g/m^2. (b) About 36.5 s and 60 s, a gap of about 24 s. (c) 81.73 g/m^2 against a sheet of 82.66 g/m^2.

3.9 (a) 7.30 %, a 1.80 points excursion, and 2.38 t. (b) 5.73 %, 0.23 points, and 4.54 t. (c) 5.74 %, 0.24 points, and 3.14 t. The stock-to-steam half, D_{21}, carries the benefit. The weight loop pays for it.

3.11 (a) 60 kPa against 8 kPa. (b) 1.33, below the value of two at which the explicit back-calculation step diverges. (c) The production setpoint is parked on the measurement while the web is down, and the loop recloses only at 95 % of a threaded web.

Chapter 4. Three-Effect Evaporator Concentration Control

4.1 (c). The entrained flow $\dot{m}_e$ is part of V_1 in Equation (4.5), and it grows with the motive flow.

4.3 (a). In Equation (4.11) T_1 contains ΔT_2 and ΔT_3 but not ΔT_1, which lifts the steam chest instead.

4.5 (a) 1.76 K. (b) 65.32 °C, a margin of 4.18 K. (c) $\theta_3 = 45.88\,^{\circ}\mathrm{C}$, a peak of about 69.30 °C, so the ceiling would not bind. The run peaks at 68.27 °C.

4.7 (a) 20, 75, 70 and 400 s. (b) 51.3 %. (c) 20, 75, 80 and 390 s. AT-301 waits for effect 3's holdup to concentrate and then for 85 s of product line.

4.9 (a) 1.49 %, 155.5 t, 67.05 °C, 4.49 kg/kg, 13 s. It fails on solids, tonnage and dry film. (b) 0.62 %, 160.0 t, 69.35 °C, 4.40 kg/kg, 0 s. It passes.

4.11 (a) 5.04 t/h, a margin of 1.48. (b) About 4.68 t/h, 0.93 of the floor. Without recirculation the film breaks and effect 3 keeps 0.961 of its film coefficient, against 0.988.

Chapter 5. Roll-to-Roll Web Tension and Coat Weight

5.1 (b). The 19 m between the die and the gauge is coated light before any trim can act, so the tuning moves the scrap by only 0.2 m.

5.3 (d). The plan scraps 40.40 m without the pump and hose lags in its model and 71.00 m with the delay counted in time, so its model carries the gain.

5.5 (a) 18.11 kg m^2 and 201 kg, against 368 kg. (b) 59.87 N m. (c) 101 mm.

5.7 (a) 8.02 g/m^2, close to $12 \times 120/180 = 8.0\,\mathrm{g/m^2}$, and 208.90 m. (b) 22.30 m, 16.6 %. (c) 12.80 m, 1.6 %. The rest is the pump's 0.8 s lag during the ramp.

5.9 (a) 21.05 m and 63.00 m. (b) 19.65 m and 24.30 m. (c) The plan's model of the pump lags and of the delay in meters. Loops with a 25 s trim and a lead scrap 20.55 m.

5.11 (a) +300, −300 and −100 N s/m. (b) 48 m/min and 28.8 m/min. Keep the diameter sensor.

Chapter 6. Storage Ring Orbit Feedback

6.1 (a). Each mode passes the noise with a gain l_m/σ_m, and the weak group has σ from 1.28 down to 0.94 µm/µrad.

6.3 (c). The reversed column makes $M = I - 2V^{\mathsf{T}} e_{38} e_{38}^{\mathsf{T}} V$, with one eigenvalue $e = -1$.

6.5 (a) 1.00, 0.88, 0.14, 0.08. (b) At $\sqrt{\mu} = 3.16\,\mu\mathrm{m/\mu rad}$, between σ_{48} and σ_{49}, so the filter falls where the cut would. (c) 2.01 at 1 Hz and 4.92 at 10 Hz.

6.7 (a) 0.198 µrad measured, 0.199 µrad predicted, orbit 0.09 µm. (b) 0.016 µrad (the chapter's 0.0164), orbit 0.21 µm. (c) About 12 times less jitter for a little over twice the orbit residual.

6.9 (a) 2.69 µm worst, 2.63 µm at the end. (b) 0.31 µm and 0.25 µm, about 8.7 times less. (c) About a tenth of the 2.6 µm bump, the gap table's error for device 2, which the BPMs cannot see either.

6.11 The tolerances are teaching values, while a real light source derives its own from how its experiments sample, integrate and scan. The model is horizontal only, with thin lenses, no coupling and no vertical plane. The loops run at 1 kHz and 100 Hz, below real fast feedback, with one lag and no rate limit per corrector, and one of seven records fails.

Chapter 7. Snowball Effect and Control Structures

7.1 (a). At a fixed holdup the reaction rate $kM_R z$ rises only with z, and the recycle $F_0(z - x_B)/(x_D - z)$ grows as z approaches x_D.

7.3 (c). The loop's manipulated variable is saturated, and no tuning creates vapor above the flooding rate.

7.5 (a) 4.61 and −3.61. (b) 11.5. (c) $\mathrm{d}D^*/\mathrm{d}x_D = -F_0(z - x_B)/(x_D - z)^2$, about 18.8 kmol/h more recycle per mol % lost.

7.7 (a) +62.3 %, 3.1 % to flood, 811.9 kmol/h against the curve's 817.5 kmol/h. (b) 4.61 rises to 5.84 % per percent. (c) +93.5 % with the recycle at 91.83 mol % A.

7.9 (a) Production passes at 238.6 kmol/h, but the product falls to 79.57 mol % B and the boilup sits at the steam limit for 434.9 min. (b) 234.9 kmol/h fails and 236.0 kmol/h passes. (c) 70.1 %.

7.11 (a) $2^{0.6} = 1.52$. (b) The column and its reboiler, condenser and recycle pumps, which are sized for a recycle and boilup less than half as large. The recycle then moves less for a given production change or catalyst loss.

Chapter 8. Tennessee Eastman Challenge Process

8.1 (d). The trip time follows the reading of the constants, and the integration step changes it by less than 0.001 h.

8.3 (b). A buffer held tight passes the disturbance on to the flow it sets.

8.5 (a) 0.0138, 72 steps. (b) 0.025 kPa. (c) 0.116 kPa, a tenth of the largest excess.

8.7 (a) 3.59 h, high reactor pressure. 2705.0 and 2689.6 kPa g. (b) 4.85 h, 5.84 h, 37.0 %. (c) 2751 kPa g, 45.4 %. Inventory loops make the plant stable, and only the composition loops above them ride a feed change.

8.9 (a) 6.68 h, high reactor pressure. (b) 215.7 $/h, target missed on cost. (c) Target met at 187.3 $/h, with the pressure peaking at 2944 kPa g near 5.53 h.

8.11 (a) 0.25 to 0.5 h, mean 0.375 h, and 0.1 to 0.2 h, mean 0.15 h. (b) 20.2 h and 4.2 h, with 1.04 and 5 mol % of setpoint travel inside the oldest reading.

Chapter 9. GLP-1 Peptide Cuts Judged at the Filter

9.1 (c). The late variant in the pool rises from 0.358 to 0.939 %, the particles shrink from 11.6 to 9.0 µm, and the cake needs 79.3 min of the 60 allowed.

9.3 (a). $Q_f \propto \Delta P^{1-s}$, and $(3.5/1.5)^{0.4} = 1.40$.

9.5 (a) 3.92 % and 33.0 %. (b) 33.0 g and 18.1 g. (c) 2.5 diavolumes, with about 0.0077 % acetonitrile and 31.4 g.

9.7 (a) 48 s, pool 97.84 %, solid 98.26 %, rejected. (b) 98.40 %, 98.81 %, 55.4 g. (c) 96.90 %, 97.40 %, rejected. The soft sensor carries the detector signal forward by the declared 1.2 L, not the 2.4 L the line holds.

9.9 (a) The cake does not leave the filter in time, and two diavolumes leave 0.0273 % acetonitrile. (b) 59.9 g, 10.0 g/h, met. (c) Pool 97.14 %, solid 97.70 %, rejected.

9.11 (a) The filtration starts 226.3 min after the transfer, so the result arrives about 211 min ahead. (b) The gap of 0.097 points is 3.2 standard deviations. The assay cannot correct the cut, but it can set the filter pressure as a feedforward.

Chapter 10. Space Station Life Support

10.1 (b). The total is the sum of the partial pressures, and the pressure loop is the only one that reads it.

10.3 (d). Eight people exercising give 18.0 mol/h, and a fresh cool bed at 0.5 kPa takes at most 10.29 mol/h.

10.5 (a) 2.456 kWh; 61.6 %. (b) 1.176 kW; 60.3 %, 52.6 %. (c) 10.10 mol against 9.58 mol, so the oxygen can wait for the sun and the heater can run in the dark.

10.7 (a) 1.46, 4.83 and 40.75 h. (b) 1.47, 4.85 and 39.20 h (the crossings fall at 1.46, 4.84 and 39.19 h). The leak takes about 0.24 mol/h of oxygen. (c) 102.13 kPa, a rise of 0.83 kPa.

10.9 (a) Four failures: carbon dioxide over its limit for 494 min, the battery below its floor for 901 min, 520 mol of nitrogen left and the leak not flagged. (b) Met, with 664 mol left and the leak flagged after 3.0 h. (c) The balance is necessary and the feedforward is not.

10.11 (a) 28.7 mol and 0.83 kg a day, about nine times the module figure. (b) About 14.6 days, each excursion taking about 41 hours' worth of it; about 67 h with the step open.

Chapter 11. Bullwhip Effect in a Perishable Dairy Chain

11.1 (b). With $\beta = 1$ the supply-line term subtracts the milk on order as fast as the stock term adds the gap, so the gap is ordered once.

11.3 (d). Equations (11.9) and (11.10): $P_h = \min(K_d, 29) = 20$ and $W_h = 29 - 20 = 9$.

11.5 (a) At $z = 1$, $Q = 0$ and $\Gamma = a_S/a_S = 1$: a lasting change passes upstream whole. (b) Denominator one, $\Gamma(-1) = -13/7$, $|\Gamma| = 1.857$, period 2 days. (c) $H_F = 0$, so $|\Gamma| = 1$.

11.7 (a) Day 16; day 17. (b) 96 h; service 0.9365, 80 kL lost. (c) 24 h; service 0.9948, 6.5 kL lost.

11.9 (a) \$898.2k, 44 h, not met, 3.72 %. (b) \$791.3k, 6 h, met, 1.28 %. (c) \$774.4k, 0 h, met, 0.47 %, 182 solves without a failure.

11.11 (a) $468.1/0.25 = 1872.4$ and $511.2/0.25 = 2044.8$ kL, a difference of 172.4 kL. (b) The shipped chain bottles its overshoot and pours 3.72 % of bottling away, and drains raw milk in the troughs. Each liter costs \$0.60 instead of \$0.25.

Chapter 12. Summary and Outlook

12.1 A good answer notes that the relative gain is close to one whenever either off-diagonal gain is close to zero, however large the other is. The reverse osmosis train has $\lambda_{11} = 1.00005$ while the pump moves the recovery 0.703 %/%, and a production step still moved the recovery 0.937 points. The cabin's oxygen and total pressure loops have a relative gain of one, yet 1 kPa more oxygen shut the nitrogen valve 2.52 h later. The checks are both off-diagonal gains read one by one, gains computed by a central difference at the operating point and across the operating range, a step of each loop's setpoint with the other loop in automatic, and a one-term feedforward if the coupling runs one way.

12.3 A good answer notes that the column floods 10.7 % above design production, so 15 % more fresh feed on CS1 reaches the flooding rate at 3.96 h, sends A into the product (11.4 kmol/h at 12 h) and leaves the product at 95.04 mol % B. The snowball elasticity of 4.61 explains why, since 10 % more feed already raised the recycle 62.3 % in 12 h. It proposes CS2, production set by the reactor holdup, which moved the recycle down while production rose, at 14 % more holdup for each 10 % of production. The check is the vapor margin at the new production, computed before the change.

12.5 A good answer pairs the relative gain from the wrong derivative with recomputing each number from the model at the operating point, the pairing switch with changing the structure in automatic in a run, the two handles with a comparison at equal loop gain, the production loop that read the truth with a probe that changes a hidden value while the measurements stay fixed, the UV rule with varying the threshold and running many seeds, the claim about every structure with a reading against the case table, and the clamping offset with a disturbance run. The procedure gives read-only access by default, computes the structural measures from a validated model, simulates the proposal on disturbances at equal loop gain, plans the initialization of every loop for a bumpless switch, has a person approve the change under management of change with a way back, commissions one loop at a time and logs every step.

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.

active constraint A limit, such as a purity or temperature bound, at which the best operation holds a variable, which uses up one degree of freedom. *Chapter 1*.

antisolvent precipitation Making a dissolved product come out of solution as solid particles by adding a liquid in which it does not dissolve. *Chapter 9*.

averaging level control Tuning a level loop loosely so that the vessel absorbs changes of flow instead of passing them straight to the next unit. *Chapter 1*.

back-calculation An anti-windup method that resets the integral of a controller to the output its final element actually delivered. *Chapter 7*.

basis weight The mass of a sheet of paper per square meter. *Chapter 3*.

beam position monitor (BPM) A sensor on the vacuum chamber of a storage ring that reports where the electron beam passes. *Chapter 6*.

beer distribution game A classroom game in which retailer, distributor and factory players order goods through delays, and their orders swing far more than demand. *Chapter 11*.

beta gauge A gauge that reads the basis weight of a sheet from how strongly the sheet absorbs beta radiation. *Chapter 3*.

bottleneck The unit whose limit is reached first as production rises, and which therefore caps the production of the whole plant. *Chapter 7*.

bullwhip effect The growth of the variation of orders from stage to stage up a supply chain, from the shop toward the factory. *Chapter 1*.

bumpless transfer Changing the mode or the pairing of a controller without a jump of the actuator, by starting the controller from its present output. *Chapter 2*.

cake filtration Filtration in which the collected solids form a cake that the liquid must drain through, so the size of the particles sets the time. *Chapter 9*.

carbon dioxide removal assembly Two beds that take turns: one adsorbs carbon dioxide from the cabin air while the other is heated to release it. *Chapter 10*.

cascade control An arrangement in which one controller sets the setpoint of another, as a tank level sets a production flow setpoint. *Chapter 2*.

closed inventory An inventory with no outlet that a controller can open, so an excess leaves only by use or by a leak. *Chapter 10*.

coat weight The mass of dry coating per square meter of web. *Chapter 5*.

concentration polarization The rise of the salt concentration at the membrane surface above that of the bulk feed, which raises the pressure the membrane needs. *Chapter 2*.

condition number The ratio of the largest to the smallest singular value of a gain matrix. A large value means that some directions are weak. *Chapter 1*.

consistent inventory control A structure in which every inventory, and every group of inventories, is held by some loop. *Chapter 1*.

control structure The choice of what to measure, what to manipulate and which loops link them, made before any loop is tuned. *Chapter 1*.

corrector magnet A small steering magnet that bends the beam of a storage ring slightly to correct its orbit. *Chapter 6*.

CS1, CS2 The two structures of the recycle example: the fresh feed sets production, or the effluent is fixed and the reactor level setpoint sets production. *Chapter 1*.

dancer roll A roll on a loaded arm in the web path, whose position reports a mismatch of speeds while its load sets the tension. *Chapter 5*.

dead band A range of command change within which an element does not move, as a sticking valve holds until the command has moved enough. *Chapter 8*.

decentralized control Control by many single loops, each reading one variable and moving one actuator. *Chapter 8*.

decision calendar The times at which an ordering rule is allowed to act, such as each midnight or each Monday. *Chapter 11*.

decoupler A correction added to each actuator from the output of the other controller, so that each controller sees only its own loop. *Chapter 1*.

degrees of freedom The number of handles, such as valves, pump speeds or heater duties, that can be set independently. *Chapter 1*.

diafiltration Washing a product on a membrane by adding fresh buffer while liquid passes through the

membrane. *Chapter 9*.

distributed control system (DCS) A supervisory level of control over a group of localized controllers that together run the loops of a plant. *Chapter 1*.

disturbance propagation The way a control structure passes an upset from the unit it enters to other units and variables. *Chapter 1*.

dryer section The steam-heated cylinders of a paper machine over which the sheet passes to dry. *Chapter 3*.

effluent The stream that leaves the reactor for the separator. *Chapter 1*.

energy recovery Returning the pressure energy of the concentrate to the feed, as a pressure exchanger does. *Chapter 2*.

equal-percentage trim A valve plug shaped so that the flow capacity changes by the same fraction for each equal step of travel. *Chapter 2*.

external reset Feeding the integral of a controller from the value actually applied, so that a controller that loses a selector cannot wind up. *Chapter 2*.

feedforward Moving an actuator from a measured disturbance or a command before an error appears. *Chapter 5*.

gain matrix The table of steady-state gains from each input of a plant to each of its outputs. *Chapter 1*.

GLP-1 receptor agonist A class of peptide medicines that mimic a gut hormone which stimulates the release of insulin. *Chapter 9*.

IAE Integral of the absolute error: the area between a variable and its setpoint over a run. *Chapter 1*.

ill-conditioned plant A plant whose gain matrix has a weakest direction much weaker than its strongest, so inverting it amplifies noise. *Chapter 6*.

inert A component that neither reacts nor leaves with a product, so it accumulates until a purge removes it. *Chapter 1*.

interaction The effect of the actuator of one loop on the variable that another loop controls. *Chapter 1*.

internal model control (IMC) A controller built around a model of the plant. The storage ring of this volume runs one for each singular mode. *Chapter 6*.

inventory control Holding the material in each vessel, usually by its level or pressure, so that the vessel neither fills nor empties. *Chapter 1*.

leak detection Finding a slow loss of gas from the balance of what the pressure loop must add to the cabin. *Chapter 10*.

line speed The speed at which the web travels through a coating line, set by the speed master. *Chapter 5*.

load cell A sensor under a roll that reads the tension of the web as a force. *Chapter 5*.

manipulated variable (MV) The actuator signal that a controller moves, such as a valve opening or a pump speed. *Chapter 1*.

mass spectrometer An analyzer that reports the partial pressures of the gases of a cabin by sorting molecules by their mass. *Chapter 10*.

membrane fouling Deposits on a membrane that lower its permeability over time. *Chapter 2*.

model predictive control (MPC) A controller that predicts the plant with a model and chooses its moves by solving an optimization at each step. *Chapter 1*.

multiple-effect evaporator A train of boiling bodies in which the vapor from each heats the next, so the steam is used several times. *Chapter 4*.

Niederlinski index The determinant of the gain matrix over the product of the paired gains. A negative value means instability for every tuning. *Chapter 1*.

noise amplification The growth of measurement noise into actuator moves when a controller corrects weak directions. *Chapter 6*.

order-up-to policy An ordering rule that closes the whole gap between a target and the stock plus the goods already on order. *Chapter 11*.

override control A selector that lets a protective controller take over an actuator when a limit comes near. *Chapter 2*.

pairing The choice of which actuator holds which controlled variable. *Chapter 1*.

partial pressure The share of the total pressure of a gas mixture that one gas exerts. *Chapter 10*.

perishable inventory Stock that ages and must be thrown away once it passes its date. *Chapter 11*.

permeate The fresh water that passes through a reverse osmosis membrane. *Chapter 2*.

PI controller A controller with proportional and integral action, the loop used throughout this volume. *Chapter 1*.

plantwide control The control philosophy of a whole plant, with emphasis on its structural decisions rather than the tuning of each loop. *Chapter 1*.

preparative chromatography Separating a product from impurities in a packed column at production scale and keeping the part of the outlet that is pure enough. *Chapter 9*.

production index A single setting that moves the feed flows of a plant in ratio, so one number sets production. *Chapter 8*.

purge A small stream drawn off to remove an inert or impurity that would otherwise accumulate.

Chapter 1.

ratio control Holding one flow or speed in a set proportion to another. *Chapter 8*.

recovery The fraction of the feed of a reverse osmosis train that leaves as permeate. *Chapter 2*.

recycle A stream returned from a later unit of a plant to an earlier one. *Chapter 1*.

relative gain The gain of a loop with the other loops open divided by its gain with them closed perfectly. *Chapter 1*.

relative gain array (RGA) The table of the relative gains of every possible pairing. Each of its rows and columns sums to one. *Chapter 1*.

response matrix The table of orbit shifts at every monitor that a unit kick of each corrector produces. *Chapter 6*.

reverse osmosis Pushing water through a membrane against its osmotic pressure so that most of the salt stays behind. *Chapter 2*.

sampled analyzer An analyzer that reports at fixed intervals, each result describing a sample drawn earlier. *Chapter 8*.

scanning gauge A sensor head that travels across the moving sheet and reports one average for each crossing. *Chapter 3*.

self-optimizing control Choosing controlled variables whose constant setpoints keep a plant near its optimum as disturbances change. *Chapter 1*.

SIMC rule A simple rule that sets a PI controller from a first-order model with dead time and a chosen closed-loop time. *Chapter 1*.

singular value decomposition (SVD) Splitting a gain matrix into input directions, output directions and the gain along each pair of them. *Chapter 1*.

snowball effect The growth of a recycle much faster than the feed when production rises through a reactor of fixed size. *Chapter 1*.

steam economy The mass of water an evaporator boils off for each kilogram of steam it receives. *Chapter 4*.

storage ring A ring of magnets that keeps electrons on a closed orbit so that they give off light for experiments. *Chapter 6*.

string stability The property of a chain in which a disturbance does not grow as it passes from stage to stage. *Chapter 1*.

supply line Goods that a stage has ordered but not yet received. *Chapter 11*.

tangential-flow filtration Filtration in which the liquid flows along the membrane, sweeping its surface while part of the liquid passes through. *Chapter 9*.

Tennessee Eastman process A published benchmark model of a whole chemical plant with a reactor, condenser, separator, stripper and recycle compressor. *Chapter 8*.

thermal vapor recompression A steam jet that draws vapor from an effect and compresses it with motive steam to heat the first effect again. *Chapter 4*.

throughput manipulator (TPM) The flow, speed or setpoint that sets the production rate of a plant. *Chapter 1*.

Tikhonov regularization Adding a penalty on the size of the moves when inverting a matrix, so that weak directions are corrected less. *Chapter 6*.

transmembrane pressure The pressure difference that drives liquid through a membrane. *Chapter 9*.

valve stiction Friction that makes a valve stem stick and then jump, so a loop through it can cycle. *Chapter 8*.

variance ratio The variance of the daily orders of a stage divided by the variance of daily shopper demand. *Chapter 11*.

velocity-form PID A PID controller that computes the change of its output at each step and adds it to the previous output. *Chapter 1*.

web A continuous sheet of film, foil or paper that runs through a processing line. *Chapter 5*.

web tension The pulling force in the web between two rolls. *Chapter 5*.

winder A drive that unwinds or rewinds a roll of web whose radius changes as it turns. *Chapter 5*.

Index of Control Topics

active constraint, 164
AI agents
  control structure design, 284
  engineering practice, 284
  open problems, 288
  reproduction, 287
  truth audit, 286
  verification record, 285
  while studying, 17
analyzer dead time, 40, 190, 200
anchoring and adjustment, 255
  with supply line, 261
averaging level control, 10, 93, 196

back-calculation, 61, 193
benchmark plants, 283
bottleneck, 164
branch and bound, 239
bullwhip effect, 254, 283
  variance floor, 264
bumpless transfer, 15, 30, 61, 171, 262
  pairing change, 286

cascade control, 30, 62
closed inventory, 229
code verification, 216
condition number, 9, 145
control performance assessment, 52
control structure design
  procedure, 3, 281
control structure selection, 171
controlled variable
  choice of, 208
CUSUM test, 240
cyclic scheduling, 238, 247

data export, 16
data reconciliation, 239
dead band, 238
dead time
  effective, 50
  of a scanning gauge, 59
decentralized control, 112, 193
decision calendar, 262, 272
decoupling, 7, 22, 49
  actuator limits, 286
  cost, 279
  one-way static, 30, 61, 66
  simplified static, 60
degrees of freedom, 4
delayed analyzer, 196
distributed control, 177, 200, 222, 246, 282
disturbance model, 120
disturbance propagation, 10
  along a chain, 280

economic optimum
  near a constraint, 35, 197
economics of control
  moisture giveaway, 75
equilibrium check, 270
event-driven control, 61
external reset, 30

fault detection, 186
feedforward, 109, 193
  from a gap table, 147
feedforward control, 92, 238
finite-volume method, 210
floating-point precision, 191
frequency response
  of an ordering rule, 264

gain margin, 122
gain matrix, 60
gain scheduling, 123
gain sign reversal, 29
grade change, 69, 75, 196

ill-conditioned plant, 138
IMC tuning, 65
inert accumulation, 184
inertia compensation, 118
infeasible demand, 241
information sharing, 264, 273
integrator windup
  after a lost measurement, 62
internal model control, 141
  per singular mode, 141, 146
inventory control, 5, 162, 193
  consistency, 164

Kalman filter, 16, 64, 120

limit cycle, 197
linearization, 122
Luyben's rule, 163

measurement noise
  passed into the product, 64
mid-ranging, 147
minimum-variance control, 51, 52
mixed-integer linear programming, 239
model error
  wrong sign, 148
model predictive control, 93, 103, 119, 133, 194
  ablation, 72
  against decentralized control, 281
  centralized, 262
  in the apps, 15
  linearization error, 74
  move blocking, 63, 119, 194, 263
  offset-free, 31, 64
  soft constraint, 120
  soft constraints, 32, 64, 94, 263
  target calculation, 32
moving horizon estimation, 16, 75
multivariable control, 1

noise amplification, 139

orbit feedback, 138, 154
order-up-to policy, 262
override control, 92, 126, 194, 269, 288
  integrating, 194
  low selector, 30
  product quality, 35

pairing, 6, 22, 49, 60, 81
  reversed, 33, 67
Parseval's theorem, 264
perishable inventory, 257
PID control
  anti-windup, 15, 61, 119
  velocity form, 15, 61, 117, 172, 193, 215, 238
PID tuning
  SIMC rule, 123
plantwide control, 1, 81, 161, 183, 207, 283
  heuristic procedure, 3, 185, 283
  open problems, 43, 287
  practice, 2, 177, 201, 223, 247, 282
population balance, 213
process gain
  drift, 35

radius compensation, 118
ratio control, 92, 119, 193
recycle, 5, 162
regulatory control layer, 12, 68
relative gain array, 6, 22, 60
  computation, 286
  limits, 278
  negative, 60
  triangular plant, 29, 237
response matrix, 140, 143

sampled analyzer, 172, 190, 215
sampled measurement, 49
  boxcar average, 56
sampled-data control, 122
sampling period
  and stability, 65
scenarios, 16
  completion rules, 16
scheduling
  integrated with control, 230
scorecard, 16
self-optimizing control, 164, 185
sensor fault detection, 147
servomechanism theory, 255
SIMC tuning, 30
simulation engine, 14
simulator framework, 14
singular value decomposition, 9, 145
  truncated, 146
  weak directions, 280
snowball effect, 5, 161
  gain definitions, 170
soft sensor, 75, 215, 224, 283
split-range control, 215
stability boundary, 122
string stability, 254
supply chain control, 253
supply-line feedback, 261

throughput manipulator, 5, 92, 162, 193, 279
Tikhonov regularization, 146, 156, 283
transport delay, 56, 89
  variable, 115

unit-operation coupling, 207
unstable steady state, 184

variance ratio, 259

winder control, 109

zero-order hold, 56, 122

Index of Applications

adsorption cycle, 235
airlock, 234
antiscalant, 28
antisolvent precipitation, 213
Arrhenius kinetics, 166

basis weight
    conditioned, 56, 74
battery
    eclipse budget, 235
battery electrode, 112, 132
beam-position monitor, 141, 154
beer distribution game, 256
benchmark, 184, 202
beta gauge, 50, 72
boiling-point elevation, 84, 86
breakthrough, 235

cake filtration, 213
carbon dioxide removal assembly, 231, 246
catalyst activity, 166
cleaning in place, 87
closed orbit, 143, 155
coat-weight gauge, 115
coating bead, 115
coating interlock, 119
coating line, 109
code date, 258, 272
code verification, 191
cold-chain depot, 257, 272
concentrate valve, 23
concentration polarization, 27
condensing heat exchanger, 235
constant molar overflow, 167
corrector magnet, 139, 154
cross-direction control, 52

dairy plant, 257
dairy processing, 81
dairy supply chain, 253
dancer roll, 111
dechlorination, 36
densitometer, 91, 101
desalination, 22, 24
diafiltration, 213
distillation column, 167
dryer cylinder, 55, 73
dryer section, 52

energy recovery, 24

falling-film evaporator, 82
flooding, 170
FODO lattice, 141
fouling, 87
fraction pooling, 215

gas chromatograph, 12, 172
gear pump, 115
GLP-1 receptor agonist, 207, 209
gradient elution, 209, 222

high-pressure pump, 23, 41

insertion device, 141
interlock, 190

Langmuir isotherm, 212
leak detection, 239
life support system, 229
line-out, 259
lithium hydroxide canister, 231
load cell, 111, 131

machine-direction control, 52
membrane fouling, 24, 42
membrane scaling, 24, 42
milk, 253
milk powder dryer, 258
moisture
    wet and dry basis, 54
moisture sensor, infrared, 50
molecular sieve, 231
multiple-effect evaporator, 81, 103

Nutsche filter, 213

operating cost, 190
orifice plate, 176
osmotic pressure, 27
oxidant damage, 36

paper machine, 49
partial pressure, 234, 245
peptide manufacturing, 207
permeate tank, 28

pH conditioning, 212
pre-metered coating, 110
preparative chromatography, 209, 221
press dewatering, 75
press felt, 68
press section, 54
pressure exchanger, 24
production rate, 62
promotion, 266
purge, 184

reactor and separator with recycle, 3, 161, 164, 178, 184
reboiler, 167, 176
recycle compressor, 186
recycle line, 164
reflux drum, 167
relative volatility, 167
reverse osmosis, 22
rewind, 112
roll-to-roll coating, 109

Sabatier reactor, 234, 247
saturation index, 28
scanning gauge, 49
scrap, 115
semaglutide, 209
sheet break, 62
shelf life, 258, 261
slack web, 114
slot die, 110
solid-phase peptide synthesis, 209
solution-diffusion model, 25
spacecraft cabin, 229
specific energy, 27
speed master, 118
spillback valve, 186
splice, 116
spot milk, 258
spray dryer, 82
spring flush, 268
steam economy, 82, 102
  marginal, 90
steam header, 55, 74
stirred tank reactor, 164
stock consistency, 53
storage ring, 138
stripper, 186
strong focusing, 140
supermarket, 257, 271
synchrotron light source, 138

tangential-flow filtration, 213
Tennessee Eastman process, 183, 185
thermal vapor recompression, 82, 101
thick stock, 53

unwind, 112

valve stiction, 197, 200

water electrolysis, 234
web, 110
web break, 116
web map, 115
web tension, 110, 131
  transport equation, 114
wetting rate, 87
winder, 112, 131
wrinkle, 115

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.